




Universidade do Minho

Escola de Engenharia

Mafalda Francisco Ramôa da Costa Alves

**Ansätze for Noisy Variational Quantum
Eigensolvers**



Universidade do Minho

Escola de Engenharia

Mafalda Francisco Ramôa da Costa Alves

**Ansätze for Noisy Variational Quantum
Eigensolvers**

Master Thesis
Engineering Physics

Work developed under the supervision of:

Ernesto Galvão
Mikhail Vasilevskiy


COPYRIGHT AND TERMS OF USE OF THIS WORK BY A THIRD PARTY

This is academic work that can be used by third parties as long as internationally accepted rules and good practices regarding copyright and related rights are respected.

Accordingly, this work may be used under the license provided below.

If the user needs permission to make use of the work under conditions not provided for in the indicated licensing, they should contact the author through the RepositoriUM of Universidade do Minho.

License granted to the users of this work



**Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International
CC BY-NC-SA 4.0**

<https://creativecommons.org/licenses/by-nc-sa/4.0/deed.en>

Acknowledgements

I would like to offer my special thanks to Raffaele Santagati for suggesting this project, and for his guidance and support throughout the course of it. I thank also my co-supervisor Ernesto Galvão, for his assistance at every stage of the project, as well as for his invaluable advice; and Mikhail Vasilevskiy, for his support.

I thank my family and friends for their encouragement. In particular, I'm deeply thankful to my sister, colleague and friend Alexandra. This year would not have been the same without her continuous support both inside and outside the academic sphere.

I want to express my gratitude also to all the teachers and professors that, throughout the years, have taught me so many different things in so many different ways. The genuine concern for students and dedication for teaching I have encountered throughout my life have certainly impacted it along the years.

Additionally, I wish to thank the Programme New Talents in Quantum Technologies of the Gulbenkian Foundation (Portugal) for the financial support, and for the opportunity of being involved in this experience. The network of students under the support of the Gulbenkian Foundation is unmatched in the enthusiasm for learning and growing.

Finally, I wish to acknowledge the use of IBMQ's systems [42], Google Colaboratory's cloud service [30], and the NOVAThesis \LaTeX template [56].

STATEMENT OF INTEGRITY

I hereby declare having conducted this academic work with integrity. I confirm that I have not used plagiarism or any form of undue use of information or falsification of results along the process leading to its elaboration.

I further declare that I have fully acknowledged the Code of Ethical Conduct of the Universidade do Minho.

Ansätze for Noisy Variational Quantum Eigensolvers

Simulating quantum mechanical systems is one of the main applications envisioned for quantum computers. In contrast with the first algorithms created for this purpose, that were devised to be implemented in a fault-tolerant quantum computer, the [Variational Quantum Eigensolver \(VQE\)](#) aims to adjust to the constraints of [Noisy Intermediate-Scale Quantum \(NISQ\)](#) devices. There is hope that the class of [Variational Quantum Algorithms \(VQAs\)](#), to which [VQE](#) belongs, will be the first to achieve quantum advantage.

The choice of ansatz can dictate the success (or lack thereof) of a [VQA](#): too deep ansätze can hinder near-term viability, or lead to trainability issues that render the algorithm inefficient. In this context, this dissertation aimed to analyse different ansätze for quantum chemistry, examining their noise-resilience and viability in state-of-the-art quantum computers. In particular, dynamic ansätze were explored, and their performance compared against predetermined ansätze, with a focus on susceptibility to noise.

Multiple variants of [VQE](#), namely [Unitary Coupled Cluster Singles and Doubles \(UCCSD\)-VQE](#) (predetermined) and [Adaptive Derivative-Assembled Pseudo-Trotter \(ADAPT\)-VQE](#) (dynamic), were implemented both in simulators and cloud quantum computers. Using noise models, the impact of several noise sources on convergence was assessed. Additionally, the importance of the operator pool in [ADAPT-VQE](#) was analysed, and strategies to manipulate the ansatz beyond the [ADAPT-VQE](#) algorithm were explored.

Several conclusions could be drawn from this work. Adapting the ansatz to the problem and system was concluded to be fundamental in avoiding trainability issues, decreasing the circuit depth required for a given accuracy, and improving noise-resilience (against circuit depth dependent and independent sources alike). Dynamic ansätze were shown to be capable of enduring significantly larger error rates than predetermined alternatives, and were thus proved to be better suited for [NISQ](#) devices. For H_2 , as far as ground state energy calculations are concerned, [ADAPT-VQE](#) was shown to tolerate a 20 times lower shot count, 150 times larger error rates in state preparations and measurements, and 850 times lower coherence times than [UCCSD-VQE](#). The difference is expected to increase with the size of the system.

Additionally, it was observed that there is still a margin for improving upon [ADAPT-VQE](#). Further manipulation of the ansatz was shown to be capable of producing yet shallower circuits for the same accuracy. Using an idea previously proposed in the literature, a more conservative selection criterion was tested. Additionally, removing operators on the fly based on available data was attempted as a new possibility. Both approaches were shown to be capable of improving upon the [ADAPT-VQE](#) ansatz, resulting in an up to 35-fold decrease in the error for a similar circuit depth within the first 10 iterations of [ADAPT-VQE](#).

Keywords: adaptive ansätze, quantum chemistry, quantum computing, variational quantum algorithms, variational quantum eigensolver

Ansätze para Eigensolvers Variacionais Quânticos Ruidosos

Simular sistemas quânticos é uma das mais relevantes potenciais aplicações dos computadores quânticos. Os primeiros algoritmos criados para tal visam implementação em computadores quânticos tolerantes a falhas; em contraste, o Eigensolver Variacional Quântico (VQE) procura adaptar-se às limitações dos dispositivos quânticos ruidosos de escala intermédia (NISQ). Há esperança que a classe de algoritmos variacionais quânticos (VQAs), à qual o VQE pertence, seja a primeira a alcançar a vantagem quântica.

O *ansatz* pode ditar o sucesso (ou falta de) de um VQA: a sua profundidade pode impedir a viabilidade a curto prazo, ou conduzir a problemas de treinabilidade que levam à ineficiência. Neste contexto, esta dissertação pretendeu analisar diferentes *ansätze* para química quântica, examinando a sua resiliência ao ruído e viabilidade em dispositivos contemporâneos. Em particular, exploraram-se *ansätze* dinâmicos, cujo desempenho se comparou ao de *ansätze* predeterminados com foco na suscetibilidade ao ruído.

Diferentes variantes do VQE, nomeadamente UCCSD-VQE (predeterminado) e ADAPT-VQE (dinâmico), foram implementadas em simuladores e em computadores quânticos. Com recurso a modelos de ruído, o impacto de diversas fontes de ruído na convergência foi analisado. Inspeccionou-se a importância do conjunto de operadores no ADAPT-VQE, e exploraram-se estratégias adicionais para manipular o *ansatz*.

Várias conclusões puderam ser retiradas deste projeto. Concluiu-se que adaptar o *ansatz* ao problema e ao sistema é fundamental para evitar problemas de treinabilidade, diminuir a profundidade dos circuitos, e melhorar a resiliência ao ruído (seja a fonte dependente da profundidade do circuito ou não). Mostrou-se que *ansätze* dinâmicos toleram taxas de erro superiores aos predeterminados, sendo portanto mais adequados para dispositivos NISQ. Quanto aos cálculos da energia do estado fundamental, mostrou-se que, para a molécula H_2 , o ADAPT-VQE tolera um número de repetições do circuito 20 vezes inferior, taxas de erro na preparação e medição de estados 150 vezes superiores, e tempos de coerência 850 vezes inferiores do que o UCCSD-VQE. É expectável que a diferença aumente com o tamanho do sistema.

Adicionalmente, observou-se que há ainda uma margem para melhorar além do ADAPT-VQE. Mostrou-se que outras formas de manipulação do *ansatz* são capazes de produzir circuitos ainda menos profundos para uma mesma precisão. Usando uma ideia previamente sugerida na literatura, tentou-se usar um critério de seleção mais conservador. Alternativamente, remover operadores com base em dados disponíveis ao longo da execução foi testado como uma nova possibilidade. Mostrou-se que ambas as estratégias são capazes de superar o *ansatz* criado seguindo o protocolo do ADAPT-VQE, resultando num erro até 35 vezes menor para profundidades de circuito comparáveis dentro de 10 iterações do ADAPT-VQE.

Palavras-chave: algoritmos variacionais quânticos, *ansätze* adaptativos, computação quântica, eigensolver variacional quântico, química quântica

Contents

List of Figures	xii
Acronyms	xx
1 Introduction	1
1.1 Context and Motivation	1
1.2 State of the Art	3
1.3 Objectives	7
1.4 Outline of the Document	8
2 Theoretical Framework	10
2.1 Quantum Computing	10
2.1.1 Basic Concepts	10
2.1.2 Measuring Pauli Strings	13
2.1.3 Quantum Simulation and Trotterization	16
2.1.4 Current Limitations	19
2.2 Quantum Chemistry	21
2.2.1 The Electronic Problem	21
2.2.2 Finding the Ground State: Classical Approaches	25
2.3 Variational Quantum Algorithms	33
2.3.1 Motivation	33
2.3.2 Outline	34
2.3.3 Variational Quantum Eigensolver	39
3 Static Ansätze for VQE	48
3.1 Problem Agnostic Ansätze	48
3.1.1 Application to Helium Hydride	48
3.2 Problem Tailored Ansätze	61
3.2.1 Application to Molecular Hydrogen	61
4 Problem Tailored, Dynamically Created Ansätze: ADAPT-VQE	72
4.1 The ADAPT-VQE Algorithm	72
4.1.1 Fermionic-ADAPT-VQE	72

4.1.2	Qubit-ADAPT-VQE	75
4.2	Original Pool Choices	76
4.2.1	Fermion Inspired Pools	76
4.2.2	Qubit Inspired Pools	79
4.3	Application to LiH	82
4.3.1	Evolution of the Optimization	82
4.4	Application to Molecular Hydrogen	83
4.4.1	Bond Dissociation Graph	84
4.4.2	Running the Algorithm on Cloud Quantum Computers	86
4.4.3	Impact of Thermal Relaxation, SPAM errors, and Sampling	88
4.4.4	Comparison with UCCSD-VQE	89
4.5	Application to H4	94
4.5.1	Qubit Pool: Effect of the Jordan Wigner String	95
4.5.2	Fermionic vs Qubit Pools: Performance Comparison	96
5	ADAPT-VQE: Exploring other Pool Options	98
5.1	Motivation	98
5.2	Analysis of the Qubit Pool Operators	99
5.3	Analysis of the Qubit-ADAPT-VQE State	100
5.4	Alternative Pools	101
5.4.1	Generalized Singles and Doubles	101
5.4.2	Eight Pool	101
5.4.3	One Pools	103
5.4.4	Two Pools	103
5.4.5	Four Pools	103
5.5	Graphical Analysis of the State Evolution	108
5.6	General Comparison of Pool Performance	111
6	ADAPT-VQE: Further Ansatz Manipulation	115
6.1	Removing Operators	115
6.1.1	Motivation	115
6.1.2	Heuristics	116
6.1.3	Application to LiH	136
6.1.4	Application to OH-	139
6.2	Conservative Ansatz Growth	143
6.2.1	Comparison with Operator Removal	144
6.3	Discussion	148

7	Conclusions and Future Work	150
	Bibliography	153
	Appendices	
A	Implementation of Unitary Pauli String Exponentials	162
B	Hamiltonian of H2	164
C	Adapt-VQE Gradient Formula	167

List of Figures

1	Representation of the state of a qubit in the Bloch sphere. The angles θ and φ correspond to the parameters from equation 2.4.	11
2	Examples of circuits that perform single-qubit Pauli measurements, drawn in the IBM Quantum Composer [41].	14
3	Examples of circuits that perform a $Z \otimes Z$ measurement, drawn in the IBM Quantum Composer [41].	15
4	Examples of circuits that apply time evolution under simple Hamiltonians H_k , consisting of Pauli strings. They apply $e^{-iH_k t}$ using a parameterized Z rotation. The circuits were drawn in the IBM Quantum Composer [41].	17
5	Representation of a molecule in a coordinate system. Upper case letters (A, B) represent nuclei, while lower case ones (i, j) represent the electrons. Similarly, upper case R represent position vectors pointing to nuclei, while lower case r represent those pointing to electrons. A single subscript is used on the position vector when the vectors point from the origin, while two subscripts are used to represent relative position vectors. In the latter case, the first and second subscripts indicate where the vector points to and from, respectively (e.g. $r_{ij} = r_i - r_j$).	22
6	Typical outline of a variational quantum algorithm.	35
7	Circuit for preparing the Hartree-Fock reference state, for the H_2 molecule. The circuits were drawn in the IBM Quantum Composer [41].	43
8	Circuit with 7 variational parameters for preparing an arbitrary 2-qubit state, drawn in the IBM Quantum Composer [41].	51
9	Circuit with 6 variational parameters for preparing an arbitrary 2-qubit state, drawn in the IBM Quantum Composer [41].	52

10	Performance of the VQE algorithm along the bond dissociation graph of HeH^+ , with and without sampling noise. The upper plot shows the VQE energy, with and without sampling noise, plotted against the exact diagonalization energy (which is exact within the basis set and up to the Born-Oppenheimer approximation and relativistic effects). The lower plot shows the overlap of the VQE state (with sampling noise) with the ground state obtained by exact diagonalization. In the cases with sampling noise, 1000 shots were used. 25 points were used for each interatomic distance; the median was used as the estimator, and the error bars show the interquartile ranges. The quantum circuits were run on the CIRQ simulator [22].	53
11	Performance of the VQE algorithm with the Constrained Optimization By Linear Approximation (COBYLA) optimizer along the bond dissociation graph of HeH^+ , with default hyperparameters. The upper plot in figure 11a compares the single-run VQE energy with the exact diagonalization energy, with and without sampling noise (1000 shots). For each interatomic distance, the same random starting point was used for both cases. The lower plot shows the overlap of the VQE state, obtained in the simulations with sampling noise, with the ground state obtained from exact diagonalization. Figure 11b compares the median and the average of the VQE energy over 15 runs with the exact diagonalization value, in the case where 1000 shots are used per energy evaluation. The quantum circuits were run in the CIRQ Simulator [22].	54
12	Performance of the VQE algorithm with the Nelder-Mead optimizer along the bond dissociation graph of HeH^+ , before (12a, 12c) and after (12b, 12d) hyperparameter optimization. The upper plots in figures 12a, 12b compare the single-run VQE energy with the exact diagonalization energy, with and without sampling noise (1000 shots). For each interatomic distance, the same random starting point was used for both cases. The lower plots show the overlap of the VQE state, obtained in the simulations with sampling noise, with the ground state obtained by exact diagonalization. Figures 12c, 12d compare the median and the average of the VQE energy over 15 runs with the exact diagonalization value, in the case where 1000 shots are used per energy evaluation. The quantum circuits were run in the CIRQ Simulator [22].	56
13	Evolution of the VQE optimization for HeH^+ at an interatomic distance of 90pm, using three different IBM Quantum [42] backends: the statevector simulator (with no noise of any kind), the QASM simulator (with sampling noise) and the Lima device (a 5-qubit Falcon processor). In the last two cases, 1024 shots were used. The energy is plotted in blue and the overlap with the true ground state in red; the pale blue line marks the true ground energy. The same random starting point was used for all backends. The chosen optimizer was COBYLA.	58
14	Evolution of the VQE optimization for HeH^+ at an interatomic distance of 90pm, for different shot counts and using the Lima device. The remaining conditions (starting point, optimizer) are the same as in figure 13.	59

15	Single-run UCCSD-VQE ground energy along the bond dissociation curve of the hydrogen molecule, obtained using IBMQ [42] simulators. The results plotted in figures 15a-15c, ordered by increasing shot count, were obtained using the QASM simulator. The results plotted in figure 15d were obtained using the state vector simulator (equivalent to infinite shots). The red and green curves represent the Full Configuration Interaction (FCI) energy and the Hartree-Fock energy, respectively.	62
16	Bond dissociation curves of figure 15, zoomed in on the region nearing the equilibrium bond length. The blue area marks the region of chemical accuracy (error of less than 1kcal/mol).	63
17	Evolution of the UCCSD-VQE optimization for H_2 at an interatomic distance of 0.74\AA , using three different IBM Quantum [42] backends: the statevector simulator (with no noise of any kind), the QASM simulator (with sampling noise) and the Belem device (a 5-qubit Falcon processor). In the last two cases, 8192 shots were used. The energy is plotted in blue; the pale blue region marks the area of chemical accuracy. The same starting point, the Hartree-Fock reference state, was used for all backends. The chosen optimizer was COBYLA. . .	65
18	Evolution of the VQE optimization under the same conditions of figure 17 (same molecule, interatomic distance and optimizer), plotted for three different shot counts and two different backends. The QASM simulator (18a-18c) portrays the behaviour of an ideal quantum processor, in contrast with IBMQ Manhattan (18d-18f), a 65-qubit Hummingbird processor.	66
19	Impact of the noise associated with different instructions on the performance of UCCSD-VQE, using Belem’s noise model from Qiskit’s [3] <i>providers.aer.noise</i> module. The model is generated automatically from the properties of IBMQ’s [42] device. Each curve represents the UCCSD-VQE optimization, simulated with a model including the noise associated with the corresponding instruction exclusively. 8192 shots were used in all curves. The conditions are the same as those from figure 17 (same molecule, interatomic distance and optimizer).	68
20	Impact of Controlled NOT (CNOT) gate noise on the performance of UCCSD-VQE. The error in the final energy is plotted as a function of the average CNOT gate error (equal to $1 - F$, where F is the average gate fidelity). Twenty runs were used per point; the median was used as the estimator, and the error bars show the interquartile ranges. The number of shots was increased to ensure that sampling noise would not be the limiting factor for any gate error. The conditions are the same as those from figure 17 (same molecule, interatomic distance and optimizer).	69

21	Plots showing the impact of different types of noise on the performance of UCCSD-VQE. For thermal relaxation, a simple model assuming average gate times and a constant ratio between the coherence times T_1 and T_2 was used. Figure 21a shows the evolution of the error in the energy as a function of these times. For State Preparation and Measurement (SPAM) noise, the ratio between the probability of measuring 0 after preparing 1 and the probability of measuring 1 after preparing 0 was fixed (and greater than one, for the model to be realistic). Figure 21b shows the evolution of the error as a function of these probabilities. Figure 21c illustrates the impact of sampling noise by plotting the error as a function of the number of shots. In all of the plots, twenty runs were used per point; the median was used as the estimator and the error bars show the interquartile ranges. All noise models were created using Qiskit [3]. The conditions are the same as those from figure 17 (same molecule, interatomic distance and optimizer).	70
22	Scheme illustrating the ADAPT-VQE algorithm. The green boxes represent processes that are exclusively classical. The orange boxes represent hybrid processes (that include state preparations and measurements in the quantum computer).	73
23	Evolution of the fermionic-ADAPT-VQE energy and gradient norm along the iterations, for the lithium hydride (LiH) molecule at an interatomic distance of 1.45\AA . The convergence criterion was set to a gradient threshold of 0.01. The Singlet Generalized Singles and Doubles (SGSD) pool and the COBYLA optimizer were used. The shaded blue area represents the region of chemical accuracy (error of less than 1kcal/mol). The simulation was done via matrix algebra, and does not include any type of noise.	83
24	Single-run, single-iteration Adaptive Derivative-Assembled Problem-Tailored (ADAPT)-VQE ground energy along the bond dissociation curve of the hydrogen molecule, obtained using IBMQ [42] simulators. The results plotted in figures 24a-24c, ordered by increasing shot count, were obtained using the QASM simulator. The results plotted in figure 24d were obtained using the state vector simulator (equivalent to infinite shots). The red and green curves represent the FCI energy and the Hartree-Fock energy, respectively.	84
25	Bond dissociation curves of figure 24, zoomed in on the region nearing the equilibrium bond length. The blue area marks the region of chemical accuracy (error of less than 1kcal/mol).	85
26	Evolution of first the ADAPT-VQE optimization for H_2 at an interatomic distance of 0.74\AA , using three different IBM Quantum [42] backends: the statevector simulator (with no noise of any kind), the QASM simulator (with sampling noise) and the Belem device (a 5-qubit Falcon processor). In the last two cases, 8192 shots were used. The energy is plotted in blue; the pale blue region marks the area of chemical accuracy. The same starting point, the Hartree-Fock reference state, was used for all backends. The chosen optimizer was COBYLA.	87

27	Plots showing the impact of different types of noise on the performance of ADAPT-VQE for H_2 at an interatomic distance of 0.74Å. For thermal relaxation, a simple model assuming average gate times and a constant ratio between the coherence times T1 and T2 was used. Figure 27a shows the evolution of the error in the energy as a function of these times. For SPAM noise, the ratio between the probability of measuring 0 after preparing 1 and the probability of measuring 1 after preparing 0 was fixed (and greater than one, for the model to be realistic). Figure 27b shows the evolution of the error as a function of these probabilities. Figure 27c illustrates the impact of sampling noise by plotting the error as a function of the number of shots. Twenty runs were used for each point; the median was used as the estimator (as it filters out aberrant runs), and the error bars show the interquartile ranges. All noise models were created using Qiskit [3]. The chosen optimizer was COBYLA.	88
28	Evolution of the optimization of the first iteration of ADAPT-VQE (figure 28a) and of the UCCSD-VQE algorithm, for H_2 at an interatomic distance of 0.74Å. IBMQ's Belem processor was used in both cases.	89
29	Comparison of the effect of different types of noise on ADAPT-VQE and UCCSD-VQE for H_2 at an interatomic distance of 0.74Å. The same noise models described in figures 21 and 27 were used.	91
30	Plots comparing the costs associated with UCCSD-VQE and ADAPT-VQE for a similar final error, in a setting with no noise other than sampling noise. For each interatomic distance, 10 runs of the algorithms were performed. The median of the errors is plotted in figure 30a. Figures 30b, 30c, and 30d present respectively the average number of energy evaluations, measured observables, and total shots per run. The number of shots per Pauli string was set to 2^{13} and 2^{18} for ADAPT-VQE and UCCSD-VQE respectively. These values were chosen so that the error was roughly matched between the two algorithms (the average error was 0.00038 a.u. and 0.00039 a.u. by the same order).	93
31	Impact of the Jordan-Wigner string on the convergence of the qubit-ADAPT-VQE algorithm. The graphs compare the evolution of the error in the energy along the 20 first iterations of the algorithm, with and without the Z-strings in the pool operators. The molecule in study is the H_4 molecule at an interatomic distance of 1.45Å.	95
32	Evolution of the error in the ADAPT-VQE energy along the evolution of the algorithm for the Spin-Complemented Generalized Singles and Doubles (SCGSD) pool, the SGSD pool, and the qubit pool. The molecule in study is the H_4 molecule.	96
33	Circuit for implementing the operator in 5.8 using the ladder-of-CNOTs method explained in subsection 2.1.3. The circuit was drawn in the IBM Quantum Composer [41].	107
34	Circuit for implementing the operator in 5.8 using the more efficient method introduced in [98]. The circuit was drawn in the IBM Quantum Composer [41].	107

35	Evolution of the 20 first iterations of the ADAPT-VQE algorithm for three different pools: the XXYX One Pool, the XXYX Two Pool, and the original Qubit Pool (without the Jordan-Wigner string). The upper plot shows the evolution of the error in the energy. The lower plot shows the evolution of the total probability amplitude of computational basis states representing Slater determinants with an altered particle number or Z spin projection. The molecule in consideration is the H_4 molecule at an interatomic distance of 1.45Å.	108
36	Evolution of the 20 first iterations of the ADAPT-VQE algorithm for different pools. The upper plot shows the evolution of the error in the energy. The lower plot shows the number of Slater determinants in the superposition state with altered particle number, S_z , or neither. In 36c and 36d, altered S_z and particle number correspond to the same colour because the curves coincide. When the green curve is not visible, it lies under the orange one. The molecule in consideration is H_4 at an interatomic distance of 1.45Å.	110
37	Comparison of the convergence of the ADAPT-VQE algorithm using several different One Pools (figure 37a) and Two Pools (figure 37b) against the full Qubit Pool (without Jordan-Wigner strings). In figure 37b, a curve corresponding to a One Pool was also plotted for reference. The error in the energy is plotted as a function of the iteration number. The molecule in consideration is H_4 at an interatomic distance of 1.45Å.	112
38	Comparison of the convergence of the ADAPT-VQE algorithm using the Generalized Singles and Doubles (GSD) Pool and the Eight Pool. The error in the energy is plotted as a function of the iteration number. The molecule in consideration is H_4 at an interatomic distance of 1.45Å.	113
39	Impact of the choice of pool (between One, Two, Four, and Eight Pools) in the ADAPT-VQE algorithm, for two different molecules. The error in the energy is plotted against the iteration number. In figure 39b, the XXYX Four Pool and the Eight Pool correspond to the same color because the curves coincide. These curves also overlap in some iterations in figure 39a.	114
40	Norm of the selected operator and change in energy in the first ten iterations of ADAPT-VQE. The molecule in consideration is the LiH molecule at an interatomic distance of 1.45Å. The SGSD pool was used.	117
41	Absolute value of the performance ratio in the first ten iterations of ADAPT-VQE, for the same pool and molecule as figure 40. The performance ratio at a given iteration corresponds to the operator added then.	119
42	Relevant data structures at iteration 3 of the ADAPT-VQE algorithm, with an additional term removal feature based on blocking and unblocking operators. The molecule in consideration is LiH , and the pool the SGSD pool. The units of energy, omitted to avoid overcrowding of the figures, are Hartree throughout the whole chapter. Because the variational parameters are dimensionless, the gradient shares the energy units.	122
43	Relevant data structures at iteration 4.	123

LIST OF FIGURES

44	Relevant data structures at iteration 7.	124
45	Relevant data structures at iteration 8.	125
46	Relevant data structures at iteration 9.	125
47	Relevant data structures at iteration 11.	126
48	Relevant data structures at iteration 12.	126
49	Relevant data structures at iteration 14.	127
50	Relevant data structures at iteration 3 of the ADAPT-VQE algorithm, with an additional term removal feature based on a performance penalty. The molecule in consideration is LiH , and the pool the SGSD pool.	131
51	Relevant data structures at iteration 4.	132
52	Evolution of the relevant data, iteration 5 through 13. On top, we have data on the evolution of the algorithm and on the selected operators. On the bottom, we have data specifically regarding operator 25, that was removed at iteration 4.	132
53	Relevant data structures at the end of iteration 11.	135
54	Plot comparing the original ADAPT-VQE and the version with term removal, for LiH at an interatomic distance of 1.45\AA and using the SGSD pool. The error in the energy (top) and the number of operators in the ansatz (bottom) are plotted against the iteration number. The shaded blue area on the upper plot marks the error values within chemical accuracy (less than 1kcal/mol).	137
55	Plots comparing the energy and energy error in the original ADAPT-VQE and the version with term removal, for LiH at an interatomic distance of 1.45\AA and using the SGSD pool. The shaded blue areas on the plots mark the values within chemical accuracy both for the energy (less than 1kcal/mol away from the FCI energy) and the error (less than 1kcal/mol).	139
56	Plot comparing the original ADAPT-VQE and the version with term removal, for OH^- at an interatomic distance of 1.45\AA and using the SGSD pool. The error in the energy (top) and the number of operators in the ansatz (bottom) are plotted against the iteration number. The shaded blue area on the upper plot marks the error values within chemical accuracy (less than 1kcal/mol).	140
57	Plots comparing the energy and energy error in the original ADAPT-VQE and the version with term removal, for OH^- at an interatomic distance of 1.45\AA and using the SGSD pool. The shaded blue areas on the plots mark the values within chemical accuracy both for the energy (less than 1kcal/mol away from the FCI energy) and the error (less than 1kcal/mol).	141
58	Plot comparing the original ADAPT-VQE and the version with term removal, for OH^- at an interatomic distance of 1.45\AA and using the Eight Pool. The error in the energy (top) and the number of operators in the ansatz (bottom) are plotted against the iteration number. The shaded blue area on the upper plot marks the error values within chemical accuracy (less than 1kcal/mol).	142

59	Plots comparing the energy and energy error in the original ADAPT-VQE and the version with term removal, for OH^- at an interatomic distance of 1.45\AA and using the Eight Pool. The shaded blue areas on the plots mark the values within chemical accuracy both for the energy (less than 1kcal/mol away from the FCI energy) and the error (less than 1kcal/mol).	143
60	Plot showing the minimum error reached for each ansatz size, for multiple ansatz growth strategies. The molecule in case is LiH at an interatomic distance of 1.45\AA . The SGSD pool was used. In the points where the green curve, corresponding to conservative ansatz growth, is not visible, it lies directly under the term removal curve.	144
61	Plot showing the minimum error reached for each ansatz size, for multiple ansatz growth strategies. The molecule in case is OH^- at an interatomic distance of 1.45\AA . The Eight Pool was used. The term removal and conservative ansatz growth strategies correspond to the same color because the curves coincide.	146
62	Plot showing the minimum error reached for each ansatz size, for multiple ansatz growth strategies. The molecule in case is H_4 at an interatomic distance of 1.45\AA . The original ADAPT-VQE and the version with term removal correspond to the same color because the curves coincide.	147

Acronyms

ADAPT	Adaptive Derivative-Assembled Pseudo-Trotter vii , xv , xvi , xvii , xviii , xix , 6 , 7 , 8 , 9 , 36 , 72 , 73 , 74 , 75 , 76 , 82 , 83 , 94 , 96 , 97 , 98 , 99 , 112 , 113 , 115 , 116 , 117 , 118 , 119 , 122 , 123 , 130 , 131 , 135 , 136 , 137 , 138 , 139 , 140 , 141 , 142 , 143 , 145 , 147 , 148 , 150 , 151 , 152
ADAPT	Adaptive Derivative-Assembled Problem-Tailored xv , xvi , xvii , 7 , 8 , 9 , 72 , 75 , 76 , 79 , 80 , 81 , 82 , 83 , 84 , 85 , 86 , 87 , 88 , 89 , 90 , 91 , 92 , 93 , 94 , 95 , 96 , 97 , 98 , 99 , 100 , 101 , 108 , 109 , 110 , 111 , 112 , 113 , 114 , 150
BCH	Baker–Campbell–Hausdorff 18 , 32
CC	Coupled Cluster 5 , 31 , 32 , 33 , 61 , 72
CCSD	Coupled Cluster Singles and Doubles 31
CI	Configuration Interaction 30 , 31
CISD	Configuration Interaction Singles and Doubles 31
CNOT	Controlled NOT xiv , xvi , 8 , 15 , 17 , 51 , 61 , 64 , 67 , 68 , 69 , 75 , 83 , 90 , 96 , 97 , 106 , 107 , 112
COBYLA	Constrained Optimization By Linear Approximation xiii , xiv , xvi , 54 , 55 , 57 , 58 , 65 , 88
FCI	Full Configuration Interaction xiv , xv , xviii , xix , 30 , 48 , 61 , 62 , 63 , 69 , 82 , 84 , 85 , 116 , 139 , 141 , 143
GSD	Generalized Singles and Doubles xvii , 101 , 112 , 113
NISQ	Noisy Intermediate-Scale Quantum vii , 3 , 4 , 5 , 7 , 8 , 19 , 20 , 34 , 75 , 86 , 96 , 150 , 152
QEB	Qubit-Excitation-Based 102 , 143
q-UCCD	Quantum Unitary Coupled Cluster Doubles 44
QAA	Quantum Adiabatic Algorithm 34
QAOA	Quantum Approximate Optimization Algorithm 34 , 38 , 39
QPE	Quantum Phase Estimation 4 , 33 , 34
QUBO	Quadratic Unconstrained Binary Optimization 34

RDM	Reduced Density Matrix 7
SAGSD	Spin-Adapted Generalized Singles and Doubles 76
SCF	Self-Consistent Field 27, 28
SCGSD	Spin-Complemented Generalized Singles and Doubles xvi, 72, 78, 79, 96, 97
SGSD	Singlet Generalized Singles and Doubles xv, xvi, xvii, xviii, xix, 76, 78, 79, 80, 83, 96, 97, 116, 117, 122, 131, 136, 137, 139, 140, 141, 144, 148
SPAM	State Preparation and Measurement xv, xvi, 8, 9, 19, 20, 21, 33, 57, 61, 70, 71, 88, 89, 90, 91, 150
SVD	Singular Value Decomposition 49, 50, 52
UCC	Unitary Coupled Cluster 5, 32, 44
UCCGSD	Unitary Coupled Cluster Generalized Singles and Doubles 44
UCCSD	Unitary Coupled Cluster Singles and Doubles vii, xiv, xv, xvi, 5, 6, 7, 8, 9, 32, 33, 36, 44, 61, 62, 63, 64, 65, 66, 67, 68, 69, 70, 71, 72, 75, 76, 82, 83, 84, 85, 86, 89, 90, 91, 92, 93, 94, 150
UpCCGSD	Unitary Paired Coupled Cluster Generalized Singles and Doubles 44
VCC	Variational Coupled Cluster 32
VQA	Variational Quantum Algorithm vii, 3, 4, 8, 34, 38, 54, 61
VQE	Variational Quantum Eigensolver vii, xiii, xiv, xv, xvi, xvii, xviii, xix, 4, 5, 6, 7, 8, 9, 33, 36, 38, 39, 43, 44, 46, 48, 49, 52, 53, 54, 55, 56, 57, 58, 59, 61, 62, 63, 64, 65, 66, 67, 68, 69, 70, 71, 72, 73, 74, 75, 76, 79, 80, 81, 82, 83, 84, 85, 86, 87, 88, 89, 90, 91, 92, 93, 94, 95, 96, 97, 98, 99, 100, 101, 102, 108, 109, 110, 111, 112, 113, 114, 115, 116, 117, 118, 119, 122, 123, 130, 131, 135, 136, 137, 138, 139, 140, 141, 142, 143, 145, 147, 148, 150, 151, 152

Introduction

1.1 Context and Motivation

Richard Feynman raised the question,

What kind of computer are we going to use to simulate physics?

, calling attention to the fact that the physical world is quantum mechanical, so *quantum* physics is what we are concerned with [25]. So arose the answer: a *new* kind of computer.

The simulation of a quantum system being a difficult task to perform classically is but a symptom of the fact that bits, as basic units of information in classical computing, are unfit to describe nature under the laws of quantum mechanics. Quantum computing was thus born with the purpose of using quantum mechanics itself to perform computation, replacing bits by their quantum counterparts: *qubits*. Making use of quantum phenomena (namely entanglement; superposition; interference) empowers quantum computers with the ability to do operations without a classical analogue.

The scaling of the description of highly entangled states using classical data inspired hope that quantum computers would be able to perform tasks that are intractable even to classical supercomputers. The memory required for storing a general state of a quantum system in a classical computer scales exponentially on the size of the system; in contrast, the scaling is linear on a quantum computer, as seems intuitive in view of it being such a system itself. In fact, an interesting reason to believe that a quantum computer may bring advantage is that there is no known classical algorithm that can simulate such a computer efficiently; representing the state of a few hundred qubits may already require more bits than the number of atoms in the visible universe [69].

However, this fact does not constitute in itself proof that quantum computers are more powerful than the classical ones. It is a misconception that quantum computers are *exponentially parallel*: even though superposition implies that qubits can be in a linear superposition of states that are exponentially many on the number of qubits, the information held by this superposition isn't necessarily accessible. Measuring the state of the qubits will cause the collapse of the wave function: transforming quantum information into classical information destroys it irreversibly, unless the measurement reveals no information about the quantum state it is inflicted upon [64].

Famously, by Holevo's theorem, the accessible classical information in n qubits is bounded above by n bits. The fact that the qubits carry more information is to no avail, unless one devises a clever way of extracting information in an efficient manner. As such, the design of techniques that allow employing

quantum resources advantageously is all but straightforward; and the quest for quantum algorithms has developed into a large field of research.

In 1980, Benioff published the first article on quantum computation, a pioneering work describing a quantum mechanical model of Turing machines [12]. Since then, there has been significant progress. Notable quantum algorithms were introduced in the following decade, such as Deutsch–Jozsa in 1992 [21] and Simon’s in 1994 [78]. Albeit not having much practical value, they hold their merit as some of the first examples of how a quantum computer could offer an exponential speedup over classical computers. Advantage of practical interest ensued when Shor, inspired by Simon’s proposal, introduced his famous factoring algorithm [76].

Shor’s algorithm enables a (hypothetical) quantum computer to factor integers in a time polynomial on their size - an exponential improvement as compared to the best-known classical algorithm for the task. While this remarkable speedup was not unprecedented, Shor’s algorithm owes its distinct success to the possible applications of the problem it was created to solve. An efficient factoring algorithm defies the security of the vastly used RSA cryptosystem, which brought attention to the possible implications of quantum computers. If resource requirements in what concerns quantum hardware were met, this algorithm would allow breaking internet security: it manages to solve in a matter of minutes a problem that would take years for the best-known classical algorithm to solve, and upon the difficulty of which the RSA algorithm relies [70]. This was a significant landmark, as well as a stronger motivation for developing the hardware that would pave the way for physically realising these theoretical proposals.

Albeit offering an unarguably less astounding quadratic speedup, Grover’s algorithm [36] for searching in an unstructured database has also drawn formidable attention. Aside from covering a wide array of applications, its strength is in offering a provable speedup: evidently, the number of evaluations required to solve the problem classically has to scale linearly (half of elements in the search space will have to be checked on average). By offering the uncanny possibility of finding a marked element after a number of operations that is sub-linear on the size of the search space, Grover’s algorithm also contributed to a build up of the public reputation of quantum computing as a curious and strange field.

While they seem compelling, these theoretical results are yet to be proven by experimental realization. Current hardware limitations hinder the implementation of large-scale instances of algorithms such as the aforementioned ones.

There has been notable progress since the concept of quantum computing was introduced. The first implementation of an algorithm on a physical quantum computer took place in 1998 [18], and several proof-of-concept demonstrations of small-scale instances of different algorithms followed. In 2012, Shor’s algorithm was used to factor $N=21$ for the first time [58].

However, practical experiments up to now have been limited to toy problems and small, classically tractable demonstrations. Building a quantum computer capable of outperforming a classical one in tasks that may be of practical use seems to be strenuous at best, impossible at worse.

The difficulty of building the physical hardware is indissociable from the very same nature on top of which the concept of quantum computing rests. Quantum mechanical states are very sensitive, causing

qubits to require close to perfect isolation from the world to store and process information reliably. Building a quantum computer thus implies granting the proper conditions, i.e. conditions that reduce the coupling with the surrounding environment to the point that useful computations can be carried. Often, this involves temperatures near absolute zero and shielding from radiation; this is aggravated by the fact that the difficulty greatly increases with the system size, preventing large-scale quantum computers to have been built so far [70].

There is a dichotomy between the scaling of the most compelling quantum algorithms, and the scaling of the difficulty of building a quantum computer capable of actually bringing them to fruition. In the words of two prominent French physicists vocally skeptical of quantum computing [37],

...the large-scale quantum machine, though it may be the computer scientist's dream, is the experimenter's nightmare.

After Shor's algorithm was discovered, concerns were raised regarding the viability of quantum computing: algorithms built under the assumption of an ideal quantum computer might be rendered useless by decoherence and other types of noise [52, 86]. In light of the *no-cloning* theorem [23, 66, 96], that forbids quantum data from being copied, error correction didn't seem as trivial as in the classical case.

In 1995, Peter Shor himself recovered some hope with a viable way of realising error correction in quantum computing [77]. Quantum error correcting codes are of the greatest importance: if error correction is important for computing in general, it is pivotal for quantum computing in particular. Unfortunately, the overhead in terms of qubit requirements and number of necessary gates is severe [31].

An error corrected, fault-tolerant quantum computer does not seem to belong to a near future. John Preskill coined the term **Noisy Intermediate-Scale Quantum (NISQ)** to describe the current era [69]. The name encompasses qubit limitations as well as noise and decoherence.

After the promising algorithms of the 1990s shed some light on the potential of a fault-tolerant quantum computer, there has been a growing effort to build such a device - but that possibility seems at best distant. In an attempt to make as good use of **NISQ** computers as possible, the attention shifted in part from the asymptotic scaling of algorithms to their near-term viability. The focus has widened, and thus, in addition to creating conceptually advantageous algorithms lacking certainty that they are ever to be implemented, we now observe a growing dedication to devising algorithms that seem implementable but aren't necessarily known to be advantageous. Variational quantum algorithms have emerged.

1.2 State of the Art

VQAs are hybrid quantum-classical algorithms that resemble machine-learning techniques typical to classical computing. The usual outline involves using a classical computer to find the variational parameters that minimize a cost function, which is evaluated by the quantum computer. An interesting review of the structure, applications and challenges of variational quantum algorithms can be found in reference [17].

The role of the quantum computer in variational algorithms is to prepare trial states using a state preparation circuit (called *ansatz*), and measure the expectation value (and possibly the gradients) of the cost function in those states. The state preparation circuit includes a parameterized portion, the parameters of which are fed to the quantum computer each iteration by the classical optimizer, which tries to adjust them in an attempt to find the minimum.

VQAs are currently gathering great interest, seeing that their very nature makes them resilient to noise: by outsourcing the optimization to a classical computer, they naturally require shallower circuits. The learning-based approach further aids noise mitigation.

In 2014 the first VQA was proposed and demonstrated experimentally with photonic qubits: the VQE [67], central to this dissertation. The proposal aimed to address quantum chemistry problems using near-term quantum computers, in alternative to longer routines (such as adiabatic state preparation in conjunction with Quantum Phase Estimation (QPE)) that are beyond the capabilities of NISQ devices.

Since it was introduced, the VQE has become an icon of NISQ algorithms, and a lot of research has been done to further increase its accuracy and decrease its demands on the quantum computer. There have been improvements on the efficiency of the evaluation of the cost function for chemistry problems in particular, involving a low rank factorization of the two-electron integral tensor [40]. Interestingly, making use of the specificity of the chemistry problem seems to allow for a better scaling of the required number of measurements as compared to general problems [17]. However, more general techniques can evidently be used to further increase the accuracy of the results; as an example, the readout-error mitigation technique presented in [13] can help eliminate the noise-induced bias from the expectation value of the cost function.

The qubit requirements can also be reduced: in reference [82] it was highlighted that it is possible to choose a more compact fermion to qubit mapping instead of the Jordan-Wigner transform employed in the original article, at the expense of more gates in the circuits.

The optimization itself is a large field of research. While not belonging to the strictly quantum part of the algorithm, it is naturally affected by quantum noise and its particularities; the performance of classical optimizers benchmarked with noiseless function evaluations cannot simply be extrapolated to a noisy scenario. Efforts have been done to find optimizers that decrease the number of calls to the quantum computer and make convergence more likely in the presence of noise. Meta-learners were used in [95] and two model-based optimization algorithms were introduced in [83]. These examples are left as a sample of a much broader set of proposals for improving the optimization.

Finally, we have the *ansatz*, which plays a fundamental role in the algorithm. It consists of the state preparation circuit that defines where we will be looking for in the Hilbert space. From the concept alone, one can conclude that the *ansatz* should satisfy two crucial aspects: contain the solution, and have a description that does not scale exponentially with the size of the system. Otherwise, it may jeopardise the accuracy and efficiency of the algorithm.

In the context of the NISQ era, this key piece assumes an even greater importance: the *ansatz* corresponds to the bulk of the depth of the utilized circuits, which in turn means that the choice of *ansatz* might allow or prevent the VQE from being successfully implemented on a given quantum computer.

Further, it has been shown that an improper ansatz might prevent the optimizer from finding the solution. As an example, randomly initialized parameterized quantum circuits were shown to cause the gradients of the cost function with respect to the parameters to vanish exponentially on the size of the system, an effect termed as barren plateaus [60]. The effect is not limited to gradient-based optimizers, as this description might seem to suggest - the result was proven to generalize to other optimization methods [10]. The vanishing gradients are merely a symptom of a deeper issue regarding the optimization landscape. As the size of the Hilbert space grows exponentially on the number of qubits, it is no surprise that a generic ansatz with polynomially many parameters proves to hinder the performance of the algorithm. The minimum becomes sharper and sharper amidst a flat landscape, as a result of too much expressibility. A relationship between expressibility and trainability was derived in [39].

Multiple strategies have been used to address this. Some involve the ansatz and the formulation of the problem itself: it was shown that shallow ansätze in conjunction with local cost functions circumvent the issue. For circuits with a depth logarithmic on the number of qubits, the gradient vanishes polynomially at most [16, 87]. Parameter initialization techniques that prevent the algorithm from starting in a barren plateau have been proposed. Reference [32] does this by limiting the initial effective depth of the circuit, while in [88] meta-learning via classical neural networks was used to find good parameter initialization heuristics. Layerwise training [79] and correlated parameters [90] were proved to aid in avoiding barren plateaus *during* the optimization. As for the classical optimizer, making use of natural evolutionary strategies in the optimization as a means to avoid barren plateaus was suggested in reference [9].

Another interesting possibility is to leverage problem-specific knowledge to prevent the issue. Considering that barren plateaus are tightly related to an excess of expressibility, leveraging what is known about the problem at hand to carefully pick where we will be looking for the solution is of great help. There is a downside: the priority is placed in tailoring the ansatz *to the problem* rather than to the quantum processor. The latter seems preferable: in the context of NISQ devices, the focus is usually set on making the circuit as shallow as possible. Because of this, a popular suggestion was building the ansatz from the interactions naturally available in the quantum processor - an idea that was introduced in [49] as the *hardware-efficient ansatz*. However, compared to these more *generic problem agnostic* ansätze, those known as *problem tailored* or *problem inspired* have the great benefit of usually being trainable, even when the parameters are initialized at random [17]. Further, the ansatz circuit being convenient is to no avail if the optimization process becomes intractable due to requiring exponential precision and function evaluations to overcome barren plateaus.

The ansatz employed in the original VQE proposal was a problem tailored one: the UCCSD ansatz, inspired on quantum chemistry and classical variational methods. States are prepared by applying an operator, consisting of an exponentiated sum of single and double fermionic excitations, to a reference state. A more detailed description of UCCSD can be found in [71]. What is relevant here is that Unitary Coupled Cluster (UCC) is the natural unitary version of the classical Coupled Cluster (CC) operator, regarded as *the gold standard* in classical quantum chemistry methods [11].

Aside from its foundations providing us a motive to believe that the UCCSD ansatz is a good choice,

both the number of variational parameters and the number of gates necessary for its implementation scale polynomially with the size of the system. However, implementing it with a quantum circuit implies using a finite-order Suzuki-Trotter decomposition, which leads to ambiguities in the operator ordering. The resulting energy variations have been proved relevant at the chemical scale, meaning that the ansatz is not chemically well-defined [43]. Additionally, it was pointed out in [93] that this ansatz has an unnecessarily large number of parameters and terms, which then reflect in an optimization more difficult and circuits deeper than required.

It is fundamental to keep the circuit as shallow as possible. In fact, while UCCSD may avoid the barren plateaus discussed above, barren plateaus of a different kind were shown to occur as a result of incoherent noise [92]. In consequence of the output state being less and less pure, the minimum of the cost function becomes less and less pronounced. While with the noise-free barren plateaus the minimum value of the cost function was still represented in the optimization landscape, noise-induced barren plateaus might erase this minimum altogether, with decoherence causing the dilution of the characteristics of the cost function.

In light of this, efforts have been done to combine problem tailoring with hardware tailoring, in order to avoid barren plateaus while keeping circuits convenient and shallow. A few suggestions for chemistry applications follow.

Reference [26] proposed efficient ansätze that respect particle number, total spin, spin projection, and time-reversal symmetries. The idea is to prevent the ansatz from covering irrelevant subsections of the Hilbert space, which allows spanning the relevant symmetry subspace with as few variational parameters as possible. As an example, we can be certain that the ground state will not be found in a region of the Hilbert space that represents a fermionic state with an incorrect number of particles. The performance of the algorithm is thus improved by excluding that region from the ansatz altogether.

An interesting hardware-efficient, but problem-tailored ansatz was proposed in [61]. The algorithm, termed ctrl-VQE by the authors, replaces the traditional quantum circuit ansatz with a quantum control routine. Directly optimizing the pulses allows for reducing the coherence time requirements without abdicating the variational freedom.

Finally, there are the *dynamically created ansätze*, which will be the main focus of this dissertation. While problem inspired ansätze *leverage* knowledge on the system, the dynamically created ones *collect* this knowledge themselves. In a way, they allow the problem to dictate its own ansatz, by growing it from scratch in a strongly system-adapted manner. Fermionic operators are added to the ansatz as ‘pieces’ along the evolution of the algorithm; the chosen ones are those that, from informative measurements, are believed to be the most beneficial.

In 2019, this idea was introduced in [35] as *Adaptive Derivative-Assembled Pseudo-Trotter (ADAPT)-VQE*, hereby denoted *Fermionic-ADAPT-VQE* for clarity against a more recent version. The name was attributed in light of the selection criterion being derivative-related and the circuits being reminiscent of a trotterization.

While the UCCSD ansatz is fixed upfront and does not allow decreasing or increasing the precision

of the result, or the circuit depth, [ADAPT-VQE](#) grows the ansatz from zero until a convergence criterion is met. Numerical results showed that this protocol can match or surpass the accuracy of [UCCSD-VQE](#) with more compact ansätze. In essence, we are trading measurements and optimizations for shallower circuits, making [VQE](#) more [NISQ](#)-friendly and bringing us a step closer to solving real problems on quantum computers.

In 2020 a follow-up article [85] was published proposing *Qubit-Adaptive Derivative-Assembled Problem-Tailored (ADAPT)-VQE*, a modified version of the algorithm. The difference is in the operators that are available to choose from. Anti-commutation and preservation of particle number and spin are dispensed with, allowing the operators to be implemented by shallower circuits. As a consequence, they no longer have an interpretation as fermionic excitations. This typically results in an increase in the number of variational parameters and decrease of circuit depth for a given accuracy, as compared to Fermionic-[ADAPT-VQE](#). [Qubit-ADAPT-VQE](#) thus burdens the classical optimizer further, in exchange for shallower circuits.

The [Qubit-ADAPT-VQE](#) article also introduced some minimal complete pools whose size scaled linearly with the number of spin-orbitals (as compared to the quartic scaling of all previous pools). This reduction in pool size implies a reduction in the number of measurements per iteration. Unfortunately, because these pools do not respect fermionic symmetries, they are not adequate when [ADAPT-VQE](#) is used to solve for the eigenstates of molecular Hamiltonians.

In September 2021, the article [75] introduced a way of creating symmetry-adapted minimal complete pools that worked in this type of problems, while also growing linearly on the size of the system. With this, the article proved that it was possible to significantly reduce the number of measurements per iteration in [ADAPT-VQE](#), from $\mathcal{O}(N^8)$ to $\mathcal{O}(N^5)$ (N the number of spin-orbitals). This brings the measurement scaling of [ADAPT-VQE](#) closer to that of typical [VQE](#) routines ($\mathcal{O}(N^4)$).

The [ADAPT-V](#) and [ADAPT-Vx](#) algorithms introduced in [55] improve the scaling of the number of measurements using a different approach, based on [Reduced Density Matrices \(RDMs\)](#). In [ADAPT-V](#) the measurement cost scales like $\mathcal{O}(N^4)$, the drawbacks being a larger number of variational parameters and deeper circuits. [ADAPT-Vx](#), in turn, leads to parameter vector lengths and circuit depths that are barely over those in [ADAPT-VQE](#), at the expense of a slight measurement overhead as compared to [ADAPT-V](#).

1.3 Objectives

This dissertation is focused on analysing increasingly [NISQ](#)-friendly ansatz proposals for variational quantum algorithms in chemistry applications. This encompasses shallower ansätze, as well as improved noise resilience.

Quantum chemistry and variational quantum algorithms will be reviewed, as well as relevant aspects of quantum computing.

Instances of [VQE](#) with predetermined ansätze, both problem agnostic and problem tailored, will be

studied, implemented, and simulated. The same analysis and simulation is to be done for dynamically created ansätze (fermionic-ADAPT-VQE and qubit-ADAPT-VQE). In all cases, numerical simulation will involve the study of the impact of noise on the performance of the algorithm. This encompasses sampling noise as well as coherent, incoherent and SPAM noise, simulated resorting to noise models, and implementation of small-scale instances in real quantum computers.

The impact of the pool in the ADAPT-VQE algorithm will be analysed by comparing the one from qubit-ADAPT-VQE, the one from fermionic-ADAPT-VQE, and other options. The impact of the choice of pool on convergence will be assessed, as well as on the conservation of expectation values of fermionic observables, which is tightly related to convergence itself.

Finally, approaches will be tested in an attempt of further reducing the ansatz depth required for a given accuracy. Heuristics for removing operators from the ansatz after their addition will be developed. Additionally, it will be attempted to grow the ansatz with a slower, more conservative method. Both strategies will be compared.

1.4 Outline of the Document

The structure of the dissertation is as follows. Chapter 2 explains the fundamental theoretical framework. Section 2.1 covers the basic concepts of quantum computing, with a focus of the aspects that are relevant for VQE (quantum expectation estimation, quantum simulation and trotterization, and the limitations of quantum computers in the NISQ era). Section 2.2 follows with the essential aspects of quantum chemistry, namely, the context in which the electronic problem arises, and the classical variational approaches that preceded (and inspired) the proposal of VQE. Finally, section 2.3 explains the context, theory, and methods that are specific to variational quantum algorithms. The general outline of VQAs is explained in this section, as well as the specific aspects concerning VQE, such as the quantum circuit representation of fermionic operators, the possible types of ansatz, the Hamiltonian averaging, and the shot requirements necessary to achieve chemical accuracy.

Chapter 3 introduces some static ansätze for VQE, with experimental results that illustrate the viability of the procedure. The first section, 3.1, covers simulations using a problem-agnostic ansatz. The bond dissociation curve of the helium hydride ion, that was obtained in the first implementation of VQE (in a photonic chip) [67], is reproduced. A circuit to span the whole state space of 2 qubits was devised, and the influence of the choice of optimizer and optimization hyperparameters studied. The implementation was tested both in simulators and real devices (IBMQ [42] backends). Section 3.2 implements VQE with the UCCSD ansatz, a problem-tailored ansatz, and presents results for molecular hydrogen. Once again, the implementation was tested in real devices as well as simulators. The influence of sampling noise was assessed, as well as the influence of CNOT gate errors, and of thermal relaxation and SPAM errors.

The concept of dynamic ansätze is explained in chapter 4, with the introduction of the ADAPT-VQE algorithm. Several operator pools used in the set of articles concerning this algorithm are explained,

analysed, and compared both in concept and in practice. The algorithm was applied to several molecules in noise-free scenarios, as well as in real quantum computers. The impact of sampling noise, thermal relaxation and SPAM errors were once again assessed, and compared with the results from the previous chapter (for UCCSD-VQE).

Other possibilities of pools, not contemplated in the original ADAPT-VQE articles, are introduced in chapter 5. The purpose of these new pools is to assess how convergence is affected by the affinity between pool operators and fermionic excitations, in multiple aspects. The operators in the original pools, and their effect on the state of the qubits, are analysed as motivation for the definitions of the new pools. The performance of ADAPT-VQE with the multiple pools is compared, and the composition of the state in terms of Slater determinants is analysed.

Chapter 6 introduces two additional methods for manipulating the ansatz beyond what is done in ADAPT-VQE, with the purpose of achieving greater accuracy with shallower circuits. Section 6.1 develops and applies heuristics for removing operators from the ansatz on the fly, and applies the strategy to multiple molecules. An alternative procedure, that attempts multiple optimizations per iteration before settling for a new operator, is explained and tested in section 6.2. The approaches are compared against each other, and against the original ADAPT-VQE algorithm. The chapter ends with a discussion of the merits and disadvantages of the different strategies (section 6.3).

Finally, chapter 7 concludes the dissertation with a few considerations about the multiple proposals of ansätze for VQE. The merits and demerits of each are discussed in light of the state-of-the-art of quantum computing hardware. Prospects of future work concerning the topics that were covered throughout the dissertation are discussed.

All the code created for this dissertation and used in obtaining the results presented in this document can be found online at the GitHub software repository [8].

Theoretical Framework

2.1 Quantum Computing

The purpose of this section is to offer a simple overview of the concepts in quantum computing that are necessary for the following chapters. A more in-depth explanation on these topics can be found on the book Quantum Computation and Quantum Information, reference [64], upon which the following was based.

2.1.1 Basic Concepts

2.1.1.1 Pure Single Qubit States

The basic unit of information for classical computing, the *bit*, undertakes two possible states, usually labeled **0** and **1**. Similarly, the *quantum bit*, or *qubit*, is a two-level system. The difference resides in the quantum mechanical nature of the qubit: while the classical bit is either in state **0** or **1**, the quantum bit can be in a superposition of these two states.

A qubit can be physically implemented by a two-level quantum-mechanical system. Irrespective of the specific implementation, the state of an isolated qubit can be represented abstractly in Dirac notation as

$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle. \quad (2.1)$$

In the expression above $|0\rangle$ and $|1\rangle$ are the *computational basis states*; they form an orthonormal basis that spans the state space of the system, a two-dimensional Hilbert space \mathcal{H} (a complex vector space equipped with inner product). By convention,

$$|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad |1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}. \quad (2.2)$$

The numbers α and β are the respective complex amplitudes, related to the respective probability of occurrence. As discussed previously, even though they characterize the state of the qubit, it is not possible to discover the values of these numbers with access to a single copy of the state: when a measurement is done on the computational basis, the outcome will be a computational basis state.

The squared absolute values of the amplitudes correspond to the probability of, upon measurement, obtaining state $|0\rangle$ or $|1\rangle$ respectively:

$$\alpha\alpha^* = |\alpha|^2 = \text{Probability of } |0\rangle,$$

$$\beta\beta^* = |\beta|^2 = \text{Probability of } |1\rangle.$$

Since the probabilities must sum to one, we have a normalization condition that reads

$$\langle\psi|\psi\rangle = |\alpha|^2 + |\beta|^2 = 1. \quad (2.3)$$

To describe the state of a perfectly isolated qubit, it thus suffices to specify the corresponding unit vector in \mathcal{H} . This can be done by specifying the values of α or β in equation 2.1 subject to condition 2.3.

The normalization removes one degree of freedom and allows the state to be parameterized by three real numbers as

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

The parameter γ is an irrelevant degree of freedom that does not impact any observable. Since the global phase associated with it is not detectable and lacks physical significance, it can be promptly ignored and removed from our parameterization altogether, leaving

$$|\psi\rangle = \cos\left(\frac{\theta}{2}\right) |0\rangle + e^{i\varphi} \sin\left(\frac{\theta}{2}\right) |1\rangle. \quad (2.4)$$

With two degrees of freedom, this parameterized expression can now be visualized as lying on the surface of a sphere. The Bloch sphere depicted in figure 1 is a popular and convenient geometric representation of the state of a (pure) qubit.




Figure 1: Representation of the state of a qubit in the Bloch sphere. The angles θ and φ correspond to the parameters from equation 2.4.

A qubit can be acted on by *operators* that are represented by two-by-two matrices. There are three matrices of particular importance, the *Pauli matrices*:

$$X \equiv \sigma_x \equiv \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y \equiv \sigma_y \equiv \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad Z \equiv \sigma_z \equiv \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \quad (2.5)$$

Frequently one considers the identity matrix I as a Pauli matrix itself, oftentimes denoted σ_0 . All Pauli matrices are unitary ($\sigma_i \sigma_i^\dagger = \sigma_i^\dagger \sigma_i = I$) and Hermitian ($\sigma_i = \sigma_i^\dagger$).

These four matrices are orthogonal to each other, and together they span the space of the *observables* of a qubit. Observables are named so because they are associated with *observable physical quantities*. Each projective measurement is associated with an observable, whose eigenvalues correspond to the possible outcomes of the measurement. These outcomes are always real-valued, which is related to the fact that all observables are Hermitian.

From definitions 2.2 and 2.5 we can see that the computational basis states are *eigenstates* of Z . Computational basis *measurements* are projective measurements associated with the observable Z , and their possible outcomes correspond to the eigenvalues of the operator. These are the measurements typically available in quantum computers.

2.1.1.2 Pure Multiple Qubit States

All these ideas generalize to larger systems: the tensor product can be used to obtain a larger vector space in which the state of the multiple qubits lives. The description grows rapidly: to specify the state of n qubits, we need 2^n complex amplitudes, that correspond to $2^{n+1} - 2$ real-valued degrees of freedom after removing those corresponding to the normalization condition and the global phase.

When we have more than one qubit, measurement outcomes can be correlated beyond what is possible in classical systems. This is because it is not always the case that an n qubit state can be written in the form

$$|\psi_0\rangle \otimes |\psi_1\rangle \otimes \dots \otimes |\psi_i\rangle \otimes \dots \otimes |\psi_{n-2}\rangle \otimes |\psi_{n-1}\rangle, \quad (2.6)$$

where the ψ_i are the states of the individual qubits. When this is in fact possible, the state is called *separable*; otherwise, we say that it is *entangled*. Entanglement plays a big role in quantum algorithms: it is an exclusively quantum-mechanical resource and a bedrock for quantum speedup.

For two qubits, the maximally entangled states are the famous *Bell states*:

$$\begin{aligned} |\psi^+\rangle &= \frac{1}{\sqrt{2}} |01\rangle + |10\rangle, & |\psi^-\rangle &= \frac{1}{\sqrt{2}} |01\rangle - |10\rangle, \\ |\phi^+\rangle &= \frac{1}{\sqrt{2}} |00\rangle + |11\rangle, & |\phi^-\rangle &= \frac{1}{\sqrt{2}} |00\rangle - |11\rangle. \end{aligned} \quad (2.7)$$

It can be readily seen that the outcome of a (computational basis) measurement done on one of the qubits fully determines the outcome of a measurement on the other one.

2.1.1.3 Mixed States

Most of what was discussed so far only holds for *pure states* - those that are known exactly, and can be written in the form of equation 2.4. During a computation, one would ideally deal with pure states exclusively. Unfortunately, this is never the case: there are unknown errors occurring during the state preparation and the computation itself. The qubits are not fully isolated, as there is always some coupling with the environment. This results in the system being in a *mixed state*, a statistical distribution over pure states. This mixture simply translates lack of knowledge about the system.

States that are not pure cannot be written as in equation 2.4 and cannot be represented in the surface of the Bloch sphere. The *density matrix* ρ provides a way to describe a system in a more general state as

$$\rho = \sum_i p_i |\psi_i\rangle \langle \psi_i|. \quad (2.8)$$

Here the p_i are the probabilities of occurrence of the respective $|\psi_i\rangle$, the possible pure state configurations. The density matrix offers a way to test the purity of the state:

$$\begin{aligned} \text{tr}(\rho^2) = 1 &\longrightarrow \text{state is pure;} \\ \text{tr}(\rho^2) < 1 &\longrightarrow \text{state is mixed.} \end{aligned} \quad (2.9)$$

There is also a way of generalizing the geometrical representation of the states: the Bloch *ball*. We rewrite the density matrix representing the state as

$$\rho = \frac{I + \vec{r} \cdot \vec{\sigma}}{2}, \quad (2.10)$$

where $\vec{\sigma} = (X, Y, Z)$. The three-dimensional vector \vec{r} , with norm lower than or equal to unity, is the Bloch vector whose coordinates are written in 2.11.

$$\left(\sum_i p_i x_i, \sum_i p_i y_i, \sum_i p_i z_i \right) \quad (2.11)$$

The sum is over the possibilities of pure states, with p_i their probabilities and (x_i, y_i, z_i) their coordinates on the Bloch sphere. It is easy to verify that, for the special case of pure states, the norm of the Bloch vector \vec{r} is 1 and the coordinates are as before. For mixed states, the norm is in $[0, 1[$: they are *inside* the Bloch sphere, in the Bloch *ball*. The lower the purity, the lower the norm of \vec{r} . The totally mixed state lies in the center of the Bloch sphere, representing a complete lack of knowledge about the system.

2.1.2 Measuring Pauli Strings

As it was mentioned before, one can perform a Z measurement on a qubit. This yields either of the operator's eigenvalues, with the respective probabilities depending on the state of the qubit. Repeated measurements allow obtaining the expectation value of this operator in this specific state, namely

$$\langle \psi | Z | \psi \rangle.$$

It is common for problems to require measurements in bases other than the computational basis; this can be done with *basis rotations*. Measuring a generic observable amounts to performing a computational basis measurement, preceded by the unitary that rotates from the eigenbasis of the desired observable to the computational basis.

Evidently, in order to be appropriate for an observable A , the unitary U needs to obey the condition


$$\langle \psi | U^\dagger Z U | \psi \rangle = \langle \psi | A | \psi \rangle$$

for all states $|\psi\rangle$, so that the requirement can be written simply as


$$U^\dagger Z U = A. \quad (2.12)$$

This generalizes for observables acting on multiple qubits.


Implementing the basis rotation unitary for a generic observable may prove to hinder efficiency. Fortunately, *Pauli measurements* in particular are very simple to do.



(a) Circuit to measure Z .



(b) Circuit to measure X .



(c) Circuit to measure Y .

Figure 2: Examples of circuits that perform single-qubit Pauli measurements, drawn in the IBM Quantum Composer [41].

Circuits for all three types of single-qubit Pauli measurements (Z, X, Y) are drawn in figure 2. The circuit in 2a implements the typical computational basis measurement: no basis rotation is necessary. For X we can use the unitary $U = H$ as the basis rotation - this is the popular Hadamard gate. The circuit is represented in figure 2b. Finally, figure 2c implements a measurement of Y , with the basis rotation being a rotation around the x axis $U = \text{Rx}(\pi/2)$.

These choices of unitaries are not unique. As an example, using $U = HS^\dagger$ for measuring Y is also a popular option. One can choose the rotation that is the most convenient, given the operations that can natively be implemented in the hardware.

As well as single qubit Pauli measurements, one can easily measure *Pauli strings*. Pauli strings are observables of the form

$$\bigotimes \sigma_i \quad i \in \{0, z, x, y\}. \quad (2.13)$$

There are multiple ways to perform such measurements. The simplest and most evident one amounts to performing the individual Pauli measurements and multiplying the outcomes. However, it is possible to include two qubit gates in addition to the single qubit rotations as a means of reducing the number of measurements that needs to be performed.

As an example, two versions of circuits to perform a $Z \otimes Z$ measurement are illustrated in figure 3.



(a) Using two single-qubit measurements.

(b) Using one single-qubit measurement.

Figure 3: Examples of circuits that perform a $Z \otimes Z$ measurement, drawn in the IBM Quantum Composer [41].

The circuit in 3a measures both qubits in the computational basis; the result is then obtained by multiplying the outcomes. The final value will be in $\{-1, 1\}$. The actual value depends on the *parity* of the computational basis state that was measured: we obtain 1 if the parity was even, meaning that the qubits were measured on the same state ($|00\rangle$ or $|11\rangle$). The outcome is -1 if the parity was odd, meaning that they were measured on opposite states ($|01\rangle$ or $|10\rangle$).

In contrast, the circuit in 3b performs only one measurement. A **CNOT** gate is used to compute and store the parity of the state of both qubits in the state of the first one alone.

Essentially, for a two-qubit observable A , we can use an unitary that obeys

$$U^\dagger(Z \otimes Z)U = A$$

and measure both qubits, or use an unitary that obeys

$$U^\dagger(Z \otimes I)U = A$$

and measure only one of them.

With this, it has become clear that any Pauli string, i.e. operator of the form of equation 2.13, can be measured in a quantum computer. This is of great importance, because these Pauli strings span the space of Hermitian operators, and all observables are Hermitian. As a consequence, all observables can be written as a linear combination of Pauli strings; and given the linearity of quantum mechanics, all *expectation values* of observables can be obtained as a weighed sum of expectation values of Pauli strings.

If we want to measure an observable \hat{O} , we can start by decomposing it in the form of equation 2.14.

$$\hat{O} = \sum_i h_i \hat{P}_i \quad (2.14)$$

Here the h_i are real coefficients and the \hat{P}_i are Pauli strings (i.e. observables of the form of 2.13). It is then straightforward to evaluate the desired expectation value: it amounts to measuring the expectation value of each Pauli string appearing in the decomposition of the observable, and plugging it into formula 2.15.

$$\langle \hat{O} \rangle = \sum_i h_i \langle \hat{P}_i \rangle \quad (2.15)$$

This procedure can be followed for any observable; however, it is only tractable if the observable can be decomposed into a sum of polynomially-many Pauli strings.

2.1.3 Quantum Simulation and Trotterization

The behaviour of an isolated quantum-mechanical system is dictated by the *Schrödinger equation* in 2.16.

$$i\hbar \frac{d}{dt} |\psi\rangle = H |\psi\rangle \quad (2.16)$$

When the Hamiltonian operator for the system, H , is time-independent, this implies that the time evolution of the wave function has the form presented in 2.17.

$$|\psi(t)\rangle = e^{-iHt} |\psi(0)\rangle \quad (2.17)$$

As it was mentioned before, the simulation of physical systems is one of the most promising applications of quantum computers, as it is generally a hard task for classical computers. As such, it is important to know how to simulate this time evolution on a quantum computer.

A general Hamiltonian is not easy to exponentiate - so the question becomes, *how do we approximate the time evolution with sufficient accuracy?*


Typically, one deals with special classes of Hamiltonians that allow easy and efficient exponentiation.

A Hamiltonian can be decomposed into a sum of several, simpler terms:


$$H = \sum_{k=1}^L H_k. \quad (2.18)$$

And the H_k may allow easier implementation on a quantum computer - for example, each one may act only on a small subsystem, or they may have a special form.

In fact, as was discussed in subsection 2.1.2, we know that any Hermitian operator (such as the Hamiltonian) can be decomposed into a linear combination of Pauli strings. Conveniently, when the H_k are Pauli strings with real coefficients, it is easy to construct the circuit that applies $e^{-iH_k t}$.



(a) Circuit for simulating the time evolution of a system under the Hamiltonian $H_k = Z \otimes Z \otimes Z \otimes Z$.



(b) Circuit for simulating the time evolution of a system under the Hamiltonian $H_k = X \otimes Y \otimes Z \otimes X$.

Figure 4: Examples of circuits that apply time evolution under simple Hamiltonians H_k , consisting of Pauli strings. They apply $e^{-iH_k t}$ using a parameterized Z rotation. The circuits were drawn in the IBM Quantum Composer [41].

Figure 4 shows two circuits that can be used to implement the exponentiation of two specific H_k .

The circuit in figure 4a is the simplest one: it implements $e^{-itZ \otimes Z \otimes Z \otimes Z}$. This amounts to applying a local phase e^{-it} to the computational basis states that have an even parity, and e^{it} to those that have an odd parity; the drawn circuit does precisely that. The first ladder of **CNOT** gates computes the parity of the state and stores it on the last qubit. Then, all that's left is to apply e^{-itZ} to the qubit that holds the parity, which can be done by rotating it around the Z axis by an angle of $2t$ (i.e. applying the gate $Rz(2t)$). Finally, a ladder of **CNOT** gates is applied to uncompute the parity before the computation proceeds.

In figure 4b, we have the circuit that implements $e^{-itX \otimes Y \otimes Z \otimes X}$. This can be done with minor adjustments to the previous circuit.

Appendix A shows that, when H_k is a Pauli string, one can write any operator of the form e^{itH_k} as

$$e^{-it \otimes_i P_i} = \bigotimes_i U_i^\dagger \left(e^{-it \otimes_i Z_i} \right) \bigotimes_i U_i, \quad (2.19)$$

where the U_i are single qubit unitaries obeying

$$U_i^\dagger Z_i U_i = P_i.$$

For any qubit i , acted on by P_i in H_k .

This is the same condition as 2.12: as we did to make Pauli measurements, we need an unitary that rotates from the eigenbasis of the operator in question to the computational basis.

The form of equation 2.19 is very convenient: it tells us that we can use the circuit for $e^{-it \otimes_i Z_i}$ (figure 4a) for implementing any operator of the form $e^{-it \otimes_i P_i}$ (where $P_i \in \{I_i, Z_i, X_i, Y_i\}$), as long as we apply the appropriate $\bigotimes_i U_i$ before and $\bigotimes_i U_i^\dagger$ after.

The circuit in figure 4b does precisely that, for the specific case of the Pauli string $\bigotimes_i P_i = X \otimes Y \otimes Z \otimes X$. The employed basis rotations were the same as those from section 2.1.2.

So now we know how to simulate the time evolution under simple Hamiltonians H_k consisting of a single Pauli string. However, the end goal is simulating the full Hamiltonian H from 2.18. This amounts to applying $e^{-it \sum_{k=1}^L H_k}$, but

$$e^{-it \sum_{k=1}^L H_k} \neq \prod_{k=1}^L e^{-it H_k}$$

in general. In fact, for two operators H_i, H_j , we can calculate $e^{H_i} e^{H_j}$ as

$$e^{H_i} e^{H_j} = e^{H_i + H_j + \frac{1}{2} [H_i, H_j] + \frac{1}{12} [H_i, [H_i, H_j]] - \frac{1}{12} [H_j, [H_i, H_j]] + \dots} \quad (2.20)$$

Equation 2.20 is known as the **Baker–Campbell–Hausdorff (BCH)** formula. We can see that the product of the exponential series is the exponential series of the sums only in the special case of commuting operators ($[H_i, H_j] = 0$).

Fortunately, we can use the *Trotter formula*:

$$\lim_{n \rightarrow \infty} (e^{iH_i t/n} e^{iH_j t/n})^n = e^{i(H_i + H_j)t}. \quad (2.21)$$

Evidently, it is not feasible to implement the circuit for $n \rightarrow \infty$. However, we can make n finite to obtain a circuit that *approximates* the desired time evolution.

$$e^{i(H_i + H_j)t} \approx (e^{iH_i t/n} e^{iH_j t/n})^n \quad (2.22)$$

Here, n is the number of *repetitions* of the basic unit of our Trotter circuit, and t/n is the *time step*. We can interpret equation 2.22 as approximating the evolution of a system under $(H_i + H_j)$ for a time t with n repetitions of the evolution under H_j followed by the evolution under H_i , both always for a time step of t/n . We switch between evolving the system under the two elements in the sum repeatedly, with the duration of this evolution being always less than the actual evolution time.

The larger the number of repetitions n , the shorter the time step and the better the approximation: in the limit of $n \rightarrow \infty$, one would get the exact result. The error in the approximation will depend on the evolution time; as an example, a Trotter approximation with a single repetition will have an error $O(t^2)$ [64].

2.1.4 Current Limitations

As discussed in the previous chapter, it will likely be a while before a fault-tolerant quantum computer (FTCQ) is available. Currently, one has to deal with noisy intermediate-scale quantum (NISQ) devices. These are not error corrected and come with a multitude of limitations, namely:

- Coherent noise.
- Incoherent noise.
- SPAM errors.

A brief analysis of each follows.

2.1.4.1 Coherent Noise

Coherent errors are unitary and don't undergo fast changes (relative to the gate time); this type of error has various causes, such as systematic control noise, global external fields, cross-talk, and unwanted qubit-qubit interactions [33].

For a single qubit, a coherent source of noise could be performing a rotation by an angle of $\theta + \delta\theta$ instead of just θ due to miscalibration. For two qubits, a well-known example is the unwanted ZZ interaction in superconducting quantum computers. It can significantly lower the fidelity of two-qubit gates; as such, it can have a great impact on the performance of the device, and consequently pose a challenge on the scalability of the architecture [100].

Coherent errors interfere constructively, and may take a significant toll on the computation. In the worst case scenario, the error rate of coherent noise can scale quadratically both in the number of qubits and in the length of the circuit [74, 44]. For that reason, it can accelerate the error accumulation rate and be a challenge in the context of error correction codes. Regardless, there is hope that error correction can be used to remove undesirable coherence from the noise channel, thus avoiding the possibility that constructive interference leads to a quadratic scaling of the error rate. For the toric code and a particular

coherent noise model, it was shown that the logical channel after error correction becomes increasingly incoherent with the length of the code [44].

2.1.4.2 Incoherent Noise

Unlike coherent noise, incoherent noise can compromise the purity of the state of the qubits, causing it to become mixed. Incoherent errors may arise for example from microscopic materials defects or unwanted interactions with the environment [50, 54]. Since it is impossible to have the qubits in perfect isolation and quantum mechanical systems are fragile, even *storing* quantum information is a difficult task.

While coherent errors are unitary and preserve the norm of the Bloch vector, incoherent errors can *shrink* the Bloch vector. This is a visual translation of the fact that we've lost information about the state the qubits are in; its purity has decreased. Decoherence can destroy superposition and entanglement and cause us to lose the quantum behaviour of the system, with the information becoming classical [68].

When the errors are incoherent, the worst case scenario error rate scales linearly with the circuit size (depth and qubits), an improvement against the quadratic scaling of coherent errors [74, 44].

2.1.4.3 State Preparation and Measurement Errors

State preparation and measurement errors are both classical errors, know collectively as **SPAM** errors.

A quantum computer needs to be able to prepare a fiducial initial state, such as the all-zero state; and in order to obtain results from the computation, measurements will be necessary as well. Initialization and measurement are thus two fundamental non-unitary operations that bridge quantum and classical information, and allow quantum computing to be of use in our classical world.

However, these operations can't be done perfectly either. The prepared state will likely not be exactly the one that was intended, and the measurement operation can be faulty as well.

As **SPAM** errors only occur at the beginning and end of the circuit, they evidently don't scale with the circuit depth, which makes them less disruptive. In spite of that, this is a relevant source of noise for near-term quantum computers, and efforts have been done to find efficient ways of partly mitigating it [27].

2.1.4.4 Other Challenges

As well as noise, **NISQ** computers come with severe limitations in qubit count and connectivity. The qubit requirements of the most promising applications are beyond what is available today; and limited connectivity demands extra routing operations or teleportation protocols, resulting in an overhead in ancillary qubits, execution time, and gate depth - each of which aggravates already critical resource requirements.

An additional challenge is that there is always some error when using a quantum computer to evaluate expectation values. This is not limited to **NISQ** computers, it is an intrinsic limitation rooted in the nature of quantum mechanics. The error is associated with quantum projection noise and the always finite number of shots, and it will always exist, even if its magnitude can be decreased. In algorithms that include

optimizing an expectation value, attention must be paid to this unavoidable source of noise that hampers the task of the optimizer.

2.1.4.5 Outlook

The performance of the quantum computers available today greatly suffers from coherent and incoherent noise, SPAM noise and the always-present shot noise.

Even if making a strong simulation of the quantum computer is hard for classical computers, producing a better approximate solution to a problem may not be so. As such, all sources of noise that contribute to corrupting the output of the computation may inhibit quantum advantage.

Errors are by no means exclusive to quantum computers; classical computations also suffer from them. Regardless, a relevant difference resides in the fact that classical information is much easier to protect and manipulate than quantum information, and error correction itself is a simpler task to perform on classical data. The viability of classical computations had been proven by as soon as 1956, with a key result by von Neumann showing that classical computers could be made robust to noise [63].

The noise in quantum computers, along with the increased difficulty of error correction that arises in dealing with quantum data, brought concern that noise robustness was simply unattainable in quantum computers [52, 86]. Optimism was brought by the introduction of several error correction codes [77, 81], followed by several proofs (comprising several error models and error correction codes) of an analogue of the result obtained by von Neumann for classical computers: the *quantum threshold theorem*.

According to the quantum threshold theorem, once the physical error rate becomes smaller than a certain threshold, quantum computations can be made arbitrarily accurate by scaling up the error correction code [15, 51, 1, 7]. This provides a path to fault-tolerant quantum computing: once the hardware error rate is under the accuracy threshold for a given quantum error correction code, fault-tolerance has been achieved.

2.2 Quantum Chemistry

This section aims at providing a light review of the aspects of quantum chemistry essential for the contents of the dissertation. The book used as a reference, Modern Quantum Chemistry [84], offers a broader and more in-depth view of the field.

2.2.1 The Electronic Problem

We are interested in studying electronic structure and the quantum mechanical *motion* of molecular systems. The objective is to understand the stable configurations of the system and its dynamics, by considering the interactions between the elements it is comprised of.

A scheme of a simple molecule, at the level of electrons and nuclei, is presented in figure 5.




Figure 5: Representation of a molecule in a coordinate system. Upper case letters (A, B) represent nuclei, while lower case ones (i, j) represent the electrons. Similarly, upper case R represent position vectors pointing to nuclei, while lower case r represent those pointing to electrons. A single subscript is used on the position vector when the vectors point from the origin, while two subscripts are used to represent relative position vectors. In the latter case, the first and second subscripts indicate where the vector points to and from, respectively (e.g. $r_{ij} = r_i - r_j$).

We are interested in finding (approximate) solutions to equation 2.23, the non-relativistic time-independent Schrödinger equation.

$$\hat{H} |\psi\rangle = E |\psi\rangle \quad (2.23)$$

This is an *eigenvalue problem*.

Here \hat{H} is the Hamiltonian operator for our molecular system. It is comprised of five terms; a brief discrimination of each follows. Atomic units are to be used throughout the whole section, so that the electronic mass and the reduced Planck constant \hbar are both equal to unity. The energy will then be in Hartrees (1a.u. = 27.2eV). N and M stand respectively for the number of electrons and the number of nuclei. The notation regarding position vectors can be understood from figure 5.

2.2.1.1 Kinetic Energy of the Electrons

The electronic kinetic energy is the sum of the kinetic energies of the individual electrons,

$$- \sum_{i=1}^N \frac{1}{2} \nabla_i^2.$$

The simple multiplicative factor is due to the fact that atomic units were used. Here ∇_i^2 is the *Laplacian operator*, involving the second partial derivatives with respect to the coordinates of the electron labeled i :

$$\nabla_i^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}.$$

This term represents the energy associated with the motion of all the electrons in the molecule.

2.2.1.2 Kinetic Energy of the Nuclei

The term associated with the motion of the nuclei, written in 2.24, is similar to the previous one.

$$- \sum_{A=1}^M \frac{1}{2M_A} \nabla_A^2 \quad (2.24)$$

The only difference is that this term comes weighed by a factor of $\frac{1}{M_A}$, of the order of 10^4 . This factor, defined in 2.25, simply accounts for the fact that the mass of the nucleus is different from the mass of an electron.

$$M_A = \frac{\text{mass of the nucleus } A}{\text{mass of an electron}} \quad (2.25)$$

2.2.1.3 Electron-Nucleus Attraction

What we have left now is the potential energy terms of the Hamiltonian, which include a negative term (attraction) and two positive terms (repulsion).

The attraction term represents the interaction between molecules of opposite charge, as per Coulomb's law; it reads

$$- \sum_{i=1}^N \sum_{A=1}^M \frac{Z_A}{r_{iA}}. \quad (2.26)$$

Here, Z_A is the atomic number of the nucleus A . The greater the number, the greater the charge, and the greater the attraction; the larger the distance between the nucleus and the electron r_{iA} , the smaller the attraction.

2.2.1.4 Electron-Electron Repulsion

This is a positive term, arising from the repulsion between particles of identically negative charge. The corresponding expression is

$$\sum_{i=1}^N \sum_{j>i}^N \frac{1}{r_{ij}}. \quad (2.27)$$

Since electrons have unit charge, the only factor here is the distance between them (r_{ij}). The double sum runs over all distinct pairs of electrons.

2.2.1.5 Nucleus-Nucleus Repulsion

Finally, we have the positive term accounting for the repulsion between positively-charged particles:

$$\sum_{A=1}^M \sum_{B>A}^M \frac{Z_A Z_B}{R_{AB}}. \quad (2.28)$$

2.2.1.6 Complete Hamiltonian and the Born-Oppenheimer Approximation

Contemplating all terms, the Hamiltonian of the system reads

$$\hat{H} = - \sum_{i=1}^N \frac{1}{2} \nabla_i^2 - \sum_{A=1}^M \frac{1}{2M_A} \nabla_A^2 - \sum_{i=1}^N \sum_{A=1}^M \frac{Z_A}{r_{iA}} + \sum_{i=1}^N \sum_{j>i}^N \frac{1}{r_{ij}} + \sum_{A=1}^M \sum_{B>A}^M \frac{Z_A Z_B}{R_{AB}}. \quad (2.29)$$

Unfortunately, solving the Schrödinger equation 2.23 with the Hamiltonian in 2.29 is not an easy task - the computational complexity grows quickly with the size of the molecule. Obtaining an exact solution is only possible in simple cases.

To address this issue, Born and Oppenheimer proposed an approximation in 1927 [14]. The idea behind the proposal was that, since the mass of the nuclei is significantly larger than that of the electrons, their motion will be much slower for a comparable kinetic energy. Thus, considering the nuclei to be stationary seems like a reasonable approximation.

The Born-Oppenheimer approximation then regards the motion of the electrons as occurring in a field created by the *fixed* nuclei. In this scenario, because they are not moving ($\nabla_A^2 = 0$), the nuclei have null kinetic energy, so that the term 2.24 in the Hamiltonian vanishes. Further, the nucleus-nucleus repulsion 2.28 becomes constant, because their relative position R_{AB} is not changing. A constant term represents only a constant shift in the eigenvalues, and does not affect the eigenstates; as such, we can ignore it as we are solving the problem, and add it back later.

Now we have

$$\hat{H}_e = - \sum_{i=1}^N \frac{1}{2} \nabla_i^2 - \sum_{i=1}^N \sum_{A=1}^M \frac{Z_A}{r_{iA}} + \sum_{i=1}^N \sum_{j>i}^N \frac{1}{r_{ij}}. \quad (2.30)$$

We call the Hamiltonian in 2.30 *electronic* Hamiltonian for obvious reasons: it now only concerns the motion of the electrons. The nuclear coordinates may affect what the solution to the Schrödinger equation looks like, but the solution itself will be an electronic wave function, that depends explicitly *only* on the electronic coordinates.

The approximation allows treating the electrons and nuclei as two decoupled systems, which greatly simplifies the problem.

2.2.2 Finding the Ground State: Classical Approaches

There is a fundamental principle in quantum mechanics which is widely used in computational chemistry, be the resources classical or quantum: the *variational principle*. This principle states that the expectation value of the Hamiltonian is lower bounded by the eigenvalue corresponding to the ground energy (E_0), i.e.

$$\langle \psi | \hat{H} | \psi \rangle \geq E_0, \quad (2.31)$$

assuming that the state is normalized ($\langle \psi | \psi \rangle = 1$). A brief sketch of the proof follows.

We consider $\{|\phi_\alpha\rangle\}$ to be the set of eigenstates of the Hamiltonian operator \hat{H} and denote the eigenvalue of $|\phi_\alpha\rangle$ by E_α , i.e.

$$\hat{H} |\phi_\alpha\rangle = E_\alpha |\phi_\alpha\rangle, \quad \alpha = 0, 1, 2, \dots$$

The eigenvalues are assumed to be ordered ($\alpha < \beta \implies E_\alpha \leq E_\beta$). Since the operator \hat{H} is Hermitian, the eigenvalues are real ($E_\alpha \in \mathbb{R} \ \forall \alpha$) and the eigenstates are orthonormal ($\langle \phi_\alpha | \phi_\beta \rangle = \delta_{\alpha,\beta}$).

We assume that this set of eigenstates is complete, so that we can write any state $|\psi\rangle$ in the eigenbasis of \hat{H} :

$$|\psi\rangle = \sum_{\alpha} c_{\alpha} |\phi_{\alpha}\rangle$$

With this, the expectation value of \hat{H} reads (using $*$ to denote complex conjugation)

$$\begin{aligned} \langle \psi | \hat{H} | \psi \rangle &= \sum_{\alpha} c_{\alpha}^* \langle \phi_{\alpha} | \hat{H} \sum_{\beta} c_{\beta} |\phi_{\beta}\rangle = \sum_{\alpha,\beta} c_{\alpha}^* c_{\beta} \langle \phi_{\alpha} | \hat{H} | \phi_{\beta} \rangle \\ &= \sum_{\alpha,\beta} c_{\alpha}^* c_{\beta} E_{\beta} \langle \phi_{\alpha} | \phi_{\beta} \rangle = \sum_{\alpha,\beta} c_{\alpha}^* c_{\beta} E_{\beta} \delta_{\alpha,\beta} = \sum_{\alpha} c_{\alpha}^* c_{\alpha} E_{\alpha} = \sum_{\alpha} |c_{\alpha}|^2 E_{\alpha}. \end{aligned}$$

It was previously stated that the eigenvalues are ordered, so that $E_{\alpha} \geq E_0 \ \forall \alpha$; further, assuming that the state $|\psi\rangle$ is normalized, we have $\sum_{\alpha} |c_{\alpha}|^2 = 1$. With this we can conclude the proof:

$$\langle \psi | \hat{H} | \psi \rangle = \sum_{\alpha} |c_{\alpha}|^2 E_{\alpha} \geq \sum_{\alpha} |c_{\alpha}|^2 E_0 = E_0$$

Typically, the methods for finding an approximate solution to the ground state of the electronic system aim to minimize $\langle \psi | \hat{H} | \psi \rangle$. They differ only in the *variational form* of $|\psi\rangle$, a concept identical to that of an *ansatz*. It is a guess that fixes the structure of the wave function, but only specifies it up to some tunable *variational parameters*. Plugging some test parameters into the variational form, we obtain a fully specified wave function in which we can calculate the expectation value of \hat{H} . The task is then reduced to finding the best parameters: given the variational principle, this is just a minimization problem.

A more thorough overview of the variational principle can be found in reference [34].

2.2.2.1 Self-Consistent Mean Field Methods

As much as the Born-Oppenheimer approximation greatly simplifies our eigenproblem, solving the Schrödinger equation with the electronic Hamiltonian (equation 2.30) exactly is still out of reach. The problematic term is now the electron-electron repulsion (expression 2.27), which couples the coordinates of *all* the electrons. Finding the exact solution implies dealing simultaneously with the coordinates of them all; decoupling these variables would allow solving the equation. Such is made possible by the *mean-field approximation*.

The essence of the mean-field approximation resides in the simplification of the effect of the electrons in each other: instead of contemplating the sum of all the individual repulsion terms, we assume that each electron experiences an *average* potential created by all the other electrons. Each one of them is considered to be moving in the *mean field* created by the others.

This gives rise to the operator in equation 2.32.

$$\hat{f}(i) = -\frac{1}{2}\nabla_i^2 - \sum_{A=1}^M \frac{Z_A}{r_{iA}} + v(i) \quad (2.32)$$

This operator is just the electronic Hamiltonian in equation 2.30, separated for each electron i . The impact of the remaining electrons is contained within $v(i)$, an average potential.

At this point we have managed to separate the Schrödinger equation into single fermion equations: instead of one complicated many-body problem, we face multiple simple one-body problems:

$$\hat{f}(i)\chi(\vec{x}_i) = \epsilon\chi(\vec{x}_i). \quad (2.33)$$

This is another eigenvalue equation, thus its form is identical to that of equation 2.23. But it is much simpler: we now have $\hat{f}(i)$, a one-body operator, instead of the many-body operator \hat{H} . Instead of the many-body wave function $|\psi\rangle$, depending on the coordinates of all of the electrons, we have single-body wave functions (spin-orbitals) $\chi(\vec{x}_i)$, each of which depends solely on the coordinates of one of them.

If we have N electrons, we have N equations of the form of 2.33. We want to find the eigenvalues ϵ and the eigenfunctions $\chi(\vec{x}_i)$ of each operator $\hat{f}(i)$. However, this set of equations is non-linear: for the i th electron this operator depends (through $v(i)$) on the wave functions of the other electrons $\chi(\vec{x}_j)$, which are the solutions of all the other $N-1$ equations. Thus, the solution must be obtained *iteratively*.

The iterative method for solving the equation 2.33 is the **Self-Consistent Field (SCF)** method. This is because the convergence condition is the self-consistency of the field.

The *variational parameters* are contained within the single-particle wave functions, the $\chi(\vec{x}_i)$. We start with an *initial guess* for the values of these parameters. With a concrete set of single-particle wave functions, we are then in place to calculate the average potential that each particle i feels ($v(i)$). Plugging the values in the expressions of the operator $\hat{f}(i)$, we get N single-particle equations of the form of 2.33. These are much simpler to solve than the full Schrödinger equation. The obtained solutions are new single-particle wave functions $\chi(\vec{x}_i)$. From these we can calculate new average potentials $v(i)$, and solve for the eigenfunctions of the operator $\hat{f}(i)$ once again. The procedure is repeated until self-consistency is achieved, i.e. the fields $v(i)$ stop changing and the trial functions $\chi(\vec{x}_i)$ used to calculate those fields are also the solutions to the equations written with them. At this point, the trial functions used to construct the operator $\hat{f}(i)$ are simultaneously its eigenvalues: the field is self-consistent.

2.2.2.2 Variational Form: Hartree and Hartree-Fock Methods

The previous discussion purposefully glossed over what the many-body wave functions look like, and how they are related to the one-body wave functions. While the essence of SCF methods is similar, they differ significantly in this one key aspect: the *ansatz*, i.e. the variational form we choose for the wave functions.

It seems like the simplest approach would be to consider that the many-body wave functions are just products of one-body wave functions:

$$\psi(\vec{x}_1, \vec{x}_2, \dots, \vec{x}_{N-1}, \vec{x}_N) = \chi_\alpha(\vec{x}_1)\chi_\beta(\vec{x}_2)\dots\chi_\pi(\vec{x}_{N-1})\chi_\rho(\vec{x}_N). \quad (2.34)$$

This is what Hartree proposed upon introducing an SCF method. However, the Hartree method does not provide satisfactory results, because it does not account for a fundamental principle: *the antisymmetry principle*.

The antisymmetry principle is a crucial axiom in quantum mechanics. It states that the many-body fermionic wave function must be antisymmetric under exchange of any two particles, which we can write as

$$\psi(\vec{x}_1, \dots, \vec{x}_i, \dots, \vec{x}_j, \dots, \vec{x}_N) = -\psi(\vec{x}_1, \dots, \vec{x}_j, \dots, \vec{x}_i, \dots, \vec{x}_N). \quad (2.35)$$

Evidently, the Pauli exclusion principle is enforced by the antisymmetry condition 2.35. The condition is rooted in the *indistinguishability* of particles. Exchanging indistinguishable, or *identical*, particles must result in a physically equivalent state. Such only allows variability up to a global phase. There are two main types of indistinguishable particles: fermions and bosons, with complex phase factors $e^{i\pi} = -1$ and $e^{i2\pi} = 1$ respectively. These factors give rise to the antisymmetric nature of fermionic wave function, and the symmetric nature of the bosonic wave function.

In either case, the simplest possible way of combining the individual wave functions, the Hartree product (definition 2.34), is inadequate. In the scenario we are concerned with (the fermionic one), this can be dealt with by using *antisymmetrized products* instead.

$$\psi(\vec{x}_1, \vec{x}_2, \dots, \vec{x}_{N-1}, \vec{x}_N) = \begin{vmatrix} \chi_1(\vec{x}_1) & \chi_2(\vec{x}_1) & \dots & \chi_{N-1}(\vec{x}_1) & \chi_N(\vec{x}_1) \\ \chi_1(\vec{x}_2) & \chi_2(\vec{x}_2) & \dots & \chi_{N-1}(\vec{x}_2) & \chi_N(\vec{x}_2) \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \chi_1(\vec{x}_{N-1}) & \chi_2(\vec{x}_{N-1}) & \dots & \chi_{N-1}(\vec{x}_{N-1}) & \chi_N(\vec{x}_{N-1}) \\ \chi_1(\vec{x}_N) & \chi_2(\vec{x}_N) & \dots & \chi_{N-1}(\vec{x}_N) & \chi_N(\vec{x}_N) \end{vmatrix} \quad (2.36)$$

Wave functions of the form of 2.36 are typically called *Slater determinants*. Matrix determinants are antisymmetric under exchange of any two rows or columns, making them a natural choice here. For bosons, one would use matrix *permanents*.

Slater determinant ansätze were used in the *Hartree-Fock* method, an updated version of the Hartree method that accounted for the antisymmetry principle.

As it was mentioned before, the choice of ansatz is what shapes a specific *SCF* method. The form of the operator $\hat{f}(i)$ as written in 2.32, and its place in the algorithm, are common to several *SCF* methods. The difference is in the potential $v(i)$: it will be *specific* to the assumptions we make about the wave function, i.e. it depends on how our ansatz looks like. This operator is then used to write Schrödinger-like one-particle equations 2.33, which means that the chosen ansatz greatly impacts the formulation of our eigenproblems.

In the Hartree-Fock method, the operator of the form of equation 2.32 is called the *Fock operator*:

$$\hat{f}(i) = -\frac{1}{2}\nabla_i^2 - \sum_{A=1}^M \frac{Z_A}{r_{iA}} + v^{HF}(i) \quad (2.37)$$

The key difference is that the $v^{HF}(i)$ term here has a superscript, referring to the fact that the potential in this operator is the *Hartree-Fock* potential. The eigenvalue equations 2.33 defined through the Fock operator are the *Hartree-Fock* equations.

As an illustration of how the potential changes depending on the ansatz, a shallow, simplified comparison of the Hartree and Hartree-Fock potentials follows.

The method proposed by Hartree used the Hartree product as the variational form. This was associated with a simple potential, consisting of a local operator that accounted solely for the electron-electron repulsion. In contrast, the potential in Hartree-Fock method includes, in addition to this local operator, an *exchange* operator. This is a non-local operator that simply results from choosing Slater determinants as wave functions: it is an artifact associated with the imposition of the anti-symmetry principle.

2.2.2.3 Limitations of the Hartree-Fock Method

The Hartree-Fock method simplifies the problem in multiple ways. Demanding that the solution is a Slater determinant is in itself a restriction that precludes exactness, as is typical of variational methods: the optimal solution is only guaranteed to be the best *within the variational form*. What the Hartree-Fock method does is to find the best solution of Slater determinantal form.

As was discussed before, the motion of the nuclei is neglected (Born-Oppenheimer approximation). Further, the method also assumes the mean-field approximation, which fails to properly account for electron correlation outside of antisymmetry.

Finally, even if none of these approximations were used, the exact solution would still be beyond reach, because for the method to be tractable the basis set must be finite. For the Slater determinants to constitute a complete N-electron basis, they would have to be built from an infinite basis set of spin-orbitals.

Regardless, the Hartree-Fock method is a cornerstone in quantum chemistry, and its approximate solution is often used as a starting point in more accurate methods, both in classical and quantum computing.

2.2.2.4 Post-Hartree-Fock Methods

Given the shortcomings of the Hartree-Fock solution, a multitude of more accurate methods has been proposed to improve upon the Hartree-Fock approximation (at the expense of extra computational cost).

The following abbreviated explanation of post-Hartree-Fock methods is based, not only on the main reference for this section ([84]), but also on references [71] and [38].

Several relevant simplifications assumed by the Hartree-Fock method were mentioned previously. Of these, one typically undone by posterior methods is the assumption that the solution is of the form of a Slater determinant.

To correct for the electronic correlation ignored by the Hartree-Fock method, a typical approach is to expand the wave function as a *linear combination of* Slater determinants, rather than a single one of them. Since the Hartree-Fock solution is a good starting point, the wave function is often described in terms of *excitation operators* that will act on this state.

It is useful to define the *cluster operator* T :

$$\begin{aligned}
 T &= \sum_{i=1}^N T_i, \\
 T_1 &= \sum_{i,a} t_a^i a_a^\dagger a_i \\
 T_2 &= \sum_{i>j, a>b} t_{ab}^{ij} a_a^\dagger a_b^\dagger a_i a_j \\
 &\vdots
 \end{aligned}
 \tag{2.38}$$

The lower case t are expansion coefficients. The indices i, j and a, b run over orbitals that are occupied and unoccupied (virtual) in the reference state, respectively. The operator a_k^\dagger is a *creation* operator, which creates a fermion in orbital k ; a_k is an *annihilation* operator, which removes a fermion from orbital k . These operators belong to the second quantization formalism.

The cluster operator consists of a sum of all possible excitation operators T_i : the one that generates single excitations from the reference state (T_1), the one that generates double excitations (T_2), and so forth.

With this operator, it is possible to write the FCI variational form in 2.39.

$$|FCI\rangle = (1 + T) |HF\rangle \quad (2.39)$$

This is just the Hartree-Fock state acted on by an operator $(1 + T)$; the expansion coefficients will be the variational parameters. The term $T |HF\rangle$ consists of a linear combination of all possible excited determinants, since the cluster operator includes all possible excitation operators. The FCI state is then a linear combination of all Slater determinants in the N-electron space formed from the chosen set of spin-orbitals. The optimal FCI energy 2.40 will be the exact solution within the N-electron subspace spanned by these determinants. Of course, other than the finite basis set, the solution is also only exact up to the Born-Oppenheimer approximation.

$$E_{FCI} = \min_{\vec{t}} \frac{\langle FCI | H | FCI \rangle}{\langle FCI | FCI \rangle} \quad (2.40)$$

In short, the inclusion of more Slater determinants increases the variational freedom and improves the solution. The name *configuration interaction* arises from the fact that each Slater determinant in the expansion is associated with a specific electronic configuration (configuration of spin-orbitals). *Full* refers to the fact that all excitation operators are included.

Unfortunately, while FCI offers accurate results, it is only tractable for small molecules: the number of determinants that need to be included in the expansion grows with the factorial of the total number of spin orbitals. Consequently, even with minimal single electron basis sets, the difficulty of calculations increases rapidly as the size of the system increases.

To make the calculations tractable, the cluster operator T must be truncated.

$$T^{(k)} = \sum_{i=1}^k T_i \quad (2.41)$$

The operator $T^{(k)}$ defined in 2.41 is just the cluster operator from definition 2.38, truncated up to excitations of order k . When a truncated version of the operator is used, only a portion of the Slater determinants formed from the one-electron basis set will appear in the expansion; the associated method is simply called *Configuration Interaction (CI)*. The variational form is written as in 2.42.

$$|CI\rangle = (1 + T^{(k)}) |HF\rangle \quad (2.42)$$

For the specific case that k is two, the method is called [Configuration Interaction Singles and Doubles \(CISD\)](#), because it includes only single and double excitations from the reference state (variational form in [2.43](#)).

$$|CISD\rangle = (1 + T_1 + T_2) |HF\rangle \quad (2.43)$$

Truncating the operator to make the computation tractable results in a few issues. Namely, if we have two independent, non-interacting subsystems (e.g. two infinitely separated molecules), the [CI](#) energy of the full system will not be the sum of the energies of the subsystems. Because it lacks this property, it is said that [CI](#) is not *size-consistent*.

A size-consistent version can be created by means of exponentiation. The corresponding method is called [CC](#); the variational form is written in [2.44](#).

$$|CC\rangle = e^T |HF\rangle \quad (2.44)$$

Once again, the cluster operator T is typically truncated at some order of excitation. For the usual choice that excitations up to order two are included, the method is called [Coupled Cluster Singles and Doubles \(CCSD\)](#) (variational form in [2.45](#)).

$$|CCSD\rangle = e^{T_1+T_2} |HF\rangle \quad (2.45)$$

Even when truncated, [CC](#) methods don't suffer from the problem of size-inconsistency. Further, they have the interesting property that even when the cluster operator T is truncated up to order k , higher-order excitations occur in the exponential series. However, they have some problems of their own. A relevant weakness in [CC](#) is that the similarity-transformed Hamiltonian \bar{H} for conventional [CC](#) (definition [2.46](#)) is not variational.

$$\bar{H} = e^{-T^{(k)}} H e^{T^{(k)}} \quad (2.46)$$

This is due to the [CC](#) expectation value

$$E_{CC} = \langle HF | e^{-T^{(k)}} H e^{T^{(k)}} | HF \rangle \quad (2.47)$$

not being symmetric, since

$$\langle HF | e^{-T^{(k)}} \neq (e^{T^{(k)}} | HF \rangle)^\dagger,$$

which is just a consequence of the operator $e^{T^{(k)}}$ not being unitary. This means that we can't use the variational principle [2.31](#), because the associated expectation value isn't of the proper form. The asymmetry of the expectation value impedes the [CC](#) energy from being an upper bound to the ground energy.

To circumvent this, **Variational Coupled Cluster (VCC)** has been proposed. This method aims to minimize

$$E_{VCC} = \frac{\langle HF | e^{T^{(k)\dagger}} H e^{T^{(k)}} | HF \rangle}{\langle HF | e^{T^{(k)\dagger}} e^{T^{(k)}} | HF \rangle}. \quad (2.48)$$

Here, the denominator is not unity because of the non-unitarity of $e^{T^{(k)}}$. The **VCC** method is obviously variational, because the numerator is of the form $\langle \psi | H | \psi \rangle$; further, much like traditional **CC**, it is size-consistent. Unfortunately, the cost of the **VCC** method scales exponentially with the system size, regardless of the truncation order k .

Another approach to make a variational version of **VCC** is **UCC**, that replaces the operator e^T with another exponential operator, now unitary:

$$|UCC\rangle = e^{T-T^\dagger} |HF\rangle. \quad (2.49)$$

The Hermitian conjugate of the operator $(T - T^\dagger)$ reads

$$(T - T^\dagger)^\dagger = T^\dagger - T^{\dagger\dagger} = T^\dagger - T = -(T - T^\dagger),$$

which means that this operator is anti-Hermitian, and the commutators between $(T - T^\dagger)$ and $(T - T^\dagger)^\dagger$ are all zero. This grants unitarity to the **UCC** operator.

$$\left(e^{T-T^\dagger}\right) \left(e^{T-T^\dagger}\right)^\dagger = \left(e^{T-T^\dagger}\right)^\dagger \left(e^{T-T^\dagger}\right) = I$$

Since the **UCC** operator is unitary, the method is variational, and the **UCC** energy is an upper bound to the ground state energy. Finding the **UCC** ground state then amounts to a minimization problem (expression 2.50).

$$E_{UCC} = \min_{\vec{t}} \langle HF | e^{-(T-T^\dagger)} H e^{T-T^\dagger} | HF \rangle \quad (2.50)$$

As usual, the cluster operator should be truncated for the computation to be tractable, so that we get an **UCC** operator of the form $e^{T^{(k)}-T^{(k)\dagger}}$. For the frequent case that only excitations up to order two are included ($k=2$), the method is called **UCCSD**. This corresponds to the variational form in 2.51.

$$|UCCSD\rangle = e^{(T_1+T_2)-(T_1^\dagger+T_2^\dagger)} |HF\rangle \quad (2.51)$$

The **UCCSD** method is variational, size-consistent, and has more Slater determinants contributing to the wave function than CISD, even though the truncation order of T is the same.

In spite of its benefits, the **UCCSD** ansatz is not used in classical computations because it can't be efficiently evaluated. Regardless of the truncation order of the cluster operator, expression 2.50 is classically intractable, because the **BCH** expansion for $e^{-(T-T^\dagger)} H e^{T-T^\dagger}$ is infinite.

What is interesting is that while originally the unitary variant of the CC operator was created with the purpose of obtaining a variational version of CC theory, this unitarity has the convenient side effect of allowing straightforward implementation of the variational form in quantum computers.

In fact, the UCCSD ansatz was used in the article that introduced the VQE (reference [67]). Further, this ansatz has inspired many other proposals of variational forms for use in chemistry applications.

2.3 Variational Quantum Algorithms

This section will present the general outline of variational quantum algorithms, as well as the context that motivated their ingress into the field of quantum computing and prompted their popularity. An overview of the subject can be found in the article [17], used as the main reference throughout the subsections to follow.

2.3.1 Motivation

The main incentives for devising variational quantum algorithms are no other than the main limitations of quantum computers. The problems mentioned in subsection 2.1.4 - reduced qubit count and connectivity; coherent, incoherent and SPAM noise - currently inhibit quantum hardware from being large enough and good enough to implement full algorithms. The prospect of a fault-tolerant era in quantum computing seems distant, if promising.

As it was mentioned before, in 2014 the VQE was proposed as the first algorithm inserting a quantum computer in an optimization loop and obtaining a solution via the variational principle. The problem it meant to solve was one familiar to us from the previous section: finding the eigenstates of a Hamiltonian representing a chemical system.

Variational algorithms were far from the first approach for using quantum resources to solve the Schrödinger equation. In fact, the problem is strongly tied with the idea of using quantum computers for the simulation of quantum mechanical systems as proposed by Feynman. Because of the exponential growth of the classical resources required to describe a molecule as its size grows, the task at hand is an example of something that quantum computers were created to tackle. And by the time quantum computers were still in the conceptual realm, algorithms were certainly not developed with noise-resilience or qubit limitations in mind - rather provable speedups.

If fault-tolerant quantum computers were available to us now, the problem could be solved by means of the QPE algorithm. QPE allows obtaining the eigenvalue of a unitary operator corresponding to an eigenstate given as input. By applying QPE to the (evidently unitary) time evolution operator e^{iHt} , it is possible to find the energy eigenstates. The only requirement is that we have a way of preparing states with high overlap with the eigenstates, which can be met by resorting to adiabatic state preparation (ASP). This procedure of using the QPE routine in combination with ASP was detailed in [94].

However, this approach does not seem feasible in the [NISQ](#) era: employing the [QPE](#) routine to solve relevant problems would require the application of millions to billions of quantum gates [67], implying a circuit depth beyond what can be endured by near-term quantum computers (given that they're not error-corrected, and coherence times are still limiting).

This is where variational hybrid quantum-classical algorithms come in. They resort to the help of a classical optimizer to spare the quantum computer from bearing the load of the algorithm in full. The goal is to provide a near-term solution for problems that are classically hard to solve but, hopefully, easy for quantum computers.

Another good example is the [Quantum Approximate Optimization Algorithm \(QAOA\)](#), a variational algorithm aiming to solve combinatorial optimization problems introduced in [24]. Among the problems suitable for [QAOA](#) are the [Quadratic Unconstrained Binary Optimization \(QUBO\)](#) problems, of which MaxCut is a famous example with several real-life applications (ranging from physics to communication networks and circuit design).

What is interesting is that an infinitely slow adiabatic evolution leading to the solution of the problem is a special case of the [QAOA](#) circuit, in the limit of infinitely many layers. If we allow the [QAOA](#) ansatz to be big enough, there is a performance guarantee based on the adiabatic theorem: the ansatz is sure to contain the solution.

The [QAOA](#) both parallels and contrast with the [Quantum Adiabatic Algorithm \(QAA\)](#), which provides a more solid solution by relying in full in the adiabatic theorem. [QAOA](#) reduces the circuit depth under what would be necessary for a truly adiabatic evolution, dispenses with the assurance of the adiabatic theorem, and inserts the circuit in a classical optimization loop, converting some gate parameters to *variational* parameters. The hope is that the extra variational freedom may make up for what is lost by using a shallower circuit.

While the parallel with [QAA](#) seems like a good reason to believe the [QAOA](#) optimization might be successful, if one employs a lower number of layers than required by the adiabatic theorem, the theorem asserts nothing. So against [QAA](#), an algorithm with solid foundations but demanding the unattainable in the [NISQ](#) era, we have [QAOA](#): a [NISQ](#)-friendly algorithm of tentative construction.

The priority shift from performance guarantees offered under the assumption of a fault-tolerant scenario to near-term viability is the bedrock of variational quantum algorithms.

2.3.2 Outline

As [VQAs](#) might differ significantly in some key aspects and cover a wide variety of applications, quantum algorithms belonging to this class exhibit great structural affinity among themselves. A high-level, greatly simplified scheme of the outline of [VQAs](#) is presented in figure 6.

The basic idea of [VQAs](#) is to use a classical computer to optimize an expectation value measured in a quantum computer. In the following, the four steps labeled in the scheme will be explained in more detail, assuming the quantum circuit model.

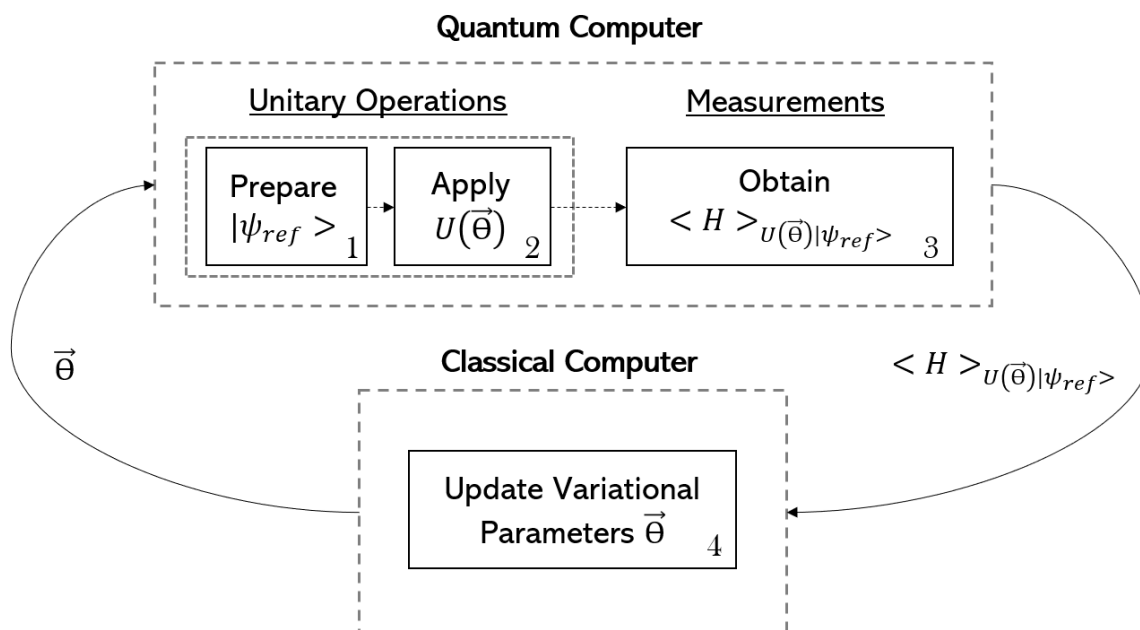


Figure 6: Typical outline of a variational quantum algorithm.

1 - Preparation of the Reference State

The quantum part of the algorithm starts with the preparation of a reference state $|\psi_{ref}\rangle$. This is the first portion of the *ansatz*, and it will typically be a shallow (often constant-depth) circuit with fixed gates and no variational freedom.

The role of this circuit is to start the computation in a state other than the typical computational basis state $|0\dots 0\rangle$. In the specific case of chemistry applications, this step might be of great importance in assuring that the computation starts in a state belonging to the correct symmetry subspace dictated by the problem. For example, it might be desired that the initial state has the correct particle number, and the correct total spin and Z spin projection expectation values.

Of course, it isn't always necessarily the case that a reference state different from the all-zero state is desired; when such occurs, the reference state preparation step is simply omitted.

2 - Application of the Parameterized Unitary

After the reference state has been prepared, the parameterized portion of the *ansatz* circuit follows. This is the part that contains the variational freedom. This is represented in the scheme by the parameterized unitary $U(\vec{\theta})$. This notation simply means that, rather than a fully-fixed circuit, a parameterized circuit is applied in this step. The parameters are considered to be organized into a parameter vector $\vec{\theta}$. After application of the *ansatz*, we are left with the parameterized state $U(\vec{\theta})|\psi_{ref}\rangle$, which will differ among iterations (via the dependence on the parameter vector, updated each round).

The complete *ansatz* then consists of the combination of the circuit that prepares the reference state

(step 1) with the parameterized circuit (step 2). The specific form of the ansatz greatly depends on the specifics of the problem and the approach. Ansätze can be labeled into three main categories.

Static, problem agnostic ansätze are those fixed upfront and completely oblivious to the problem. Ansätze of this type don't attempt to leverage any sort of previous knowledge about the problem, and don't bother spanning specific subspaces in which the solution is expected to be contained. Typically, this represents a shift of priority from problem-tailoring to hardware-tailoring. The *hardware-efficient* ansätze, that attempt to reduce depth by constructing a circuit directly from gates natural to the specific hardware, belong to this category. Despite the evident advantages of strategies like this one, searching for the solution in an exponentially big space without an informed choice of variational form creates other problems. Too much expressibility and too little problem-tailoring lead to *barren plateaus* in the optimization landscape, causing the gradients of the cost function to vanish exponentially on the number of qubits and rendering the full algorithm inefficient.

Static, problem tailored ansätze are also fixed upfront, but they do attempt to leverage information about the problem at hand. For example, in chemistry, this class includes ansätze that respect symmetries, be those symmetries generic or specific to the system in study. One can enforce time-reversal symmetry, as well as spin and particle number conservation. The [UCCSD](#) ansatz introduced in the previous section is an example of an element of this class. While these ansätze don't typically suffer from barren plateaus (aside from the noise-induced ones) [17], they have some issues of their own. They are not always easy to devise, because information about the problem can be limited or, if available, difficult to leverage. Additionally, they typically correspond to longer ideal circuits (since they're not developed with shallowness as the end goal) and even longer transpiled circuits (since the available gates or connectivity are also not the main concern).

Dynamic, problem tailored ansätze differ from the previous two in an important aspect: they are not fixed upfront, rather grown on the fly as the algorithm evolves. This implies a very special degree of tailoring: instead of human fabricated problem tailoring happening beforehand, additional measurements are used *during* the computation to gather information about which parameterized gates should be added to the ansatz at each point. The idea is to optimize, as well as the parameters, the variational form itself, which is consequently not only tailored to the problem as a whole: it is tailored to each specific instance of the problem. [ADAPT-VQE](#) was the first such ansatz to be introduced in the literature. It grows the ansatz from scratch by selecting operators from a pool, based on their gradients. Unlike [UCCSD](#), which is tailored to electronic structure problems in general, [ADAPT-VQE](#) is tailored to each specific molecule we might want to study. The main disadvantage of dynamically created ansätze is that they result in a sizeable overhead in measurement and optimization costs; however, they have multiple advantages. Their design is not an issue, because the ansatz structure is dictated by the the problem at hand - one needn't find a way to embed information about the problem into the construction of the circuit, because the algorithm learns how to do that autonomously. The resulting circuit is typically of reduced and controllable depth, due to the higher degree of customization and the existence of a tunable convergence criterion. Further, the accuracy can surpass that of fixed-depth problem tailored ansatz circuits.

3 - Cost Function Sampling

The final step to occur in the quantum computer is the evaluation of the expectation value of the cost function - i.e., the function whose minimum we want to find - in the trial state that was prepared. This is the information that will be passed on to the classical computer at the end of each iteration.

The operator whose expectation value we wish to find is typically denoted H , because it may correspond to a physical Hamiltonian. However, this ‘Hamiltonian’ can also be fabricated to encode a problem that does not have any relation to a physical system. Regardless of the interpretation of the operator, we can view this step as obtaining the expectation value of H measured in the state $U(\vec{\theta})|\psi_{ref}\rangle$, abbreviated to $\langle H \rangle_{U(\vec{\theta})|\psi_{ref}\rangle}$ in figure 6.

It was mentioned in section 2.1 that, in general, obtaining an expectation value implies several measurements. Thus, even though this is not represented on the scheme so as to privilege terseness, the estimate of the expectation value is obtained by repeating the state preparation and measurements multiple times, and averaging over the results. Even under ideal conditions, sampling entails some noise in the result: the inevitable quantum projection noise, and shot noise, deriving from the use of a finite number of repetitions. The accuracy of the result can be increased arbitrarily by increasing the number of shots (also called repetitions) of the circuit: the error goes with the inverse of the square root of the number of shots ($\epsilon \propto \frac{1}{\sqrt{N}}$).

The diagram also disregards the evaluation of more than one expectation value in each iteration. Oftentimes, the classical optimizer will make multiple calls to the quantum processor before moving on to the next iteration (i.e. updating the parameters). For example, some optimization methods may require a few gradient evaluations per iteration. In many relevant cases, this can be done by means of measurements on the quantum computer using the parameter-shift rule [73].

4 - Parameter Vector Update

Finally, we have the step performed by the classical computer: the parameter update.

Based on the expectation value of the cost function (and possibly its gradients) obtained in the state prepared by the previous parameters, a classical optimizer will decide on the parameters for the next iteration. This can be done using the method of gradient descent, among others.

After the optimizer decides how to update the parameter vector $\vec{\theta}$, this vector will be fed to the quantum computer. The effect of this is to instantiate the ansatz, binding the specific parameter values to the gates occurring in the variational circuit. This will allow the quantum computer to obtain the expectation value of H in the new trial state, which will then be passed on to the classical optimizer, and so on.

The cycle stops when some *convergence criterion* is met. This is decided by the classical computer. A common stopping criterion is the absolute change in the value of the cost function between two consecutive iterations being under a certain threshold.

2.3.2.1 Variability among VQAs

While the previous succinct explanation focused mostly on the common factors between all VQAs, this generic scheme embraces considerable diversity. As an example of this, we may compare two VQAs that were already mentioned in this section: the VQE and the QAOA.

The VQE addresses chemistry problems: its goal is to find the ground state of a system. The cost Hamiltonian is simply the electronic Hamiltonian of such system, and its expectation value is the energy. This molecular Hamiltonian only needs to be mapped onto an observable that can be measured in a quantum computer. As happens in many classical variational algorithms for chemistry, the reference state is frequently the Hartree-Fock state, and the parameterized portion of the ansatz commonly consists of circuits that mimic fermionic excitations applied to the reference determinant.

On the other hand, the QAOA aims to solve combinatorial problems; the solution is a classical bit string, and the cost function is classical. Once again, this function must be represented by a quantum observable; in this case, the cost Hamiltonian is diagonal and consists of a sum of projectors onto computational basis states multiplied by the corresponding cost values. The reference state is the homogeneous superposition of computational basis states, which is the maximum energy eigenstate of another fabricated Hamiltonian called the *mixer Hamiltonian*. The parameterized portion of the ansatz consists of p rounds of time evolution under alternating Hamiltonians - this mixer Hamiltonian and the cost Hamiltonian. These rounds are resembling of a trotterized adiabatic process, except here the evolution times of each round correspond to independent variational parameters. In the limit that p tends to infinity, we can make the Trotter steps very small, and the runtime very large (corresponding to a very slow evolution). In this limit we are certain that the QAOA ansatz contains the solution, because for some variational parameters the circuit implements the slowest possible adiabatic evolution leading the maximum energy state of the mixer Hamiltonian to the maximum energy state of the cost Hamiltonian - i.e., the solution.

VQE and QAOA thus have distinct reference states, and the parameterized portion of their ansätze is even more distinct. Notably, while the VQE ansatz is independent of the Hamiltonian, the QAOA ansatz depends on it. The cost functions have distinct origins, even though both can be represented by a quantum observable: the one used in VQE arises from a real system and the one used in QAOA is fabricated to represent a classical problem. Consequently, there is also a different interpretation for the expectation values of these functions. The VQE expectation value is the energy of the system in the fermionic state represented by the state of the qubits; in contrast, the QAOA expectation value is the (weighted) average of the cost function over the several bit strings (computational basis states) that appear in the superposition state of the qubits.

Given the classical nature of the problem, the QAOA solution is a *single* computational basis state (the one with the highest amplitude), while in VQE the *full* superposition state is the solution. This is also related to *how* we expect these quantum algorithms to bring advantage as compared to classical ones. The hope for quantum advantage in VQE is rooted in the fact that the classical description of the molecular system grows exponentially with its size, i.e. VQE attempts to make use of quantum resources for a problem that is

quantum by nature. In contrast, [QAOA](#) attempts to make use of *quantum* superposition to simultaneously explore several solutions to a *classical* problem.

2.3.3 Variational Quantum Eigensolver

As it was explained before, the variational quantum eigensolver was introduced in [67] with the purpose of finding the ground energy and ground state of a chemical system. The approach can be generalized to excited states by means of the folded spectrum method [57, 91].

The essential idea behind the algorithm is to use a classical optimizer to find the minimum eigenvalue of a Hamiltonian, whose expectation value is evaluated in a quantum computer.

The purpose of this section is to explain a few important aspects that have not yet been covered, including the representation of the problem and specific details concerning the quantum expectation estimation in the context of [VQE](#). A deeper overview of the several components of [VQE](#) can be found in [59].

2.3.3.1 The Jordan-Wigner Transform

In order to be able to solve a fermionic problem in a quantum computer, we need to be able to encode the evolution of a fermionic state in the evolution of the state of the qubits through unitary operations. This is done through a *fermion-to-qubit mapping*, of which there are several; the most widely used one, and the one employed in this dissertation, is the Jordan-Wigner Transform [48].

The first thing that should be pointed out is that the state of a fermion cannot in general be described by a binary variable. For example, the state of a bound electron in an atom depends on the principal quantum number n , the azimuthal quantum number l , the magnetic quantum number m_l , and the secondary spin quantum number m_s . Aside from m_s , these quantum numbers are not even binary variables themselves - and all four are necessary for a complete description. As such, representing the state of a fermion in the state of a qubit is not feasible.

Another problem is that qubits are distinguishable, while identical fermions are indistinguishable. Even if the state of a fermion *could* be fully specified by a binary variable, representing it in the state of a qubit would not work. As an example, while the two two-qubit states

$$|\psi_{Q,1}\rangle = |0\rangle \otimes |1\rangle, \quad |\psi_{Q,2}\rangle = |1\rangle \otimes |0\rangle$$

are distinct, they would not be if we were dealing with fermions. Because fermions are identical particles, exchanging two of them cannot affect any physical process. This is enforced by the antisymmetry requirement presented in equation 2.35: if we exchange any two fermions, the many-body fermionic wave function simply flips the sign. Ensuring the antisymmetry of the wave function guarantees that we are not discerning between states that are not physically distinguishable. For the simple two-body example, if we have a fermion in a state labeled **0** and one in a state **1** (regardless of what those states are), we can write the wave function as

$$|\psi_F\rangle \propto |0\rangle \otimes |1\rangle - |1\rangle \otimes |0\rangle,$$

which has the required property that exchanging the two fermions (putting the fermion to the left of the tensor product in the state of the one to the right and vice-versa) results in a change of sign, but does not affect any measurable properties of the wave function. It is easy to see that any expectation value is left unchanged.

$$|\psi_F\rangle \rightarrow -|\psi_F\rangle$$

$$\langle\psi_F|\hat{O}|\psi_F\rangle \rightarrow (-\langle\psi_F|)\hat{O}(-|\psi_F\rangle) = \langle\psi_F|\hat{O}|\psi_F\rangle$$

The two-body example is very simple, but as was discussed, we can write a generic N-body wave function through Slater determinants, formula 2.36.

All of this trouble to express the many-body fermionic wave function is merely a symptom of the notation in use being cumbersome for our purpose. By specifying the many-body wave function in the basis of tensor products of states of the individual particles, we are distinguishing between what isn't physically distinguishable. Because we are discriminating the state of each individual particle, it is then necessary to remove the redundancy in the notation by enforcing the antisymmetry of the wave function.

This is where the *occupation number representation* of the second quantization proves useful. In dealing with identical particles, it seems more natural to specify *how many* particles are in *each* state, rather than *which* particles are in *which* state. Instead of having all the work antisymmetrizing the wave function, it suffices to describe it by enumerating the occupation numbers of each single-particle state (also called *fermionic mode*), a formalism which naturally embraces indistinguishability and clears up the notation. In the previous example, we write

$$|\psi_F\rangle = |n_0 = 1, n_1 = 1\rangle$$

to mean that there is a fermion in the state labeled 0 (occupation number n_0 is 1) and another in the state labeled 1 (occupation number n_1 is 1), without making any reference to which fermion is in which state. In the case of a bound electron in an atom, we could specify the occupation number of each *spin-orbital* (identified by the four quantum numbers n, l, m_l, m_s). In accordance with the Pauli exclusion principle, each occupation number must be 0 or 1, which can be readily seen from the fact that the antisymmetry principle 2.35 causes the wave function to vanish if any two fermions occupy the same state.

The second quantization formalism clears up the notation by relaying the task of ensuring antisymmetry to the *fermionic ladder operators*. For any fermionic mode i , there is the creation operator a_i^\dagger and the annihilation operator a_i , which act as

$$\begin{aligned}
 a_i^\dagger |n_i = 0\rangle &= |n_i = 1\rangle, & a_i^\dagger |n_i = 1\rangle &= 0, \\
 a_i |n_i = 1\rangle &= |n_i = 0\rangle, & a_i |n_i = 0\rangle &= 0.
 \end{aligned}
 \tag{2.52}$$

Attempting to create a fermion in an occupied state or to remove a fermion from an unoccupied state quenches the wave function.

The ladder operators have to obey the *anti-commutation relations*:

$$\{a_j^\dagger, a_i^\dagger\} = 0, \quad \{a_j, a_i\} = 0, \quad \{a_j, a_i^\dagger\} = \delta_{j,i},
 \tag{2.53}$$

where $\{A, B\} = A + B$ is the *anti-commutator* of operators A and B , and $\delta_{j,i}$ is the Kronecker delta (which is 1 if $j = i$, and 0 if not). The anti-commutation relations in 2.53 enforce anti-symmetry and serve as the algebraic definition of the fermionic ladder operators.

It should be noted that there is a notion of *order* at play. For example, if we use $|0\rangle$ denote the vacuum state (no fermions) and assume $i \neq j$, the expression

$$a_j^\dagger a_i^\dagger |0\rangle = -a_i^\dagger a_j^\dagger |0\rangle$$

arising from the first anti-commutator in 2.53 states that if we create a fermion in state i , then one in state j , there is a sign flip as compared to a situation in which fermions are created in the opposite order. The ladder operators are responsible for this sign flip, which ensures that the anti-symmetry principle is obeyed; thus, it is to be expected that their mapping to qubit operators will carry this notion of order.

It is now natural to introduce the *Jordan-Wigner transform*, a mapping from fermions to qubits. We match qubit states with the occupation number of fermionic modes, and map

$$\begin{aligned}
 a_i^\dagger &\rightarrow e^{i\pi \sum_{k=1}^{i-1} \sigma_k^+ \sigma_k} \cdot \sigma_i^+, \\
 a_i &\rightarrow e^{i\pi \sum_{k=1}^{i-1} \sigma_k^+ \sigma_k} \cdot \sigma_i^-,
 \end{aligned}
 \tag{2.54}$$

where the σ_i^+ , σ_i^- are the qubit raising and lowering operators acting on qubit i , defined by means of Pauli operators as

$$\begin{aligned}
 \sigma_i^+ &= \frac{1}{2}(X_i - iY_i), \\
 \sigma_i^- &= \frac{1}{2}(X_i + iY_i).
 \end{aligned}
 \tag{2.55}$$

Evidently, the action of σ_i^+ , σ_i^- on qubit states $|0\rangle_i, |1\rangle_i$ is identical to the action of a_i^\dagger, a_i on the fermion states $|0\rangle_i, |1\rangle_i$ as defined in 2.52.

$$\begin{aligned}
 \sigma_i^+ |0\rangle_i &= |1\rangle_i & \sigma_i^+ |1\rangle_i &= 0 \\
 \sigma_i^- |1\rangle_i &= |0\rangle_i & \sigma_i^- |0\rangle_i &= 0
 \end{aligned}
 \tag{2.56}$$

For a single fermionic mode, identifying σ_i^+ , σ_i^- with a_i^\dagger , a_i respectively would suffice. However, the antisymmetry requirement would not be satisfied, so this simpler mapping would not be fitting in a multi-mode scenario. That is the purpose of the phase that appears in 2.54.

The operator $e^{i\pi \sum_{k=1}^{i-1} \sigma_k^+ \sigma_k}$ is typically called *the Jordan-Wigner string operator*. It can be readily seen that it adds a phase $e^{i\pi N_{k<i}}$ to the operators acting on qubit i , where $N_{k<i}$ is the number of the occupied fermionic modes k for which $k < i$. This is crucial for the anti-symmetry requirement, and it is where the previously mentioned notion of order comes in. We have to establish and be consistent with a sequence, i.e. choose an order $\dots < i < j < \dots$ relating all fermionic modes.

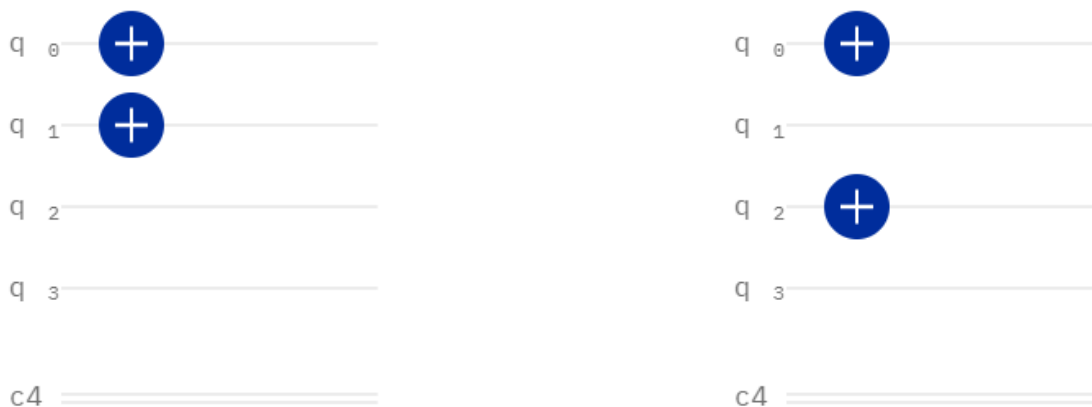
It is easy to see what is the effect of the Jordan-Wigner string by considering the successive creation of fermions in two different unoccupied modes. When the first fermion is added to the mode with the lowest index, it affects the string operator of the operator that creates the second fermion, but when the first fermion is added to the mode with the highest index, the string operator of the operator that creates the second fermion is left unchanged. This causes the wave function in these cases to differ by a phase factor $e^{i\pi} = -1$, as antisymmetry requires: whenever we swap the order of creation of fermions in two modes, we get a flipped sign.

It is convenient to write the Jordan-Wigner transform 2.54 using Pauli operators. This is easily done using 2.55 and noting that the effect of the Jordan-Wigner string on operators acting on mode i amounts to adding a minus sign when the parity of the occupation numbers of the modes under i is odd. This parity calculation can be done through a string of Pauli Z operators.

$$\begin{aligned}
 a_i^\dagger &\rightarrow \frac{1}{2} \prod_{k=1}^{i-1} Z_k \cdot (X_i - iY_i) \\
 a_i &\rightarrow \frac{1}{2} \prod_{k=1}^{i-1} Z_k \cdot (X_i + iY_i)
 \end{aligned}
 \tag{2.57}$$

The formulation of the Jordan-Wigner transform in 2.57 is very convenient for quantum computing applications.

Finally, it is worth noting that this definition of the Jordan-Wigner mapping is not unique. Often the roles of the operators σ_i^+ and σ_i^- in 2.54 are reversed, causing the empty and occupied fermionic modes to be represented by qubit states $|1\rangle$ and $|0\rangle$ respectively (the opposite of what the definitions presented here lead to).



(a) OpenFermion [2] orbital ordering.

(b) Qiskit [3] orbital ordering.

Figure 7: Circuit for preparing the Hartree-Fock reference state, for the H_2 molecule. The circuits were drawn in the IBM Quantum Composer [41].

2.3.3.2 Preparation of the Reference State

The reference state used in VQE is typically the Hartree-Fock approximation to the ground state. Because the computational basis states usually represent the Slater determinants obtained as approximate eigenstates of the system in performing the self-consistent mean field Hartree-Fock calculations, the reference state preparation state simply amounts to a $\mathcal{O}(1)$ circuit with N_e parallel Pauli X gates, where N_e is the number of electrons.

Preparing the reference state amounts to applying Pauli X gates on the qubits that, under the Jordan-Wigner mapping (see 2.3.3.1), correspond to occupied spin-orbitals.

The exact circuit depends on the convention for organizing the spin-orbitals. Figure 7 contrasts the reference state preparation for the hydrogen molecule in the two software packages used for creating the circuits: OpenFermion [2], used in conjunction with CIRQ [22], and Qiskit [3].

In the basis set that was always used in the simulations (the STO-3G minimal basis set), H_2 has four spin-orbitals, so that the electronic states are represented by four qubits.

OpenFermion alternates between α (up) and β (down) type orbitals. With this ordering, each spatial orbital is represented by two qubits that are adjacent on the register. Figure 7a shows how to create the state $|0011\rangle$, assuming little endian ordering. Under OpenFermion's convention, this corresponds to creating two fermions on the two spin-orbitals associated with the lowest-energy spatial orbital (the orbitals are assumed to be organized in ascending energy from right to left).

In contrast, Qiskit uses block-wise orbital ordering: all α orbitals come first, followed by all β orbitals. This means that the first qubit of the register corresponds to the same spatial orbital as the one in the middle of the register, with the two representing opposite spins. Figure 7b shows how to create the state $|1010\rangle$ (assuming little endian ordering once again). In spite of the differences in convention, the states

prepared by circuits 7a and 7b represent the same fermionic state. Evidently, the ansatz and Hamiltonian averaging must be done in accordance with the chosen convention.

2.3.3.3 Choice of Ansatz

As was discussed, the typical ansätze for VQE are chemistry-inspired. Even though they may not lead to the shallowest circuits, the fact that they're tailored to chemistry problems in particular helps create trial states in the appropriate symmetry subspace. As a consequence, the optimization is helped by the fact that subspaces where the ground state is sure not to be are entirely avoided. Ansätze that don't leverage previous knowledge about the system typically lead to barren plateaus. These trainability issues have been shown to be a consequence of the choice of ansatz that cannot be helped through the choice of optimizer [60, 10, 39].

UCC theory enjoys great popularity due to its natural genesis in the context of previous computational chemistry methods. UCC proposals are also fairly broad, and this class has been expanding as a consequence of research efforts to make the ansätze more compact and more accurate.

The article [80] explored singlet and pair Quantum Unitary Coupled Cluster Doubles (q-UCCD) approaches (where the q stands for *quantum*), which reduce the double excitations appearing in the ansatz to a smaller subset. Singlet q-UCCD splits the double excitations into singlet and triplet components, and makes use of symmetry / anti-symmetry to reduce the number of excitation operators. Pair q-UCCD limits the excitations to those which excite a *pair* of electrons with opposite spins from one spacial orbital to another. Both of these ansätze were tested in several systems with results of compelling accuracy, while reducing the number of two-qubit gates in the ansatz circuit by up to 75%.

A different ansatz termed k-Unitary Paired Coupled Cluster Generalized Singles and Doubles (UpCCGSD) was introduced in [53]. k-UpCCGSD consists of k products of unitary paired coupled-cluster double excitations, as well as all generalized single excitations. Here generalized indicates that virtual-to-virtual and occupied-to-occupied excitations are also allowed (orbitals are considered virtual or occupied if they are so in the reference state). The K-UpCCGSD ansatz increases the variational freedom as compared to the simple pair Unitary Coupled Cluster Generalized Singles and Doubles (UCCGSD) (pUCCGSD) ansatz, while still enjoying a linear scaling of the circuit depth with the system size. The k-UpCCGSD was shown to be capable of achieving chemical accuracy with a better resource scaling than UCCSD and UCCGSD (the generalized version of the former), both of which require an $\mathcal{O}(N^4)$ circuit depth, where N is the number of orbitals.

Changing the form of the ansatz is not the only option. Low-rank decomposition of the basic UCC operator was used in [62] to reduce the gate complexity of each Trotter step from $\mathcal{O}(N^4)$ to $\mathcal{O}(N^3)$. By exploiting the structure and symmetries of the cluster operator and truncating small terms, the depth of the circuits can be reduced without precluding chemical accuracy.

2.3.3.4 Hamiltonian Averaging

The part of the algorithm that consists of obtaining the expectation value of the cost Hamiltonian (in this case a physical Hamiltonian) is typically called *Hamiltonian averaging* or *quantum expectation estimation*.

In subsection 2.1.2, a procedure to measure arbitrary observables was outlined; it consisted of decomposing them as a linear combination of Pauli strings (formula 2.14) and calculating the expectation value as a weighed average of the corresponding expectation values (formula 2.15). The expectation value of the Hamiltonian can be obtained exactly this way; the question is whether this can be done efficiently.

Luckily, it is often the case that the Hamiltonian of a physical system can be decomposed as a linear combination of a polynomial number of Pauli terms. This includes the electronic structure Hamiltonian (formula 2.30), as well as the Hamiltonians of the Ising and Heisenberg models [67]. The problems addressed in this dissertation concern electronic Hamiltonians with a number of terms that scales like $\mathcal{O}(N^4)$ on the number of spin-orbitals/qubits N [75]. An example of such a Hamiltonian, for the very simple case of the hydrogen molecule, can be found in appendix B.

While the scaling is polynomial, obtaining the expectation value can still imply a large number of measurements, especially in chemistry problems which typically require exceptional accuracy. One way this can be improved is by organizing the Pauli strings appearing in the decomposition of the operator into commuting groups, that can be measured simultaneously.

As a simple two-qubit example, we can consider the following set of observables.

$$II, ZI, IZ, ZZ$$

The expectation values of all four terms can be obtained from the same circuit. There are multiple ways of arranging Pauli strings into commuting groups; for example, if we have the set of observables

$$ZI, IZ, ZZ, XI, IX, XX, XZ, ZX,$$

we can chose to split the Pauli strings into groups as

$$\{ZI, IZ, ZZ\}; \{XI, IX, XX\}; \{ZX\}; \{XZ\},$$

but also as

$$\{ZI, IX, ZX\}; \{XI, IZ, XZ\}; \{ZZ\}; \{XX\}.$$

These divisions only exploit qubit-wise commutativity. However, two Pauli strings also commute if they have non-commuting Pauli operators in an even number of indices. This general commutativity allows us to further reduce the circuit repetitions that are required to obtain the final expectation value. In our example, it brings them down to $\frac{3}{8}$ as compared to when no grouping is used; we can organize the operators as

$$\{ZI, IX, ZX\}; \{XI, IZ, XZ\}; \{ZZ, XX\}.$$

Even though this is evidently not the only possibility; we can also have

$$\{ZI, IZ, ZZ\}; \{XI, IX, XX\}; \{ZX, XZ\}.$$

In this example, both options reduce the necessary circuit repetitions by the same amount (assuming these repetitions are to be distributed equally). However, this is not always necessarily the case. What is more, the task of finding the optimal partitioning becomes harder the larger the Hamiltonian is, and taking into account sampling noise complicates the task further. Proposals of grouping strategies can be found in references [59, 93, 49, 72, 45, 47, 65, 46, 97, 28, 29, 89, 19, 40]. The last article among these proposed an approach that allows for a cubic reduction in term groupings over previous proposals, with the drawback of requiring the execution of a linear-depth circuit before the measurement. A table comparing all these strategies can be found in this same reference.

2.3.3.5 Estimating Shot Requirements

The number of terms in the Hamiltonian decomposition, and the number of commuting groups, are not the only factors weighing in the shot requirements of VQE. The impact of sampling noise in the error, which typically also increases with the size of the system, is another relevant aspect. The following estimation of shot requirements was based on reference [93].

Assuming that the Hamiltonian operator has been decomposed in the form of formula 2.14 and that for each Pauli string P_i we use M_i shots, the error coming from the finite number of shots in the expectation estimation can be written

$$\epsilon^2 = \sum_i \frac{|h_i|^2 \text{Var}(\hat{P}_i)}{M_i}.$$

The variances depend, not only on the Pauli operators appearing in the Hamiltonian, but also on the particular state. We can bound them as $\text{Var}(O_i) \leq 1$, since any Pauli operator has eigenvalues ± 1 and so will their product.

Without further knowledge about their variance, the best choice is to distribute the shots by the operators in proportion to the norm of their coefficients: $M_i \propto |h_i|$. Assuming this distribution, one can estimate the necessary number of shots to hit a precision ϵ as

$$M \approx \frac{(\sum_i |h_i|)^2}{\epsilon^2}.$$

The term $(\sum_i |h_i|)^2$ will depend strongly on the molecule in study: it will be larger for larger molecules, requiring a larger number of shots for the same precision if the system size is increased.

Given that chemical accuracy is $1 \text{ kcal/mol} \approx 1.59 \times 10^{-3} \text{ Hartree}$, achieving it requires

$$M \approx 0.4 \times 10^6 \times \left(\sum_i |h_i| \right)^2$$

shots. This is already of the order of 10^7 to 10^8 for small molecules like helium and lithium hydride; Fe_2S_2 , for instance, would require around 10^{13} shots.

This number of shots needs to be performed every time one needs to evaluate the energy, which happens many times throughout the optimization. The shot requirements are heavier for larger molecules, since there are more Pauli strings in the Hamiltonian, larger coefficients $|h_i|$, and more parameters to optimize (resulting in more iterations and more function evaluations per iteration being required).

Further, this shot count doesn't guarantee proper functioning of the optimizer: it ensures that the difference between the energy estimate and the true energy *in a state* are within chemical accuracy, but the final state isn't necessarily close to the ground state. The error might still be enough to confuse the optimizer into outputting a wrong guess for the ground state (e.g. by tricking it into an early convergence).

Static Ansätze for VQE

This chapter aims to present the results of VQE simulations with static ansätze. In the first section, an entirely problem-agnostic ansatz will be applied in searching for the ground state of the helium hydride ion. The impact of the optimizer, optimization hyperparameters, and noise on performance will be assessed. In the second section, a problem-inspired ansatz will be applied in searching for the ground state of the hydrogen molecule. The impact of several types of noise on performance will be analysed.

For creating and executing quantum circuits, CIRQ [22] and Qiskit [3] were used in alternation; the CIRQ simulator was used as well as IBMQ's simulators and real quantum computers. Which package and backend was used in obtaining which results will be clear along the chapter.

When analysing the results, the final VQE energy and state will be compared with those obtained from performing exact diagonalization on the FCI Hamiltonian. The FCI solution is exact up to the Born-Oppenheimer approximation, relativistic effects, and the chosen basis set.

3.1 Problem Agnostic Ansätze

As was explained before, problem agnostic ansätze come with several advantages: they don't demand previous knowledge about the problem or system, nor a way of leveraging such knowledge. Additionally, taking the focus away from the problem allows for a more hardware-friendly circuit implementation, which may decrease the depth of the circuits. The main disadvantage of this type of ansatz is that randomly initialized parameterized quantum circuits often lead to barren plateaus in the optimization landscape [60].

This section will present simple a example of a problem agnostic ansatz, applied to the helium hydride ion. An ansatz spanning the whole Hilbert space was created and employed in the simulations; such an approach is feasible given the dimensionality of the problem. The section contains results from running the algorithm on the CIRQ simulator and on IBMQ's backends (simulators or real quantum processors).

3.1.1 Application to Helium Hydride

The purpose of the following is to implement VQE as in the article that originally introduced it [67], which applied the algorithm to the helium hydride ion using a photonic quantum processor.

In addition to the CIRQ and Qiskit implementations, an independent noise-free version of VQE, based on matrix algebra, was implemented to verify correctness.

The Hamiltonians for the different bond distances were taken from the article (which in turn obtained them via the PSI3 computational package [20]), as was the parameterization of an arbitrary 2-qubit state.

The ansatz employed here searches through the whole Hilbert space, much like the ansatz in the original article. However, while in the article the parameterization was rewritten into a convenient form that directly dictated the values of six phase shifters in the photonic chip, here the state preparation was abstracted of the specifics of any physical implementation and instead mapped to a generic circuit.

3.1.1.1 Preparing an Arbitrary 2 Qubit State

An arbitrary (pure) two qubit state can be written as in equation 3.1.

$$|\psi\rangle = \alpha |00\rangle + \beta |01\rangle + \gamma |10\rangle + \delta |11\rangle \quad (3.1)$$

The computational basis coordinates α , β , γ and δ are complex numbers that fully determine the state.

Since optimizers typically deal with real variables, it is convenient to rewrite the state using real parameters. Further, the degrees of freedom corresponding to the normalization and the global phase should be removed. With this, the two qubit state can be described by six real variational parameters. A possible parameterization is presented in 3.2.

$$|\psi\rangle = \cos \frac{\theta_0}{2} \cos \frac{\theta_1}{2} |00\rangle + \cos \frac{\theta_0}{2} \sin \frac{\theta_1}{2} e^{i\omega_1} |01\rangle + \sin \frac{\theta_0}{2} e^{i\omega_0} \cos \frac{\theta_2}{2} |10\rangle + \sin \frac{\theta_0}{2} e^{i\omega_0} \sin \frac{\theta_2}{2} e^{i\omega_2} |11\rangle \quad (3.2)$$

Here, $\theta_i \in [0, \pi]$ and $\omega_i \in [0, 2\pi]$. This was the parameterization used in the original VQE article [67], and it will be used here as well.

We wish to transform this parameterized description of the state into a parameterized circuit that we can use as our ansatz. The concept of Schmidt decomposition, presented in theorem 2.7 of reference [64], is useful here.

If we have a pure state of a composite system AB , $|\psi_{A,B}\rangle$, then the Schmidt decomposition allows us to write it as

$$|\psi_{A,B}\rangle = \sum_i \lambda_i |i_A\rangle |i_B\rangle, \quad (3.3)$$

where the $|i_A\rangle$, $|i_B\rangle$ are orthonormal states of the systems A , B respectively. They form the *Schmidt basis* of the respective system. The λ_i , called *Schmidt coefficients*, are real non-negative numbers that satisfy $\sum_i \lambda_i^2 = 1$.

We can find the Schmidt decomposition of a two qubit state written in the computational basis (formula 3.1) by using *Singular Value Decomposition (SVD)*. Firstly, we must organize the coordinates into a two-by-two matrix M .

$$M = \begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix} \quad (3.4)$$

This is done by identifying the coordinate of computational basis state $|ij\rangle$ with the matrix entry $M_{i,j}$: the left qubit, A , is associated with row indices, and the right qubit, B , is associated with column indices. Calculating the scalar product of the state of the composite system with a state of system A or B ($|\psi_A\rangle$ or $|\psi_B\rangle$) corresponds to left- or right-multiplying matrix M by the bra $\langle\psi_A|$ or the ket $|\psi_B\rangle$, which results in a linear superposition of the rows or columns of matrix M , respectively. Calculating the scalar product of the state of the composite system AB with a computational basis state of the same system amounts to doing both simultaneously, i.e.

$$\langle ij|\psi_{A,B}\rangle \equiv \langle i|\begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix}|j\rangle.$$

We can now use [SVD](#) decomposition to write M in the more convenient way of [3.5](#).

$$M = U \begin{pmatrix} \lambda_0 & 0 \\ 0 & \lambda_1 \end{pmatrix} V \quad (3.5)$$

The λ_i are the *singular values* of matrix M , non-negative real numbers that we can easily identify with the Schmidt coefficients of formula [3.3](#). U and V^T (note the transpose) are two-by-two unitary matrices, whose column vectors form the Schmidt basis for A and B respectively.

It will be necessary to create a circuit to implement U and V^T . A generic single-qubit unitary can be written by means of three rotations, along with a global phase. This is done in [3.7](#).

$$U = e^{i\theta_0} R_Z(\theta_1) R_Y(\theta_2) R_Z(\theta_3) = e^{i\theta_0} e^{-\frac{i\theta_1 Z}{2}} e^{-\frac{i\theta_2 Y}{2}} e^{-\frac{i\theta_3 Z}{2}} \quad (3.6)$$

$$= e^{i\theta_0} \begin{pmatrix} e^{-\frac{i(\theta_1+\theta_3)}{2}} \cdot \cos\left(\frac{\theta_2}{2}\right) & -e^{-\frac{i(\theta_1-\theta_3)}{2}} \cdot \sin\left(\frac{\theta_2}{2}\right) \\ e^{-\frac{i(-\theta_1+\theta_3)}{2}} \cdot \sin\left(\frac{\theta_2}{2}\right) & e^{-\frac{i(-\theta_1-\theta_3)}{2}} \cdot \cos\left(\frac{\theta_2}{2}\right) \end{pmatrix} \quad (3.7)$$

The rotation angles can easily be found from this; for example, they can be obtained using the expressions in [3.8](#), where $\angle U_{ij}$ denotes the phase ϕ of the complex number $\rho e^{i\phi}$ in line i and column j of matrix U , and $|U_{ij}|$ denotes its modulus ρ .

$$\begin{aligned} \theta_0 &= \frac{\angle U_{00} + \angle U_{11}}{2} & \theta_1 &= -\angle U_{00} + \angle U_{10} \\ \theta_2 &= 2 \arccos |U_{00}| & \theta_3 &= -\angle U_{00} + \angle(-U_{10}) \end{aligned} \quad (3.8)$$

It should be noted that θ_0 corresponds to a global phase with no relevance. The only parameters that are important and should be found are θ_1 , θ_2 and θ_3 .

A procedure to create a circuit that prepares a state from its four complex coordinates in the computational basis can now be easily described. Once U , V , λ_0 and λ_1 have been found by [SVD](#) decomposition (allowing us to rewrite the coordinate matrix as in [3.5](#)), it can be divided into the three following steps:

1. Create the state $\lambda_0 |00\rangle + \lambda_1 |10\rangle$ (from $|00\rangle$). This can be done by rotating qubit A around the Y axis by an angle of $2 \arccos |\lambda_0|$. Because $\lambda_0^2 + \lambda_1^2 = 1$, this is the same angle as $2 \arcsin |\lambda_1|$.
2. Create the state $\lambda_0 |00\rangle + \lambda_1 |11\rangle$ (from $\lambda_0 |00\rangle + \lambda_1 |10\rangle$). This can be done by applying a **CNOT**, with qubit A as the control and qubit B as the target.
3. Create the state $U \otimes V^T(\lambda_0 |00\rangle + \lambda_1 |11\rangle)$ (from $\lambda_0 |00\rangle + \lambda_1 |11\rangle$). This can be done by decomposing the single qubit unitaries U and V^T into single qubit rotations, and applying them to qubits A and B respectively.

Steps 1 and 2 create the state $\lambda_0 |00\rangle + \lambda_1 |11\rangle$, which is the desired state *in the Schmidt basis*. Step 3 rotates each qubit from its Schmidt basis to the computational basis, creating the final state. The corresponding unitaries are those formed by the Schmidt vectors of the respective qubits.

The outlined procedure leads to the state preparation circuit presented in figure 8. Such an ansatz has 7 variational parameters: one from the rotation on the first qubit in step 1, and 3 from the decomposition of each of the unitaries in step 3.




Figure 8: Circuit with 7 variational parameters for preparing an arbitrary 2-qubit state, drawn in the IBM Quantum Composer [41].

One thing that stands out is that there were six variational parameters in the parameterization 3.2. In fact, there are six degrees of freedom left in a two-qubit state once we remove those corresponding to the normalization and the global phase. This means that there is a redundant degree of freedom in our parameterized circuit.

Luckily, this can easily be solved by noting that after the two first steps (i.e. after the **CNOT** gate), we have a symmetric state $\lambda_0 |00\rangle + \lambda_1 |11\rangle$. This means that the first two parallel Z rotations, with parameters $p2$ and $p5$ in figure 8, can be merged into a single Z rotation with angle $p2 + p5$ on either of the qubits, with exactly the same effect.

The final 6-parameter ansatz is presented in figure 9, where the parameter labels were redefined to run from $p1$ to $p6$.

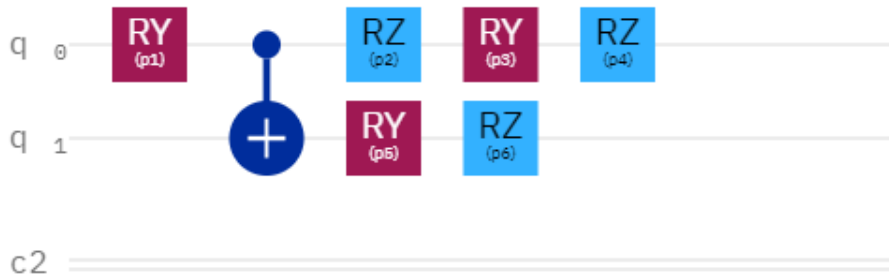


Figure 9: Circuit with 6 variational parameters for preparing an arbitrary 2-qubit state, drawn in the IBM Quantum Composer [41].

In this project’s VQE implementation, the parameters chosen by the optimizer were the ones corresponding to parameterization 3.2. At each step, the parameters output by the optimizer were treated as follows:

- Calculate the complex coordinates α , β , γ , δ (formula 3.1) from the parameterization 3.2. Organize them into a two-by-two matrix M as in formula 3.4.
- Use SVD to obtain the U , V , λ_0 , λ_1 that allow rewriting the matrix M as in formula 3.5. This was done using the `svd` function from SciPy’s [6] linear algebra module.
- Use the three-step procedure outlined before to create the 6-parameter circuit that prepares the desired state.

3.1.1.2 Bond Dissociation Graph

Before delving into a deeper analysis of the impact of the optimizer and of generic noise in the VQE results, a simple illustration of the performance of the algorithm is presented in figure 10. This figure shows the VQE results for the helium hydride ion along its bond dissociation curve, for simulations with and without sampling noise (and no other types). The Nelder-Mead simplex method from SciPy [6] was used as the optimizer, with properly adjusted hyperparameters.

It can be seen that the results are mostly consistent from run to run, and even with only 1000 shots in the expectation estimation the energy is consistently close to the ground energy obtained from exact diagonalization.

It must be noted that the employed ansatz was overly generic (and overly expressible as a consequence), and it was used along with a completely random starting state. In most runs, the starting state did not have the correct particle number (among other symmetries) nor appreciable overlap with the ground

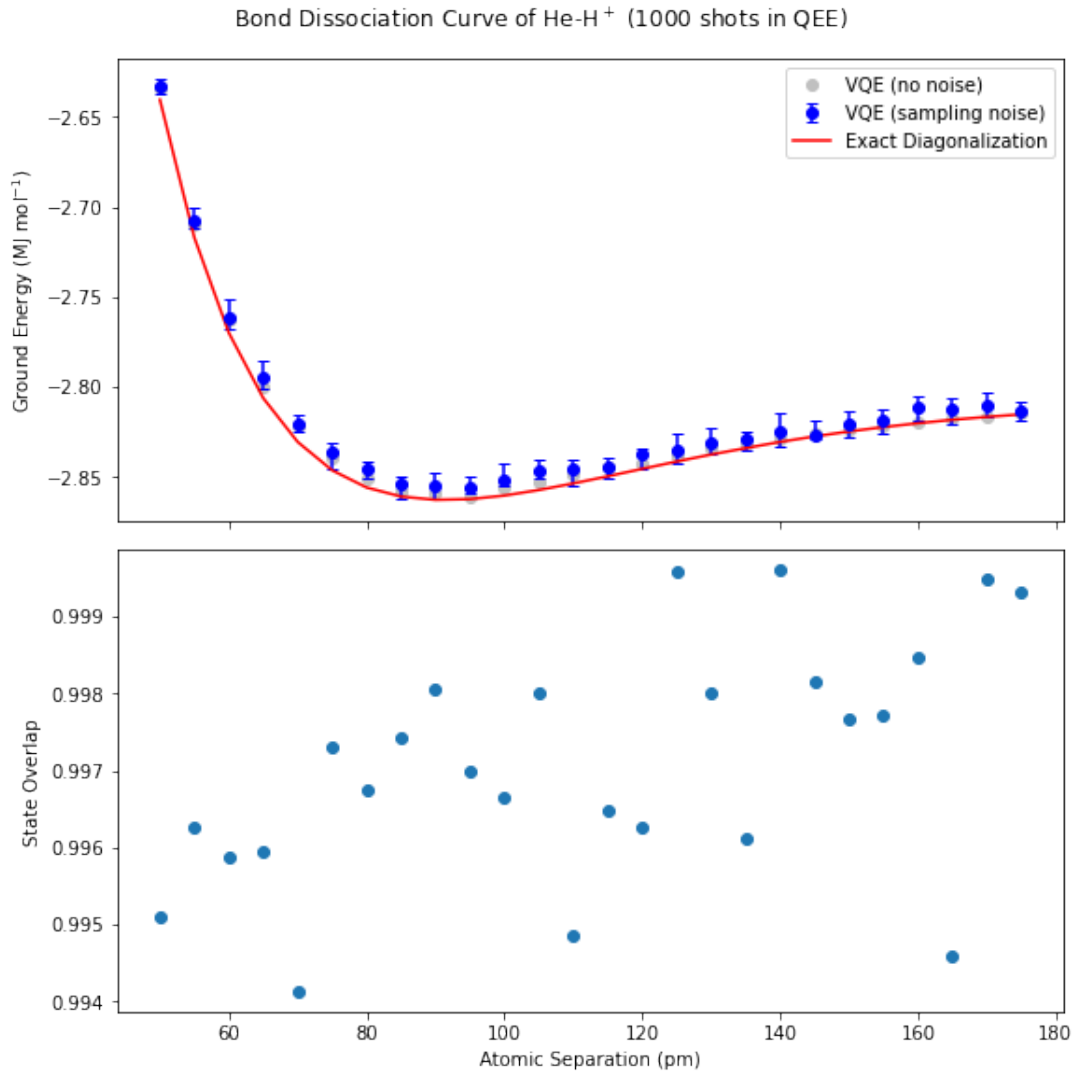


Figure 10: Performance of the VQE algorithm along the bond dissociation graph of HeH^+ , with and without sampling noise. The upper plot shows the VQE energy, with and without sampling noise, plotted against the exact diagonalization energy (which is exact within the basis set and up to the Born-Oppenheimer approximation and relativistic effects). The lower plot shows the overlap of the VQE state (with sampling noise) with the ground state obtained by exact diagonalization. In the cases with sampling noise, 1000 shots were used. 25 points were used for each interatomic distance; the median was used as the estimator, and the error bars show the interquartile ranges. The quantum circuits were run on the CIRQ simulator [22].

state. Regardless, the VQE energy points clearly trace out the bond dissociation curve. Further, at any given point, the final VQE state has an overlap of over 99% with the true ground state.

3.1.1.3 Effect of the Optimizer and Hyperparameters

The role of the classical optimizer is pivotal not only in VQE, but in VQAs in general: a bad optimizer might hinder the performance of the algorithm. Further, what is a *good* and a *bad* optimizer is specific to the quantum scenario, since noise (be it sampling or others) implies an added difficulty, and one that is not typically contemplated in benchmarking classical optimizers.

Another relevant factor in the optimization are the hyperparameters - the parameters that control the learning process. Because optimization-based algorithms are themselves computationally heavy, it is typically not easy to find the optimal hyperparameters, and they may not generalize from small-scale instances of the problem.

To illustrate the relevance of all of this, the performance of two different optimizers in obtaining VQE results (under the same conditions as figure 10) will be compared.

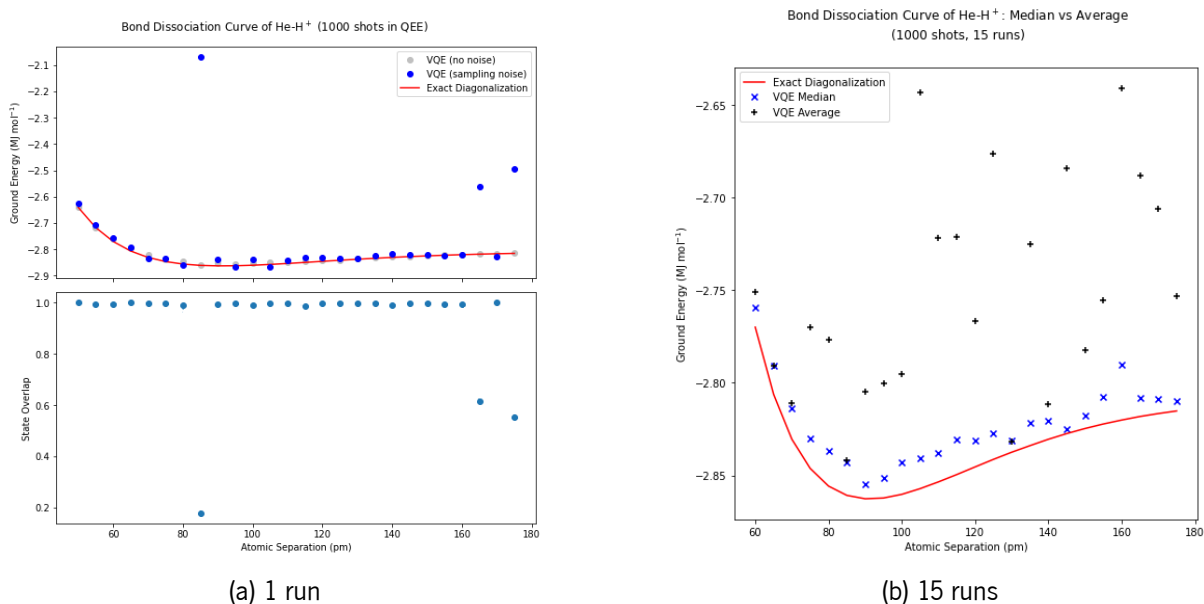


Figure 11: Performance of the VQE algorithm with the COBYLA optimizer along the bond dissociation graph of HeH^+ , with default hyperparameters. The upper plot in figure 11a compares the single-run VQE energy with the exact diagonalization energy, with and without sampling noise (1000 shots). For each interatomic distance, the same random starting point was used for both cases. The lower plot shows the overlap of the VQE state, obtained in the simulations with sampling noise, with the ground state obtained from exact diagonalization. Figure 11b compares the median and the average of the VQE energy over 15 runs with the exact diagonalization value, in the case where 1000 shots are used per energy evaluation. The quantum circuits were run in the CIRQ Simulator [22].

Figure 11 shows the VQE results with the COBYLA (Constrained Optimization BY Linear Approximation) optimization method. Figure 11a shows the single-run performance of this optimizer with the default

hyperparameters in the SciPy [6] implementation. In this case, the relevant hyperparameters are the maximum number of iterations, the final accuracy required for convergence, and the magnitude of initial changes to the variables. Remarkably, even with the default hyperparameters, the optimizer converges to a reasonable approximation to the ground state ($> 90\%$ overlap) in all but three points. It can be seen that sampling noise represents a significant challenge to successful convergence: in the few points whose bad results stand out, the completely noise-free version of the algorithm still outputs an energy that essentially coincides with the exact value, even though the starting point was the same.

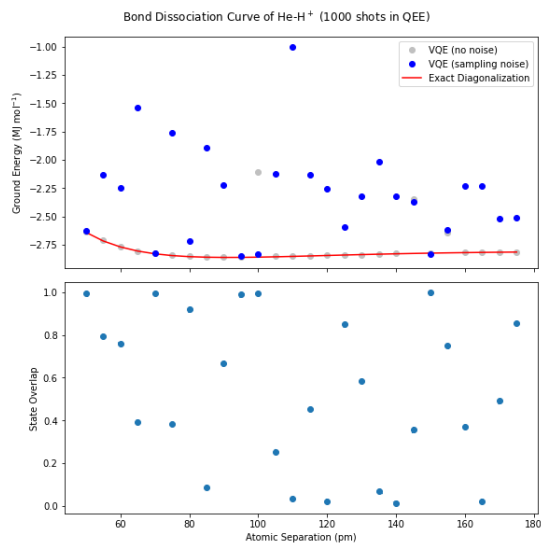
Figure 11b compares the median of the VQE energy against the average for 15 runs, in simulations with 1000 shots in the quantum expectation estimation. While in some cases they almost coincide, it is more frequent that the median is a significantly better estimator. This is caused by aberrant runs in which the optimizer doesn't successfully converge to the ground state: they are promptly excluded by the median, while they still weigh in on the average. The robustness of the median causes it to be a suitable choice here, and this is the estimator that was used in this project whenever the gathered data included more than one run per point.

Taking the median over as few as ten runs proved to be enough for the COBYLA optimizer (with default hyperparameters) to output states with over 90% overlap with the solution throughout the whole bond dissociation curve, which is a solid result considering that the whole Hilbert space is being searched through and the starting state is entirely random.


In figure 12a, we can see that the performance of the Nelder-Mead optimization method is very different. With the default hyperparameters, Nelder-Mead optimized single-run VQE often fails to converge to the ground state even without sampling noise. With sampling noise, this optimizer only successfully converges to the ground state ($> 90\%$ overlap) in six of the points. The noise in the data does not allow statistical filtering to improve upon the single-run behaviour: in figure 12c we can see that neither the average nor the median allow recovering a reasonable approximation to the ground state. In fact, the two estimators are always quite close, suggesting that nothing but noise is being recovered.

This poor performance of Nelder-Mead can greatly improve through an adjustment of the hyperparameters. In this case, the initial simplex is the most relevant. However, the error in parameters/function deemed acceptable for convergence, and the maximum number of function evaluations, may also pose too strict or too relaxed requirements for successful convergence. All of these hyperparameters of the Nelder-Mead optimizer were optimized using COBYLA. In adjusting the initial simplex, a single parameter affecting the simplex size was used. The impact of this 'meta-optimization' is reflected on figure 12b, which shows that each single run posterior to hyperparameter tuning resulted in an overlap with the ground state over 98.6%. In 12d, we can confirm that after hyperparameter optimization, Nelder-Mead VQE allows for statistical treatment: the median is filtering out the aberrant runs, and for each interatomic distance we obtain an energy that is visibly close to the true ground energy. The median VQE energies clearly trace out the bond dissociation curve.


Evidently, conclusions regarding which among these two is the best optimizer cannot be drawn from such scant results. In particular, this example has an overly expressible ansatz that might favour optimizers



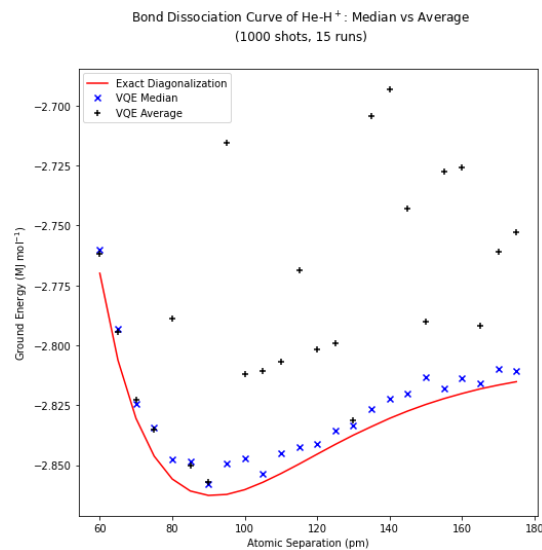
(a) Single run, default hyperparameters



(b) Single run, optimized hyperparameters



(c) 15 runs, default hyperparameters



(d) 15 runs, optimized hyperparameters

Figure 12: Performance of the VQE algorithm with the Nelder-Mead optimizer along the bond dissociation graph of HeH^+ , before (12a, 12c) and after (12b, 12d) hyperparameter optimization. The upper plots in figures 12a, 12b compare the single-run VQE energy with the exact diagonalization energy, with and without sampling noise (1000 shots). For each interatomic distance, the same random starting point was used for both cases. The lower plots show the overlap of the VQE state, obtained in the simulations with sampling noise, with the ground state obtained by exact diagonalization. Figures 12c, 12d compare the median and the average of the VQE energy over 15 runs with the exact diagonalization value, in the case where 1000 shots are used per energy evaluation. The quantum circuits were run in the CIRQ Simulator [22].

with greedier default initial changes to the variables. In fact, it was observed that when narrower ansätze (covering smaller regions of the Hilbert space) were used, the tendency was reversed. In that case, the **COBYLA** optimizer frequently failed unless the initial changes to the variables were greatly decreased.

These examples serve as a demonstration of how the optimization process can affect the performance of **VQE**, and how the choice of optimizer and hyperparameters is of the utmost importance.

3.1.1.4 Running the Algorithm on Cloud Quantum Computers

Up to this point, only *sampling* noise was included in the simulations. Unfortunately, this is not realistic: real quantum computers suffer from many other sources of noise. A brief examination of the impact of generic noise on **VQE** follows.

Figure 13 compares three distinct scenarios: noise-free (13a), sampling noise only (13b), and generic noise (real quantum computer; 13c).

As expected, in the noise-free scenario, the optimizer has no issue in converging to the ground state. It stabilizes the energy around the ground energy, and after an uneventful convergence **VQE** outputs a state whose overlap with the ground state is around 100%. The effect of adding sampling noise (1024 shots) can be seen on figure 13b. The curves show less stability, and oscillate more as convergence is reached. Regardless, the final energy still visibly nears the exact diagonalization value. When the scenario is shifted to a real quantum processor (figure 13c), all types of noise come into play (decoherence; **SPAM** noise; coherent gate errors;...).

The difference between the ideal performance of a quantum processor and a real one is the difference between figures 13b and 13c. When **VQE** is ran using IBM Quantum's Lima device, the energy does not only oscillate: even when convergence is reached, it seems to have a constant offset from the line marking the exact diagonalization energy. This is a symptom of noise-induced barren plateaus. Incoherent noise is capable of decreasing the magnitude of the minimum of the cost function. In the limit that we are left with the fully mixed state by the end of the circuit, the cost landscape is evidently flat. Before that limit is reached, noise will simply prevent the optimizer from ever finding the exact diagonalization ground energy, with a shift that will be the larger the greater the impact of noise.

To better illustrate this, figure 14 shows how the **VQE** optimization changes as the number of shots is increased. The same quantum computer was used (IMBQ's Lima).

In theory, an arbitrary accuracy could be reached by increasing the number of shots. However, these plots suggest that a plateau is reached after a certain point. In figure 14d, we can see that even with a 4-fold increase in the number of shots, the Lima backend cannot perform as well as an ideal quantum processor would (13b). While with only 256 shots (14a) the shot count seemed to be the limiting factor in accuracy, that is no longer the case when it is increased to 8192 shots. At this point, the energy seldom oscillates from iteration to iteration, and it is evident that the issue goes beyond sampling noise. Evidently, increasing the shot count only allows us to go so far. A limit will be reached where only improving the characteristics of the quantum computer, or resorting to a quantum error correction will further improve

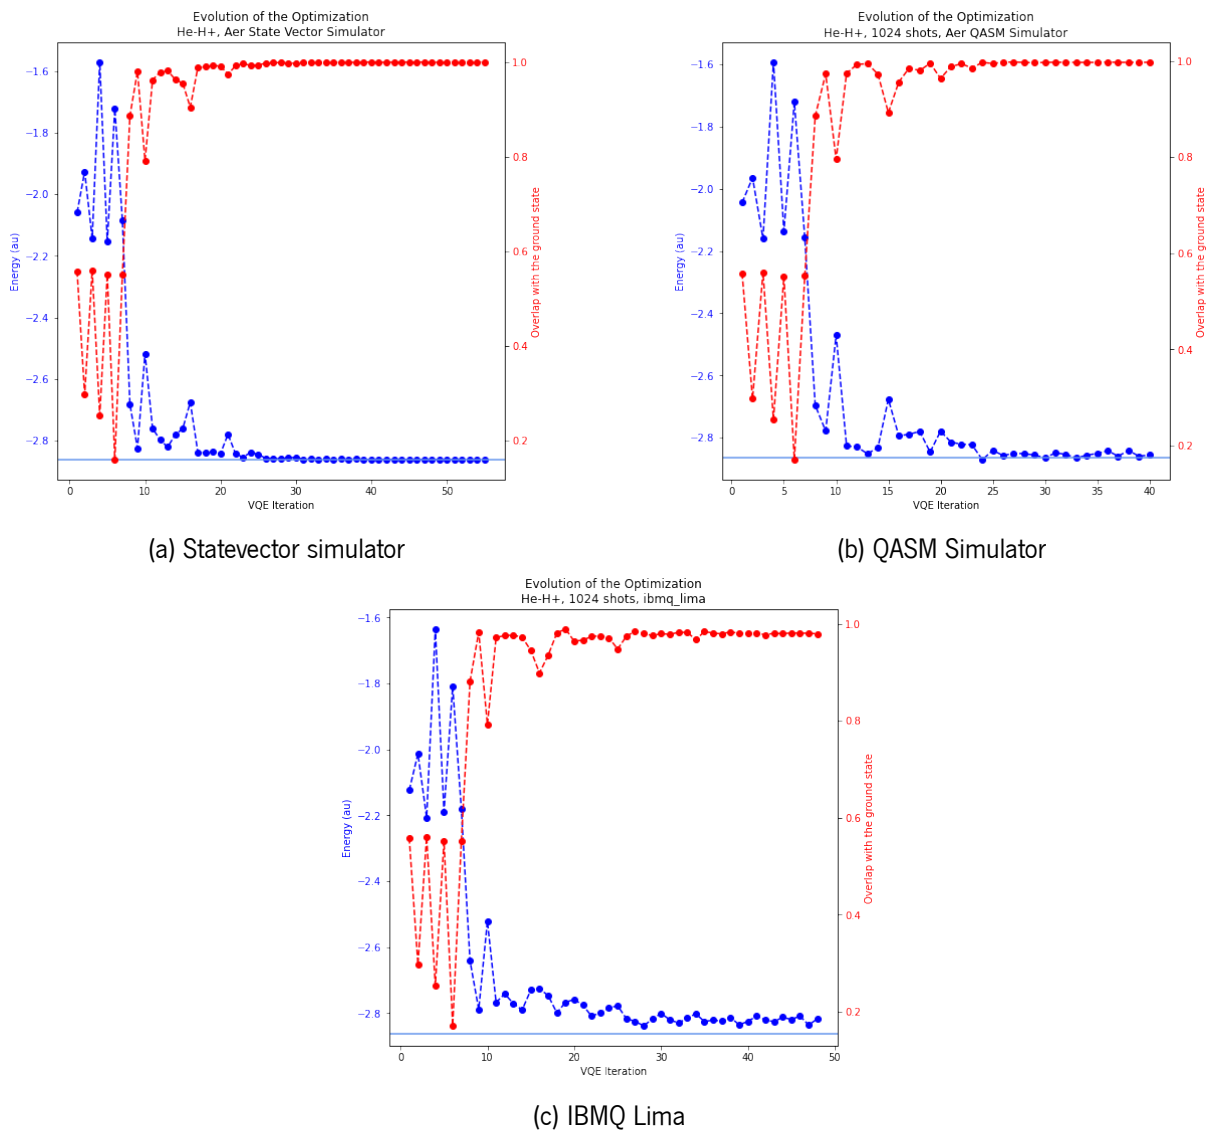


Figure 13: Evolution of the VQE optimization for HeH^+ at an interatomic distance of 90pm, using three different IBM Quantum [42] backends: the statevector simulator (with no noise of any kind), the QASM simulator (with sampling noise) and the Lima device (a 5-qubit Falcon processor). In the last two cases, 1024 shots were used. The energy is plotted in blue and the overlap with the true ground state in red; the pale blue line marks the true ground energy. The same random starting point was used for all backends. The chosen optimizer was COBYLA.




Figure 14: Evolution of the VQE optimization for HeH^+ at an interatomic distance of 90pm, for different shot counts and using the Lima device. The remaining conditions (starting point, optimizer) are the same as in figure 13.

accuracy.

3.2 Problem Tailored Ansätze

Leveraging problem-specific knowledge to design the ansatz allows narrowing the search space by excluding regions that are not expected to contain the solution. Such techniques reduce the expressibility of the ansatz without abdicating precision, because the solution is still contained within the variational form. Since problem-tailoring is of great importance in preventing barren plateaus, and this trainability issue renders variational quantum algorithms inefficient, a great part of the research surrounding VQAs is aimed at developing precise and shallow ansätze of this type.

This section contains the results obtained from applying VQE to molecular hydrogen (H_2) both in simulators and real quantum computers, using a problem-tailored ansatz. The employed ansatz was UCCSD, introduced in section 2.2. As was explained then, this variational form arises in the context of conventional CC theory as a natural modification to the non-variational wave form of the latter, which in turn is the natural size-consistent successor to the Configuration Interaction variational form. While UCCSD is not classically tractable, the unitarity that forces variationality allows for convenient implementation in quantum circuits. Due to being unitary and having solid roots in older computational chemistry methods, the UCCSD ansatz has become a popular choice in VQAs addressing chemistry applications.

3.2.1 Application to Molecular Hydrogen

The UCCSD-VQE algorithm was implemented in this project using CIRQ [22] and Qiskit [3], alternative software libraries for manipulating and running quantum circuits. Tasks such as obtaining the Hamiltonians and performing fermion-to-qubit mappings were done via OpenFermion [2] in the first case and Qiskit's chemistry module in the second. In both, the underlying electronic structure information was extracted from PySCF [5]. Additionally, an independent noise-free UCCSD-VQE based on matrix algebra was implemented to verify correctness.

Simulations were done to analyse the effect of sampling noise on UCCSD-VQE, as well as of CNOT gate errors, and of thermal relaxation and SPAM errors, using noise models created in Qiskit. Density matrix simulations were employed to analyse the purity of the final state, using noise models created from the specifications of IBMQ [42] quantum processors to mimic their behaviour.

3.2.1.1 Bond Dissociation Graph

Much like in the previous section, we will start with the bond dissociation plots to get a general sense of performance before delving into the evolution of individual optimizations.

In figure 15, we can see the VQE energies tracing out the bond dissociation curve of H_2 . Here, the simulations include no noise other than sampling.

The effect of the number of shots clearly shows in the results, with the VQE energy getting closer and closer to the exact diagonalization curve from figure 15a to 15d. In the limit of infinite shots, the VQE energy essentially matches the FCI result (the observed error was of the order of 10^{-9}).

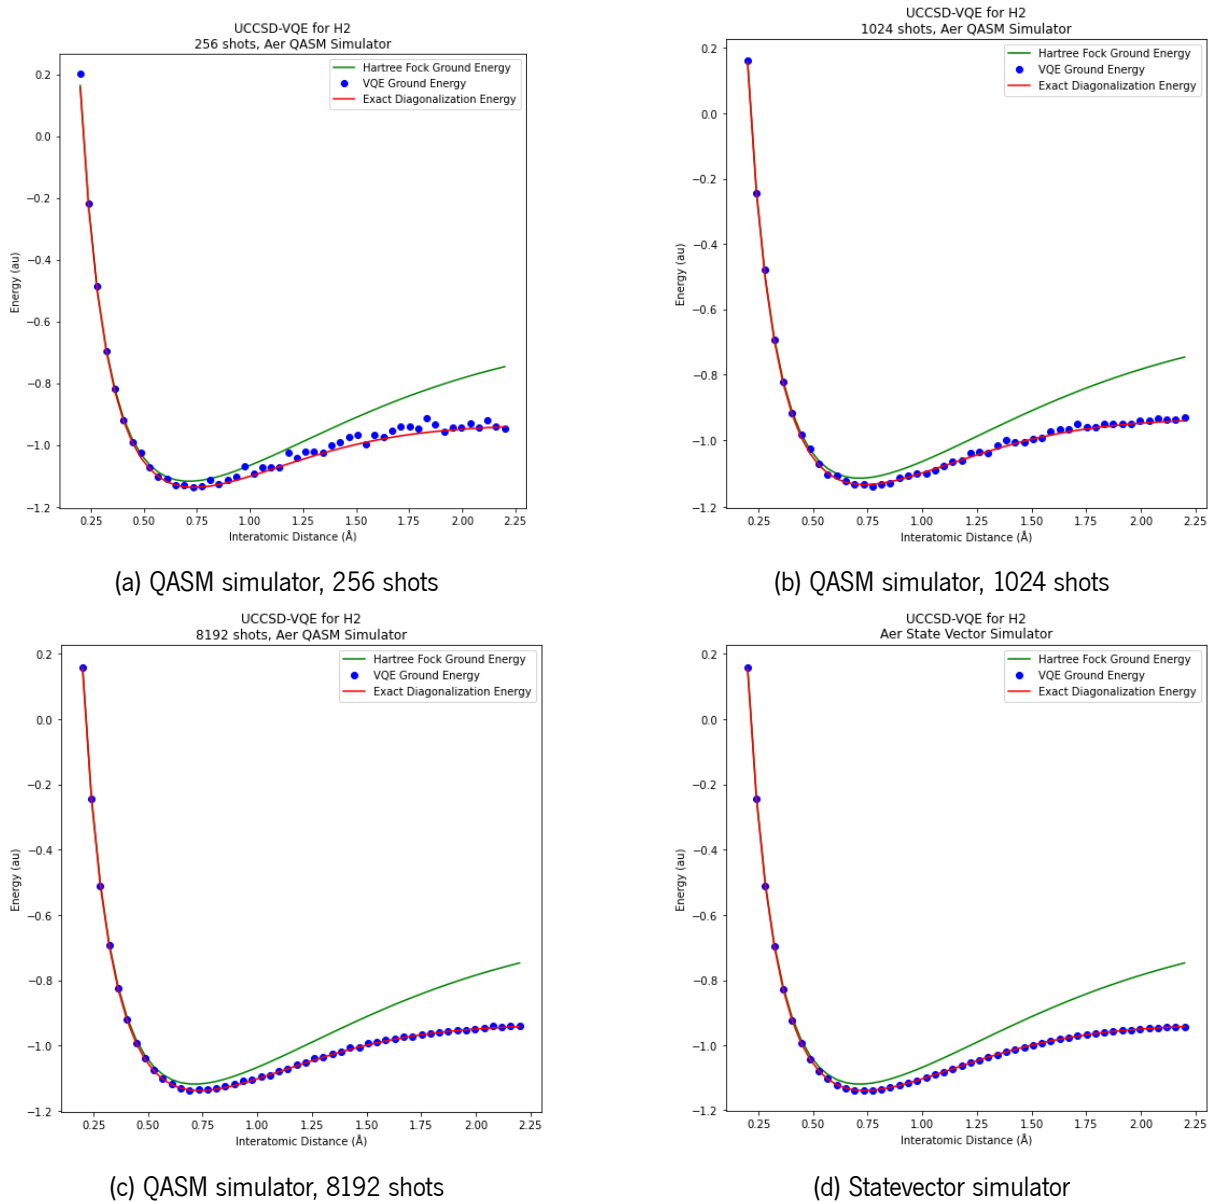


Figure 15: Single-run [UCCSD-VQE](#) ground energy along the bond dissociation curve of the hydrogen molecule, obtained using IBMQ [42] simulators. The results plotted in figures 15a-15c, ordered by increasing shot count, were obtained using the QASM simulator. The results plotted in figure 15d were obtained using the state vector simulator (equivalent to infinite shots). The red and green curves represent the [FCI](#) energy and the Hartree-Fock energy, respectively.

It is interesting to compare the VQE energy with that obtained by performing the Hartree-Fock self-consistent field calculations. The state resulting from these calculations is the reference state in VQE, upon which the UCCSD ansatz performs a ‘correction’. It becomes evident that, in the presence of sufficiently weak noise, optimizing the UCCSD variational form is enough to bring the energy away from the Hartree-Fock approximation and into the FCI value. The variational freedom takes care of electronic correlation effects that go unaccounted for in the Hartree-Fock approximation.

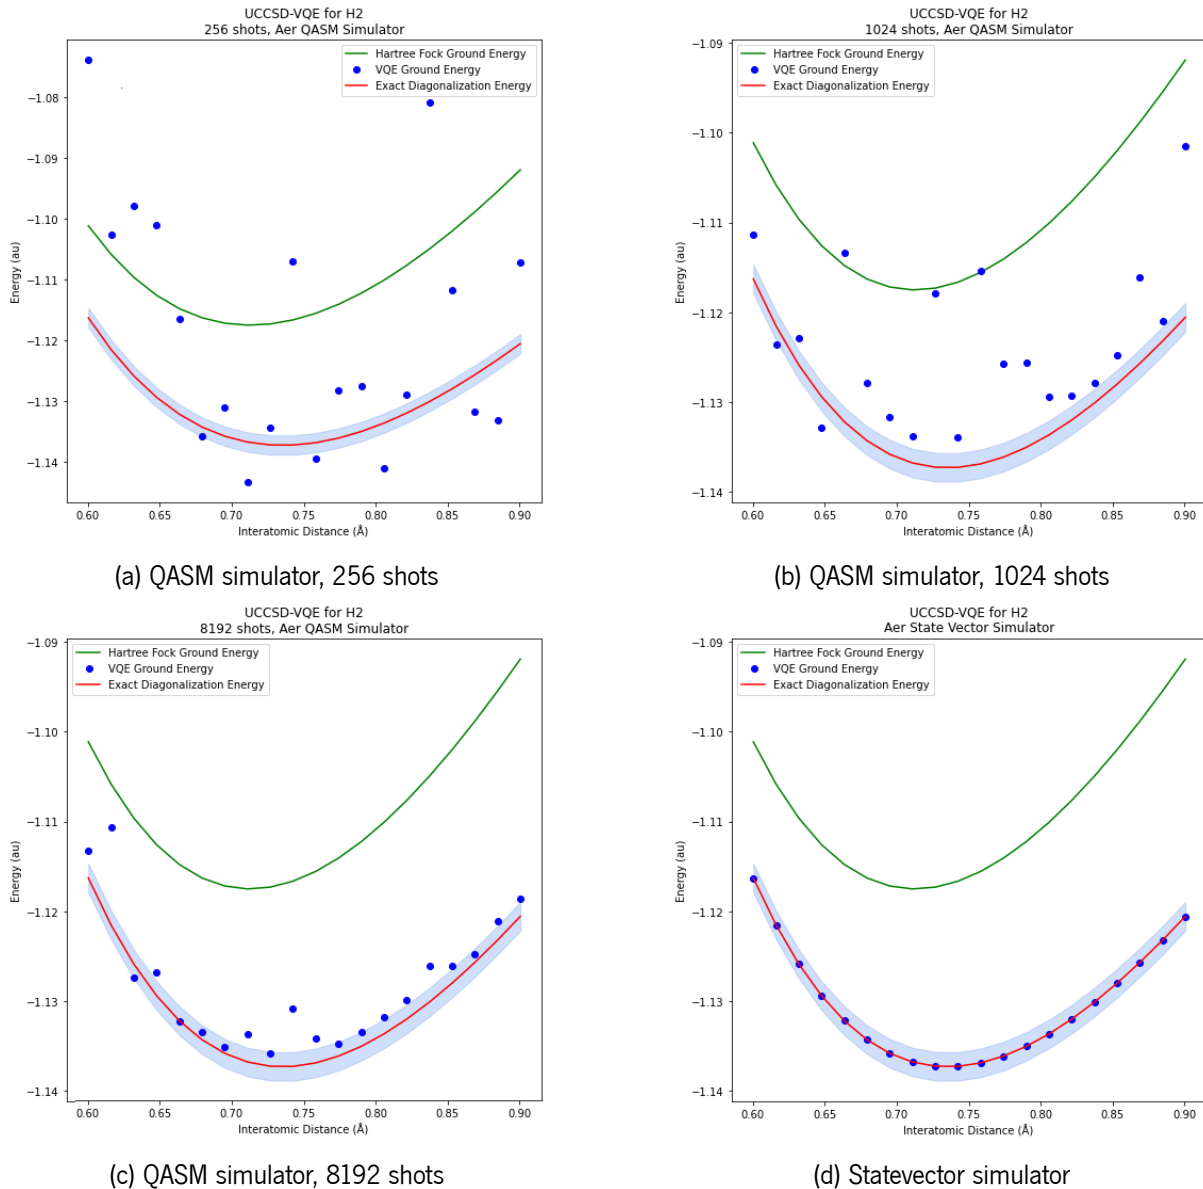


Figure 16: Bond dissociation curves of figure 15, zoomed in on the region nearing the equilibrium bond length. The blue area marks the region of chemical accuracy (error of less than 1kcal/mol).

To better analyse the results, figure 16 shows the same plots as figure 15, now magnified around the equilibrium bond length (the minimum of the curve). In these plots, it is visible that reaching chemical accuracy requires a considerable number of shots. With only 256 shots (figure 16a), UCCSD-VQE is often

outperformed by the Hartree-Fock approach; when this happens, we know that the UCCSD correction didn't at all improve upon the reference state, and the optimization was entirely in vain. 1024 shots (figure 16b) seem to be enough for UCCSD-VQE to consistently beat the Hartree-Fock method, but the single-run energy is still far from reaching chemical accuracy. In figure 16c, we can see that with 8192 shots, the energy is already often inside or nearing the region of chemical accuracy, even with a single run (taking the median over several runs would further refine the result). Of course, in the limit of infinite shots (figure 16d), the performance is nearly perfect.

3.2.1.2 Running the Algorithm on Cloud Quantum Computers

So far, the simulations have included at most sampling noise - but they all assumed ideal behaviour from the quantum processor. In the previous section, we saw that the performance of the algorithm in real backends could prove significantly worse because of the other types of noise they inevitably suffer from. With UCCSD the difference becomes even more relevant, because the circuit is deeper: the 2-qubit ansatz from the previous section had a single CNOT and a total of only 7 gates. Even for H_2 , the simplest molecule possible, implementing the UCCSD operator requires significantly longer circuits.

Figure 17 showcases the evolution of the VQE optimization for a single point (interatomic distance) in three different scenarios: noise-free (figure 17a), sampling noise only (figure 17b), real quantum processor (figure 17c). Here, the optimization of a single run was plotted instead of the final results of multiple runs (as in the bond dissociation graphs) because of the time required to run hybrid algorithms on cloud quantum computers with fair-share queuing.

The difference between the noise-free and sampling noise only scenarios is already remarkable. While in the former the energy is fast to stabilize inside the region of chemical accuracy, in the latter the energy oscillates enough that it's in and out of the region throughout the whole optimization. Even towards the last iterations, the energy does not meet chemical accuracy consistently: it is outside the region in the second-to-last iteration.

However, if sampling noise is significant, it does not come close to the impact of other types of noise. Figure 17c shows that running the algorithm in IBMQ's Belem backend, a real quantum processor, does not allow recovering any valid results. The VQE output energy oscillates about 1 Hartree over the true ground energy, without being remotely close to reaching chemical accuracy (of the order of the millesimals of Hartree). This is a sign that quantum information has been washed out by noise. The UCCSD ansatz for H_2 , obtained using Qiskit's variational form and transpiled onto Belem, has 130 CNOTs; the circuit depth is approximately proportional to the CNOT count. The results show that the circuit in question is too deep for any meaningful data to be recovered by the end.

An approximate density matrix representing the state at the end of the noisy UCCSD ansatz was obtained by simulating the algorithm on the QASM simulator, now with a noise model aiming to mimic the behaviour of the Belem processor. The noise model was obtained from Qiskit [3] and is built from data concerning the specific backend. This procedure allowed obtaining estimates to the purity of the final state

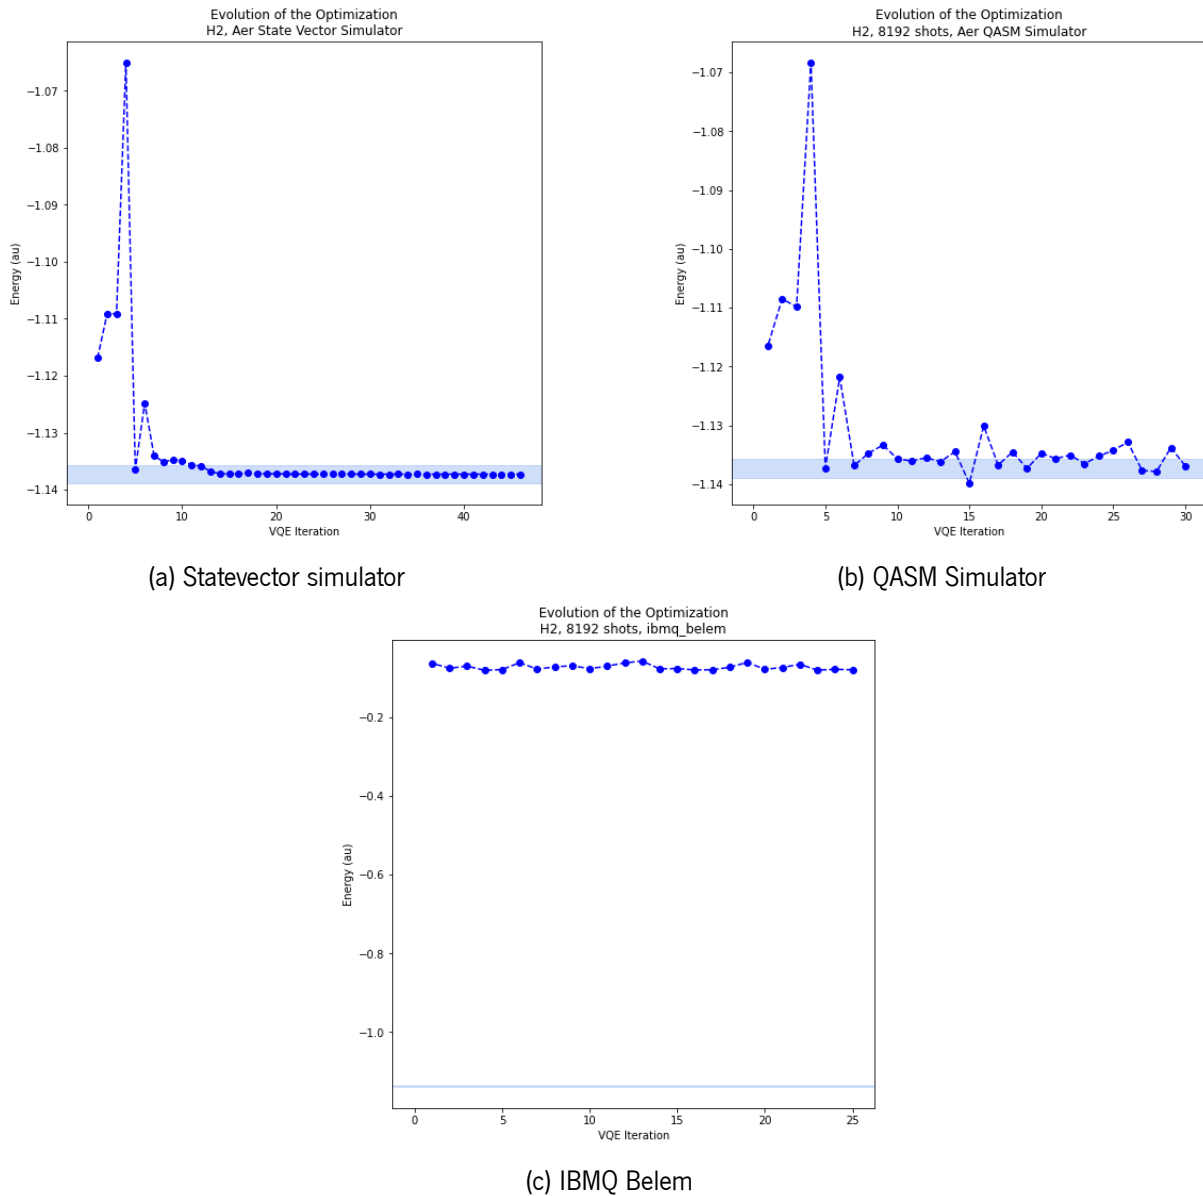


Figure 17: Evolution of the [UCCSD-VQE](#) optimization for H_2 at an interatomic distance of 0.74 \AA , using three different IBM Quantum [42] backends: the statevector simulator (with no noise of any kind), the QASM simulator (with sampling noise) and the Belem device (a 5-qubit Falcon processor). In the last two cases, 8192 shots were used. The energy is plotted in blue; the pale blue region marks the area of chemical accuracy. The same starting point, the Hartree-Fock reference state, was used for all backends. The chosen optimizer was [COBYLA](#).

and its fidelity with the ideal noiseless output.

The fidelity of the final state with the ideal state was found to be only 0.081, in contrast with its fidelity with the fully mixed state, which was 0.40. The final purity of the state was found to be around 0.16. The STO-3G basis set was used; in a minimal basis set (as is the case), H_2 has four spin-orbitals, so that fermionic states of the molecule were mapped into states of four qubits. This implies a Hilbert space of dimension 16, so that the purity is lower bounded by 0.0625.

The expectation value of the energy calculated in the fully mixed state was found to be -0.097 Hartree. In figure 17c, we can see that the energy of the state at the end of the UCCSD circuit in the Belem backend remains close to that value throughout the whole optimization. The optimizer spends 25 iterations trying to adjust the parameters that will lead to a minimum of the energy, but it is to no avail: the energy shows no more than slight deviations from the mean value, with no sign of approaching the ground energy.

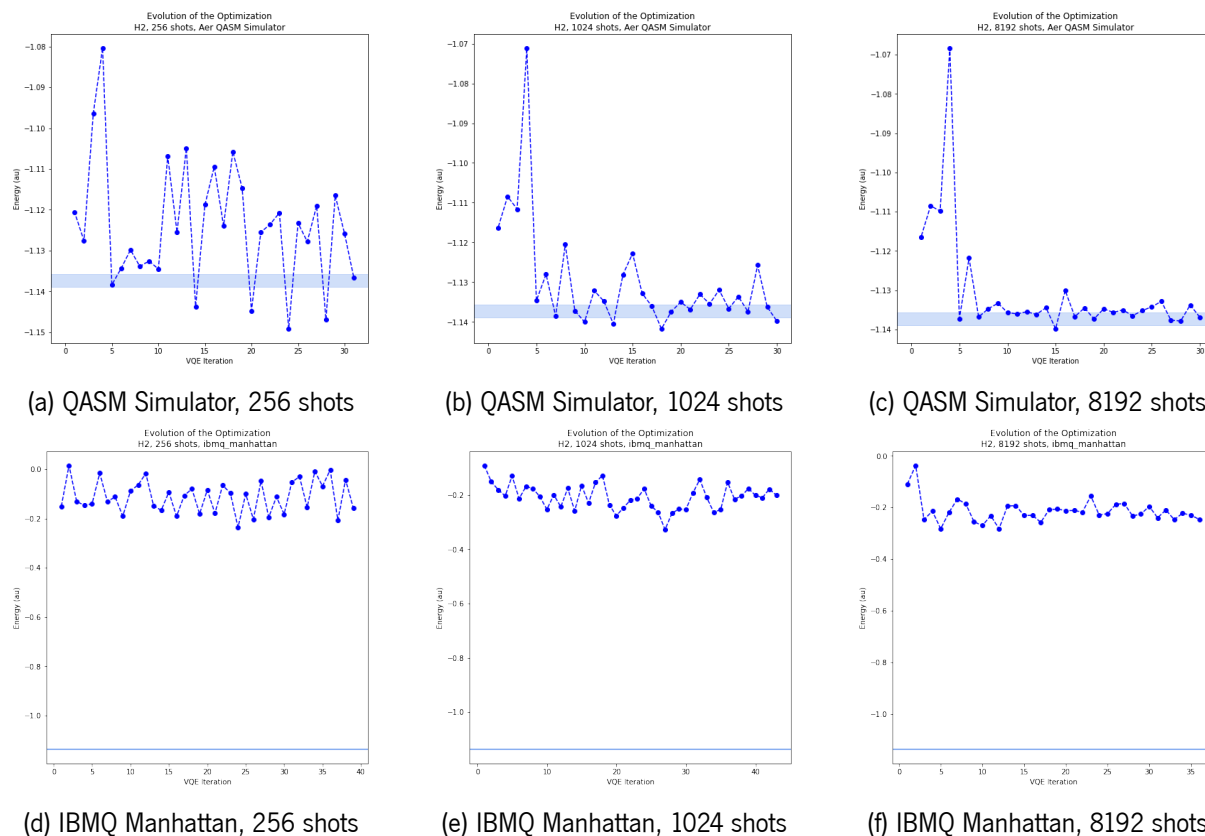


Figure 18: Evolution of the VQE optimization under the same conditions of figure 17 (same molecule, interatomic distance and optimizer), plotted for three different shot counts and two different backends. The QASM simulator (18a-18c) portrays the behaviour of an ideal quantum processor, in contrast with IBMQ Manhattan (18d-18f), a 65-qubit Hummingbird processor.

From figure 17, we can already conclude that sampling noise is far from being the limiting factor here. To further illustrate this, figure 18 compares VQE optimizations using the QASM simulator (with sampling noise only) and using the Manhattan backend (another real device) for 256, 1024, and 8192 shots.

The improvement of the VQE performance with increased shot count is evident when the QASM simulator is used: the energy is visibly faster to stabilize. This is to be expected, since here sampling is the only source of noise. With 256 shots only (figure 18a) the energy shows changes of the order of the centesimals of Hartree up until the last iterations. With 8192 shots (figure 18c) the energy change is of the order of the millesimals of Hartree in every iteration from the seventh to the last one (30).

When the Manhattan processor is used, the energy shows the same tendency for stabilization: from 256 shots (figure 18d) to 8192 shots (figure 15c), there is also a decrease in the order of magnitude of the oscillations in the energy towards the last iterations. However, here, this does not translate into any significant improvement in accuracy, because the energy is still very far from the ground energy, with an error of around 1 Hartree. Much like it happened with the Belem backend, the state at the end of the UCCSD circuit in the Manhattan backend is too corrupted for the optimizer to converge to the ground state: the energy evaluations are too far from the exact value. The purity is so low that changing the parameters is barely reflected on the expectation value. We can conclude that the circuits are too deep to be implemented in these quantum computers.

3.2.1.3 Analysing Noisy Instructions

To better understand which sources of noise were most affecting the algorithm, the noise model emulating the behaviour of the Belem backend was decomposed into several noise models, each considering only the noise associated with a specific instruction. UCCSD-VQE was then simulated with each of the noise models; the results are plotted in figure 19.

The curve that stands out is the one associated with the noise model that isolated CNOT noise. Even though this model ignored all other sources of noise, the performance is remarkably close to that of the complete Belem noise model (figure 17c). In here, the error in the final energy is once again close to 1 Hartree; further, the energy barely moves away from the starting energy throughout the whole optimization. One again, the impact of noise is felt too strongly for any meaningful results to be recovered.

It is not surprising that this is the noisiest gate: it is the only two-qubit gate. In this particular backend (Belem), the average CNOT gate time is around 550 nanoseconds, while the single-qubit gate time is typically of the order of tens of nanoseconds. Among the available gates, the CNOT thus corresponds by far to the biggest cost in circuit depth.

Since the noise associated with CNOT gates is the one that most impacts UCCSD-VQE, it will be interesting to analyse the relation between the magnitude of CNOT gate errors and the performance of the algorithm. A graphical representation of the impact of the average CNOT gate error on the final UCCSD-VQE error can be found in figure 20. To obtain these results, a noise model was created to simulate CNOT gate errors. This model consisted of a two-qubit depolarizing error, succeeded by single qubit thermal relaxation errors on the involved qubits.

This is a similar approach to the one employed in Qiskit's backend noise models, mentioned previously and used in obtaining the results presented in figure 19. In these noise models, made available by the Qiskit




Figure 19: Impact of the noise associated with different instructions on the performance of [UCCSD-VQE](#), using Belem’s noise model from Qiskit’s [\[3\] providers.aer.noise](#) module. The model is generated automatically from the properties of IBMQ’s [\[42\]](#) device. Each curve represents the [UCCSD-VQE](#) optimization, simulated with a model including the noise associated with the corresponding instruction exclusively. 8192 shots were used in all curves. The conditions are the same as those from figure [17](#) (same molecule, interatomic distance and optimizer).

package, single qubit thermal relaxation errors are created for both qubits using the relaxation parameters of the backend, and a two-qubit depolarizing error is then added so that the total gate error matches the experimental value. In our case, the purpose was precisely to vary the gate error, which prevented us from using it to dictate the characteristics of the depolarizing channel. Instead, the ratio between the gate error due to the depolarizing channel and due to thermal relaxation was assumed constant. This constant was set to the average value of this ratio for a real device, Belem (calculated from available data). For simplicity, all [CNOT](#) gates were assumed to behave identically.

Figure [20](#) suggests that in the case of H_2 , the relation between the average [CNOT](#) gate error and the error in the final energy (in atomic units) is roughly linear in the interval of values contemplated in the plot. When the magnitude of the [CNOT](#) error is of the order of 10^{-2} , as is the case for the quantum processors made available on the cloud by IBMQ [\[42\]](#), the error can reach almost 1 Hartree. For the error in the [UCCSD-VQE](#) energy for this molecule to be under 0.1 Hartree, the [CNOT](#) gate error must be of the order of 10^{-4} .

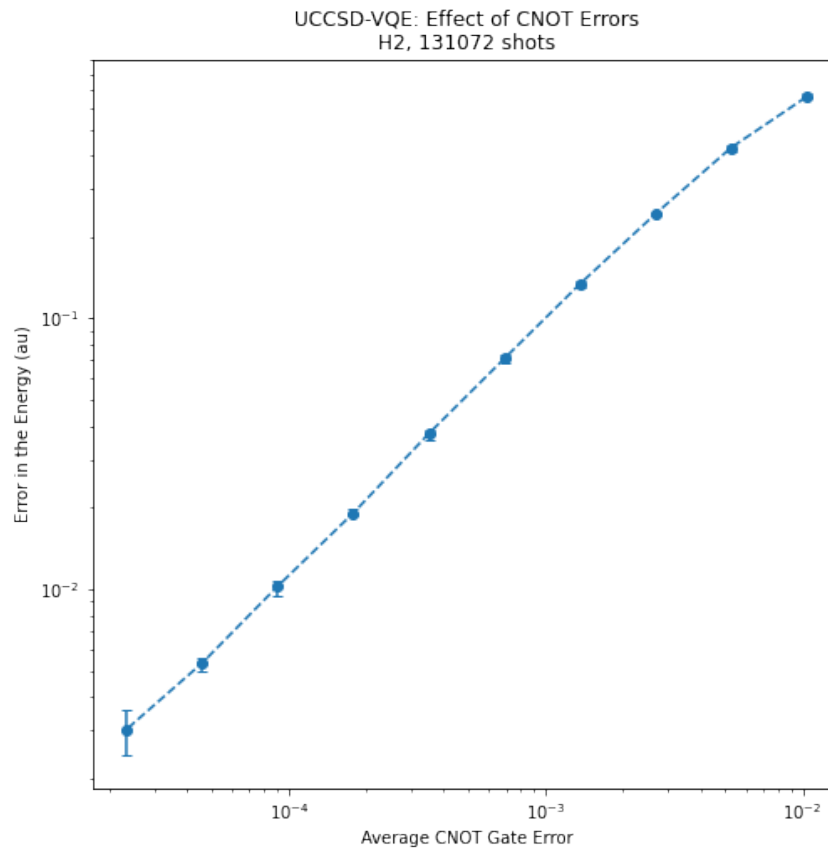


Figure 20: Impact of **CNOT** gate noise on the performance of **UCCSD-VQE**. The error in the final energy is plotted as a function of the average **CNOT** gate error (equal to $1 - F$, where F is the average gate fidelity). Twenty runs were used per point; the median was used as the estimator, and the error bars show the interquartile ranges. The number of shots was increased to ensure that sampling noise would not be the limiting factor for any gate error. The conditions are the same as those from figure 17 (same molecule, interatomic distance and optimizer).

The noise associated with each of the remaining instructions in figure 19 has a significantly lower impact on the algorithm than the noise associated with the **CNOTs**. The second worst performance, already quite better, was associated with the noise model that isolated measurement errors. Such errors could be partly helped by employing measurement error mitigation techniques. The remaining instructions considered - namely, reset (preparing the qubits in state $|0\rangle$), the identity operation, and the **SX** and **X** gates - are associated with significantly less noise. The final **UCCSD-VQE** energies obtained with noise models contemplating each of these instructions are exceptionally close the **FCI** value as compared to those isolating **CNOT**-related noise (mostly) and measurement-related noise (to a lesser degree).

3.2.1.4 Impact of Thermal Relaxation, SPAM errors, and Sampling

Figure 19 analysed noise by dividing it into the instructions it is associated with. However, it is also interesting to analyse it by nature - as it was briefly covered in 2.1.4, there are different types of noise affecting quantum computers. That is the intent of figure 21: it shows how scaled down thermal relaxation effects

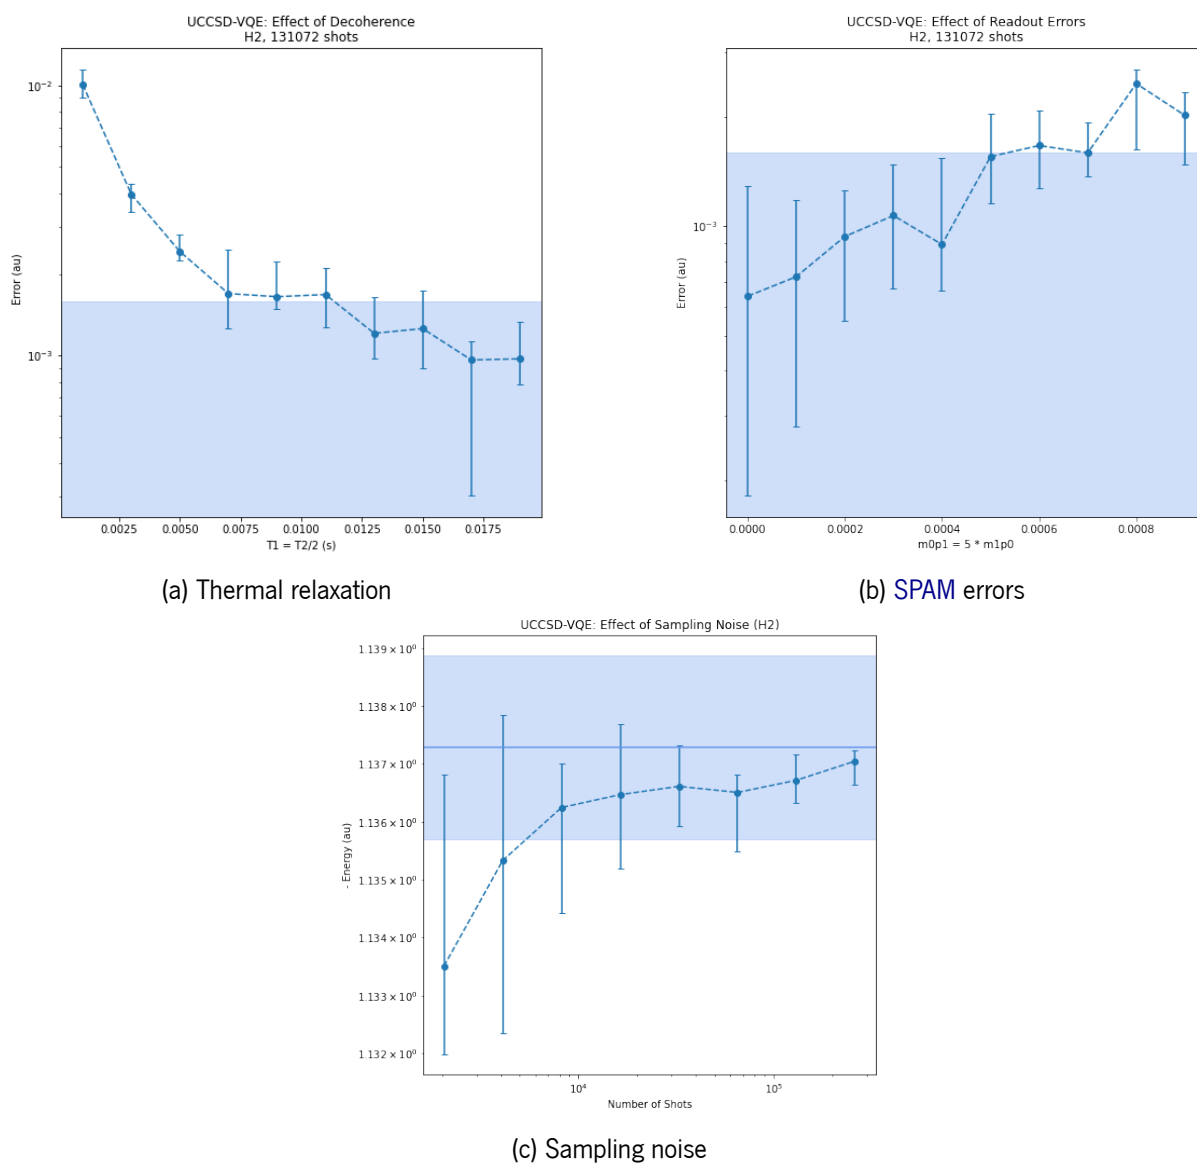


Figure 21: Plots showing the impact of different types of noise on the performance of UCCSD-VQE. For thermal relaxation, a simple model assuming average gate times and a constant ratio between the coherence times $T1$ and $T2$ was used. Figure 21a shows the evolution of the error in the energy as a function of these times. For SPAM noise, the ratio between the probability of measuring 0 after preparing 1 and the probability of measuring 1 after preparing 0 was fixed (and greater than one, for the model to be realistic). Figure 21b shows the evolution of the error as a function of these probabilities. Figure 21c illustrates the impact of sampling noise by plotting the error as a function of the number of shots. In all of the plots, twenty runs were used per point; the median was used as the estimator and the error bars show the interquartile ranges. All noise models were created using Qiskit [3]. The conditions are the same as those from figure 17 (same molecule, interatomic distance and optimizer).

(from interaction with the environment), **SPAM** errors (from faulty state preparations and measurements), and sampling noise (from using a finite number of shots) should individually be for chemical accuracy to be confidently reached. In each case, all noise sources other than the one being analysed were ignored, except for sampling noise, that was included in each of the simulations due to being intrinsic to quantum computing. When fixed, the number of shots was set to 2^{17} , chosen to be high enough not to impede chemical accuracy.

Figure 21a shows the impact of thermal relaxation on **UCCSD**, by plotting the final energy error as a function of T1 (the relaxation time) and T2 (the dephasing time). The process of thermal relaxation leads to the loss of quantum coherence (*decoherence*), causing quantum information to be irreversibly lost. Currently, on the quantum computers made available on the cloud by IBMQ, the coherence times are of the order of 10^{-4} seconds. In contrast, **UCCSD-VQE** seems to require them to be of the order of 10^{-2} seconds to reach chemical accuracy in the middle 50% runs (the runs yielding the energies between the first and the third quartiles).

Figure 21b showcases the effect of **SPAM** errors. Chemical accuracy was secured in all runs within the interquartile range only when the probability of measuring $|0\rangle$ after preparing $|1\rangle$ reached 4×10^{-4} (with the probability of measuring $|1\rangle$ after preparing $|0\rangle$, made proportional, at 8×10^{-5}). For reference, the recently used Belem backend has average error probabilities around 0.04 and 0.009, respectively - which amounts to around 100 times more than that.

Finally, figure 21c shows the effect of sampling noise, plotting the error as a function of the number of shots. The simulations showed that for the 50% center runs to fall inside the region of chemical accuracy, the number of shots would have to be roughly 80000. Currently, IBMQ limits the shot count on cloud quantum computers to 8192.

Concluding, when running **UCCSD-VQE** on IBMQ cloud quantum computers, chemical accuracy is beyond reach. Achieving it would require a significant improvement of readout and state preparation errors, thermal relaxation effects, etc. Even the shot count necessary for such an accuracy is 10-fold what is allowed; of course, this is not a relevant hardware limitation because it would not exist if a dedicated quantum processor was available. Regardless, it is a good show of how demanding chemistry problems can be.

It is also worth noting that, throughout the section, we were dealing with the smallest possible molecule. For bigger molecules, the situation is even worse: they would require more qubits, more shots for the same accuracy (see 2.3.3.4), deeper circuits, and more variational parameters (leading to a harder optimization).

Problem Tailored, Dynamically Created Ansätze: ADAPT-VQE

If predetermined, problem-tailored ansätze already bypass several of the issues that problem-agnostic ansätze suffer from, dynamic ones reach farther. They are not simply problem-tailored, they are *system-tailored*, with all the advantages that arise from that.

This section will introduce and compare two similar dynamically created ansätze (fermionic-ADAPT-VQE and qubit-ADAPT-VQE), and compare the noise resilience and measurement cost of ADAPT-VQE against UCCSD-VQE.

The ADAPT-VQE algorithm was implemented in CIRQ [22], in Qiskit [3], and independently using matrix algebra. In all cases, the OpenFermion [2] software package was used for creation and manipulation of fermionic operators, along with the corresponding plugin with PySCF [5], an electronic structure package. The Qiskit implementation used the VQE class from the Aqua module with a personalized ansatz, while in the CIRQ implementation not only the ansatz but also the Hamiltonian averaging were done from scratch, with the optimization loop having also been implemented independently up to SciPy's [6] minimization function. Qiskit's noise module was used for the creation of noise models meant to analyse the impact of noise in the algorithms via simulations.

4.1 The ADAPT-VQE Algorithm

4.1.1 Fermionic-ADAPT-VQE

In 2019, ADAPT-VQE was introduced in [35] with the purpose of creating a more accurate, compact, and problem-customized ansatz for VQE. This version will hereby be denoted fermionic-ADAPT-VQE to distinguish it against a more recent version that will be covered shortly.

The idea behind the proposal is to let the molecule in study 'choose' its own state preparation circuit, by creating the ansatz in a strongly system-adapted manner.

Figure 22 shows a schematic outline of the algorithm. The only human role contributing to the creation of the ansatz is the selection of an *operator pool*. Any operator in this pool may or may not be a part of the final ansatz, in its exponentiated and parameterized version. In the numerical simulations presented in the original article, SCGSD excitations were used. In this context, much like in the unitary version of CC theory, *excitations* is used to mean anti-Hermitian sums of excitation and de-excitation operators, that are mapped to parameterized unitary operators upon exponentiation. The convenience of this form goes beyond unitarity, as will be seen shortly.

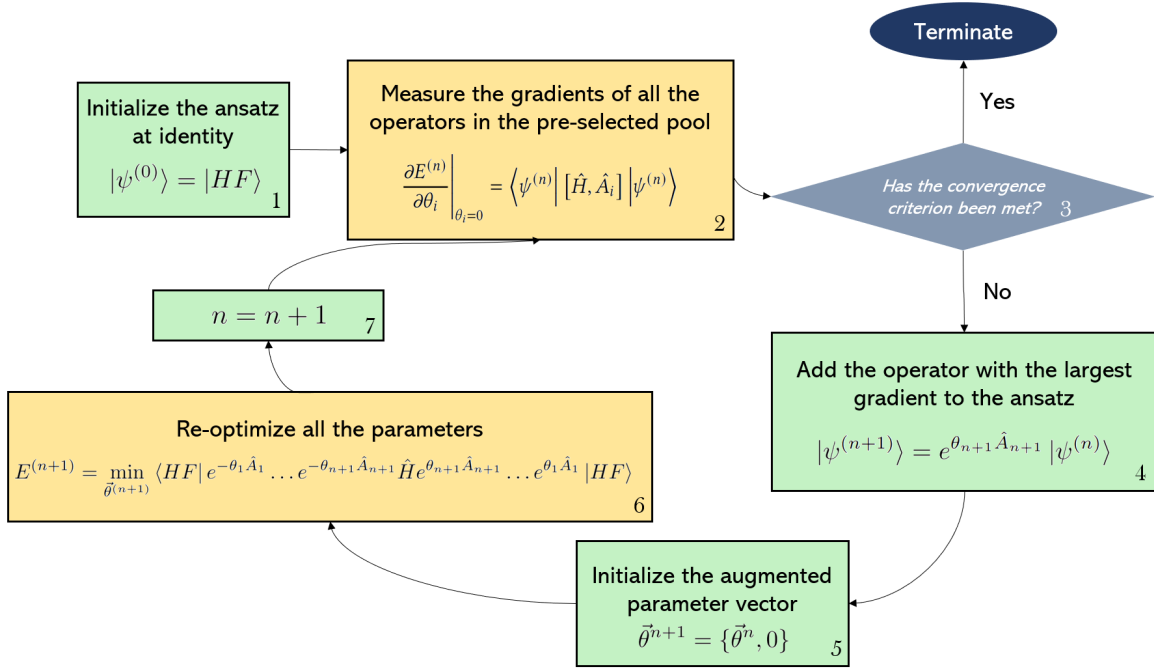


Figure 22: Scheme illustrating the **ADAPT-VQE** algorithm. The green boxes represent processes that are exclusively classical. The orange boxes represent hybrid processes (that include state preparations and measurements in the quantum computer).

The ansatz is initialized at identity: in the first iteration, the prepared state is simply the reference Hartree-Fock state. Each iteration adds an operator to the ansatz, along with the corresponding variational parameter, initialized at zero. Thus, the state preparation circuit and the parameter vector grow from iteration to iteration.

The criterion for operator selection evidently must tend to choose operators that are more beneficial for the wave function (i.e. lower the energy more) at a given point. Similarly to a wide class of optimization methods, the *gradient* is used as the indicator.

In the beginning of an iteration n , **ADAPT-VQE** determines the instantaneous rate of change of the energy with respect to each candidate for variational parameter, associated with each candidate for operator, evaluated at the point where the parameter will be initialized (zero). The selected operator is the one which, if added along with the corresponding variational parameter θ_i , has the derivative $\frac{\partial E^{(n)}}{\partial \theta_i}$ of greatest magnitude.

The question is, how can these derivatives be evaluated? Luckily, the convenient form of the operators offers a solution. For each aspirant ansatz operator $e^{\theta_i \hat{A}_i}$ associated with a pool operator \hat{A}_i , we can calculate the energy slope at point $\theta_i = 0$ as

$$\left. \frac{\partial E^{(n)}}{\partial \theta_i} \right|_{\theta_i=0} = \langle \psi^{(n)} | [\hat{H}, \hat{A}_i] | \psi^{(n)} \rangle. \quad (4.1)$$

A sketch of the proof can be found on appendix C.

The only thing left to explain is when the algorithm terminates: in order to know when to stop adding operators and growing the ansatz, and finally output an attempted solution, a *convergence criterion* is necessary. There may be different possibilities, but the one employed in the original article (and reproduced in this project) was to terminate when the total gradient norm falls below a certain threshold ϵ . At the beginning of each iteration, after calculating all of the derivatives, they are used to calculate this norm. If the norm is still over the threshold, the operator with the highest derivative is added to the ansatz; if not, the algorithm terminates. Evidently, the lower ϵ is, the stricter the convergence requirement, the larger the final ansatz, and the higher the precision.

We can then summarize the [ADAPT-VQE](#) algorithm in 7 steps (as labeled in figure 22).

1. Initialize the ansatz at identity, so that the initial state is simply the Hartree-Fock solution. The corresponding state preparation circuit consists of N parallel Pauli X gates, where N is the number of electrons.
2. Measure the gradient of each operator in the pool on the quantum computer. This can be done using formula 4.1.
3. Calculate the square norm of the gradient vector. If it is under a predetermined threshold ϵ , terminate. If not, proceed to the following step.
4. Using the information obtained in step 2, select the operator with the largest gradient and append it to the ansatz.
5. Initialize the parameter vector with the values from the previous iteration. The newly added variational parameter should be initialized at zero.
6. Re-optimize *all* the parameters. This is a typical [VQE](#) optimization, with the current ansatz.
7. Move on to the next iteration, going back to step 2.

It should be noted that before these steps take place, one must have done the necessary classical computations (common to the original [VQE](#)) in preparation, to obtain the Hamiltonian operator and the Hartree-Fock spin-orbitals. Further, one must have chosen the constitution of the pool, and the convergence threshold ϵ .

Steps 2 and 6 require preparing the [ADAPT-VQE](#) state multiple times, and measuring the expectation value of the Hamiltonian \hat{H} as well as of the commutators $[\hat{H}, \hat{A}_i]$. A brief explanation on how this can be done follows.

In iteration n , the [ADAPT-VQE](#) state is written $|\psi^{(n)}\rangle = e^{\theta_n \hat{A}_n} \dots e^{\theta_1 \hat{A}_1} |HF\rangle$. Under the Jordan-Wigner transform 2.57 the excitation operators \hat{A}_i are mapped to linear combinations of Pauli strings with imaginary coefficients. The operators $e^{\theta_i \hat{A}_i}$ can then be implemented as explained in subsection 2.1.3. Along with the preparation of the Hartree-Fock state presented in 2.3.3.2, this tells us how to create the state preparation circuit.

The measurements are also straightforward: the Hamiltonian \hat{H} also consists of a linear combination of Pauli strings, so that its expectation value can be obtained using the procedure outlined in subsection 2.1.2 (possibly improved with the strategies for grouping commuting operators mentioned in 2.3.3.4). Since the product of two Pauli strings will also be a Pauli string, measuring $[\hat{H}, \hat{A}_i]$ amounts to defining a new observable that can be measured in a similar way. In all cases, the different strings can be measured in parallel, if multiple quantum processors are available.

It is important to note that the pool operators \hat{A}_i are basic spin-complemented one and two-body excitations, that are Jordan-Wigner transformed into a linear combination of at most 16 Pauli strings (8 from each of the two anti-Hermitian two-body excitations with opposite spins), independently of the system size. Further, the Hamiltonian itself has been decomposed into a sum of polynomially-many Pauli strings (the number being $\mathcal{O}(N^4)$ on the number of spin-orbitals N). The observable corresponding to the commutator $[\hat{H}, \hat{A}_i]$ can then be efficiently measured in a quantum computer. Thus, the extra measurements of the ADAPT-VQE algorithm do not in principle hinder efficiency, as long as the scaling of the size of the operator pool is reasonable. In the case considered in the original article, including only excitations up to order two, this scaling is evidently also quartic on the number of spin-orbitals of the chosen basis set. The measurements per iteration in ADAPT-VQE thus scale like $\mathcal{O}(N^8)$ (instead of $\mathcal{O}(N^4)$ like in VQE with static ansätze). It should be noted that, as mentioned previously, modified variants such as ADAPT-V and ADAPT-Vx manage to decrease this measurement overhead at the expense of an acceptable increase in the number of variational parameters and circuit depth [55].

While polynomial, the measurement overhead of ADAPT-VQE as compared to UCCSD-VQE or other predetermined ansätze is considerable. This is the price to pay for growing the ansatz using a procedure this customized: to choose the new operator, the gradients of all operators in the pool must be measured. ADAPT-VQE also implies an overhead in optimizations: instead of a single optimization with a fixed ansatz, we have multiple optimizations with a growing ansatz. The larger number of optimizations implies in itself a larger number of energy evaluations, as well as a bigger burden on the optimizer. All this sacrifice is done in favour of shallower circuits, and in some cases, greater accuracy. The ADAPT-VQE ansatz was designed to be a NISQ-friendly, high-accuracy ansatz for chemistry applications.

4.1.2 Qubit-ADAPT-VQE

In 2020, Qubit-ADAPT-VQE was proposed in [85] with the purpose of improving upon the previous ADAPT-VQE algorithm. The entirety of the modification lies in the operator pool.

Spin-complemented fermionic single and double excitations can consist of a linear combination of up to 18 Pauli strings, and spin-adapted ones of up to 48 Pauli strings. Using the ladder-of-CNOTs implementation of Pauli exponentials [64], the circuit for one of such operators can require already thousands of CNOTs when as few as a dozen spin-orbitals are used to describe a molecule, even assuming all-to-all connectivity.

Qubit-ADAPT-VQE drastically reduces the circuit depth required for the implementation of each operator, by dispensing with the rigorous representation of fermionic excitations. Instead of using a linear combination of Pauli strings that respects fermionic symmetries (e.g. particle number), the operators are broken down into the individual Pauli strings. Because of this, the pool is called *qubit* pool, whereas the pool with the proper excitations is called *fermionic* pool. In addition to the one directly arising from the decomposition fermionic excitations, the original Qubit-ADAPT-VQE article introduced a few more pools, which will be covered in the next section.

4.2 Original Pool Choices

This section provides a more thorough description of the different pools used in the original articles: as it was stated before, they can be of *fermion* or *qubit* inspiration. The former take inspiration from fermionic excitations, and the latter opt for more hardware-friendly operators, that are easier to implement in quantum computers.

4.2.1 Fermion Inspired Pools

The first ADAPT-VQE article [35] proposed pools with operators of fermionic inspiration. The ansätze resulting from pools of this type are remarkably similar to the UCCSD operator, in that they also typically contain exponentiated single and double excitations. Of course, the ordering will usually be different, as it is dictated by the system in study; further, the number of operators in the ADAPT-VQE will vary (not only with respect to UCCSD-VQE, but also within the algorithm, depending on the molecule and on the convergence criterion).

The operators in fermionic pools respect fermionic symmetries: namely, particle number, S_Z and S^2 preservation, and the anticommutation requirement.

A simple N -order fermionic excitation consists of exciting N electrons from some N orbitals to other N orbitals. Often, *generalized* excitations are used: this means that we include operators that excite fermions from occupied orbitals to occupied orbitals, for example, rather than limiting ourselves to occupied-to-virtual excitations.

The fermionic pools used in this project were taken or adapted from the simulation code of the original ADAPT-VQE article [35], that can be found in the public GitHub repository [4].

4.2.1.1 Spin Adapted Generalized Singles and Doubles

One option is the SGSD pool, also called *Spin-Adapted Generalized Singles and Doubles (SAGSD)*. This pool was compared against the qubit-inspired pool of qubit-ADAPT-VQE in the article that introduced the latter [85].

In this pool, the double excitation operators are split into their singlet and triplet components.

In the first quantization formalism, spin-adapted single τ_1 and double τ_2 excitation operators can be written in the basis of eigenstates of the total S_z and S^2 operators as:

$$\begin{aligned}
\hat{\tau}_1 &= |s = 1/2, s_z = 1/2\rangle_p \langle s = 1/2, s_z = 1/2|_q \\
&\quad + |s = 1/2, s_z = -1/2\rangle_p \langle s = 1/2, s_z = -1/2|_q \\
&\quad - h.c., \\
\hat{\tau}_{2,A} &= -\frac{2}{\sqrt{12}} (|s = 1, s_z = 1\rangle_{rs} \langle s = 1, s_z = 1|_{pq} \\
&\quad + |s = 1, s_z = 0\rangle_{rs} \langle s = 1, s_z = 0|_{pq} \\
&\quad + |s = 1, s_z = -1\rangle_{rs} \langle s = 1, s_z = -1|_{pq}) \\
&\quad - h.c., \\
\hat{\tau}_{2,B} &= -\frac{1}{2} |s = 0, s_z = 0\rangle_{rs} \langle s = 0, s_z = 0|_{pq} \\
&\quad - h.c.,
\end{aligned} \tag{4.2}$$

where $\tau_{2,A}$ is the triplet operator and $\tau_{2,B}$ is the singlet operator, and *h.c.* denotes the Hermitian conjugate of the expression before. Of course, not all sets of four spin-orbitals allow all these excitations. If p, q or r, s are spin-orbitals with the same spatial part, then there is only the singlet operator, since the triplet wave function is symmetric and thus can't be the spin part of the wave function of two fermions with the same spatial part (as the antisymmetry principle would not be respected).

Rewriting these operators in the basis of tensor products of single particle S_z and S^2 eigenstates, and using the convention that an overbar on the spatial orbital index implies a β (down) spin (while no overbar implies α , or up, spin), we get:

$$\begin{aligned}
\hat{\tau}_1 &= |p\rangle \langle q| + |\bar{p}\rangle \langle \bar{q}| - h.c., \\
\hat{\tau}_{2,A} &= -\frac{2}{\sqrt{12}} \left(|rs\rangle \langle pq| + \frac{1}{2} (|r\bar{s}\rangle + |\bar{r}s\rangle) (\langle p\bar{q}| + \langle \bar{p}q|) + |\bar{r}\bar{s}\rangle \langle \bar{p}\bar{q}| \right) - h.c., \\
\hat{\tau}_{2,B} &= -\frac{1}{2} (|r\bar{s}\rangle - |\bar{r}s\rangle) (\langle p\bar{q}| - \langle \bar{p}q|) - h.c.
\end{aligned} \tag{4.3}$$

And finally, we can transform this into the second quantization formalism to obtain the operators in 4.4.

$$\begin{aligned}
\hat{\tau}_1 &= a_p^\dagger a_q + a_{\bar{p}}^\dagger a_{\bar{q}} - h.c., \\
\hat{\tau}_{2,A} &= \frac{1}{\sqrt{12}} \left(2a_r^\dagger a_p a_s^\dagger a_q + 2a_{\bar{r}}^\dagger a_{\bar{p}} a_{\bar{s}}^\dagger a_{\bar{q}} + a_r^\dagger a_p a_s^\dagger a_{\bar{q}} + a_{\bar{r}}^\dagger a_{\bar{p}} a_s^\dagger a_q + \right. \\
&\quad \left. a_r^\dagger a_{\bar{p}} a_{\bar{s}}^\dagger a_q + a_{\bar{r}}^\dagger a_p a_s^\dagger a_{\bar{q}} \right) - h.c., \\
\hat{\tau}_{2,B} &= \frac{1}{2} \left(a_r^\dagger a_p a_s^\dagger a_{\bar{q}} + a_{\bar{r}}^\dagger a_{\bar{p}} a_{\bar{s}}^\dagger a_q - a_r^\dagger a_{\bar{p}} a_s^\dagger a_q - a_{\bar{r}}^\dagger a_p a_s^\dagger a_{\bar{q}} \right) - h.c.
\end{aligned} \tag{4.4}$$

For convenience, the fermionic anticommutation relations 2.53 were used to reorganize the order of the operators; for this purpose, they are more conveniently formulated as in 4.5.

$$a_i^\dagger a_j^\dagger = -a_j^\dagger a_i^\dagger, \quad a_i a_j = -a_j a_i, \quad a_i a_j^\dagger = \delta_{ij} - a_j^\dagger a_i \quad (4.5)$$

The formulation of the pool operators using creation and annihilation operators 4.4 is particularly useful, as they are directly mapped to qubit operators via the Jordan-Wigner transformation 2.57. The OpenFermion [2] library allows creating instances of the FermionOperator class directly from expressions like this one.

Under the Jordan-Wigner transformation, it would appear that each product of four fermionic ladder operators is mapped to a linear combination of 16 Pauli strings; however, the Pauli strings with real coefficients don't change under conjugation and are thus canceled out by the Hermitian conjugate, which is why adding it forces unitarity in the first place. The Pauli strings with imaginary coefficients simply switch sign under conjugation, so that each anti-Hermitian sum of four body operators results in a linear combination of only 8 Pauli strings. As a consequence, the operators in the SGSD pool consist of a linear combination of at most 48 Pauli strings. Even though the order of the excitations is the same, because of the anticommutation string, these Pauli strings act on more qubits for larger systems: their length grows on average linearly with the size of the system. The number of operators in this pool scales as $\mathcal{O}(N^4)$, where N is the number of spin-orbitals.

4.2.1.2 Spin Complemented Generalized Singles and Doubles

In the SCGSD pool, each excitation is grouped along with its spin-complement (i.e. the excitation acting on the opposite spin-orbitals of the same spatial orbitals).

In the first quantization formalism, they can be written

$$\begin{aligned} \hat{t}_1 &= |p\rangle\langle q| + |\bar{p}\rangle\langle \bar{q}| - h.c., \\ \hat{t}_{2,A} &= -(|rs\rangle\langle pq| + |\bar{r}\bar{s}\rangle\langle \bar{p}\bar{q}|) - h.c., \\ \hat{t}_{2,B} &= -(|r\bar{s}\rangle\langle p\bar{q}| + |\bar{r}s\rangle\langle \bar{p}q|) - h.c., \\ \hat{t}_{2,C} &= -(|r\bar{s}\rangle\langle \bar{p}q| + |\bar{r}s\rangle\langle p\bar{q}|) - h.c. \end{aligned} \quad (4.6)$$

And if we transform them into the second quantization formalism, we obtain (using equations 4.5 to reorder the expressions once again)

$$\begin{aligned} \hat{t}_1 &= a_p^\dagger a_q + a_{\bar{p}}^\dagger a_{\bar{q}} - h.c., \\ \hat{t}_{2,A} &= a_r^\dagger a_p a_s^\dagger a_q + a_{\bar{r}}^\dagger a_{\bar{p}} a_{\bar{s}}^\dagger a_{\bar{q}} - h.c., \\ \hat{t}_{2,B} &= a_r^\dagger a_p a_{\bar{s}}^\dagger a_{\bar{q}} + a_{\bar{r}}^\dagger a_{\bar{p}} a_s^\dagger a_q - h.c., \\ \hat{t}_{2,C} &= a_r^\dagger a_{\bar{p}} a_{\bar{s}}^\dagger a_q + a_{\bar{r}}^\dagger a_p a_s^\dagger a_{\bar{q}} - h.c. \end{aligned} \quad (4.7)$$

The operators in this pool consist of a linear combination of at most 16 Pauli strings, whose length also scales linearly (on average) with N , the number of spin-orbitals; and once again, the number of operators in the pool scales like $\mathcal{O}(N^4)$. While the [SCGSD](#) pool has more operators consisting of less Pauli strings on average than the [SGSD](#) pool, the asymptotic scaling with the system size remains unchanged between the two.

4.2.2 Qubit Inspired Pools

The first pools of qubit inspiration to be introduced broke the fermionic excitation operators down to their constituent Pauli strings. From that, other qubit pools were proposed with the goal of further decreasing the circuit depth per operator and the size of the pool.

In general, operators in qubit pools don't respect fermionic symmetries. They usually abdicate of particle number, S_Z and S^2 preservation, and may also forgo anticommutation.

The pools to follow were presented in the article introducing [Qubit-ADAPT-VQE](#) [85].

4.2.2.1 Pauli String Pool

The first pool to be considered in [Qubit-ADAPT-VQE](#), then called simply Qubit Pool, was the one consisting of all individual Pauli strings appearing in the [SGSD](#) pool (or equivalently, the [SCGSD](#) one: they result in the same set of Pauli strings, even if doubles might differ).

As a simple example, we can consider the following (fermionic) single excitation.

$$0.5a_0^\dagger a_4 + 0.5a_1^\dagger a_5 - 0.5a_4^\dagger a_0 - 0.5a_5^\dagger a_1$$

This excitation exists (when there are more than 6 spin-orbitals) both in the [SCGSD](#) and [SGSD](#) pool, since they don't differ in single excitations. The two last terms are the Hermitian conjugates of the former two, and OpenFermion's [2] orbital ordering was used, so that 0, 1 and 4, 5 are pairs of spin-complements corresponding to the same spatial orbital. Under the Jordan-Wigner transformation, this becomes

$$\begin{aligned} &0.25i \cdot X_0 \otimes Z_1 \otimes Z_2 \otimes Z_3 \otimes Y_4 + 0.25i \cdot X_1 \otimes Z_2 \otimes Z_3 \otimes Z_4 \otimes Y_5 \\ &-0.25i \cdot Y_0 \otimes Z_1 \otimes Z_2 \otimes Z_3 \otimes X_4 - 0.25i \cdot Y_1 \otimes Z_2 \otimes Z_3 \otimes Z_4 \otimes X_5. \end{aligned}$$

This excitation then gives rise to four operators in the Qubit Pool:

$$\begin{aligned} i \cdot X_0 \otimes Z_1 \otimes Z_2 \otimes Z_3 \otimes Y_4, & & i \cdot X_1 \otimes Z_2 \otimes Z_3 \otimes Z_4 \otimes Y_5, \\ i \cdot Y_0 \otimes Z_1 \otimes Z_2 \otimes Z_3 \otimes X_4, & & i \cdot Y_1 \otimes Z_2 \otimes Z_3 \otimes Z_4 \otimes X_5. \end{aligned}$$

The norm of the coefficients no longer matters because now each Pauli string will have its own variational parameter.

The number of Pauli strings per fermionic excitation remains approximately the same on average as the size of the system grows. Thus, the size of the pool is still $\mathcal{O}(N^4)$, N the number of spin-orbitals - albeit with a significantly larger prefactor that implies a larger number of measurements per iteration. The Qubit Pool operators also act on a number of qubits that grows like $\mathcal{O}(N)$. However, the circuit depth associated with each operator is significantly changed (in this case, improved). While [SGSD](#) operators contain up to 48 Pauli strings, each operator in the Qubit Pool consists of a single one. This implies an up to almost 50-fold reduction in the circuit depth associated with a single operator, and is the motive behind the [Qubit-ADAPT-VQE](#) proposal.

4.2.2.2 Pauli String Pool Without the Jordan-Wigner String

The [Qubit-ADAPT-VQE](#) article also introduced a modified version of the Qubit Pool, consisting of the same operators, now cleared of the Jordan-Wigner string. This is the $\prod_{k=1}^{i-1} Z_k$ factor that appears in the Jordan-Wigner mapping (formula 2.57) to account for the anticommutation of fermions. The four operators introduced by the same single fermionic excitation used as an example before would simply be

$$i \cdot X_0 \otimes Y_4, \quad i \cdot X_1 \otimes Y_5, \quad i \cdot Y_0 \otimes X_4, \quad i \cdot Y_1 \otimes X_5.$$

The scaling of the total number of operators in this pool is once again $\mathcal{O}(N^4)$, with N being the number of spin-orbitals. However, the scaling of the length of the Pauli strings has changed from $\mathcal{O}(N)$ to $\mathcal{O}(1)$. Without the Jordan-Wigner string, only qubits representing spin-orbitals directly involved in the original excitations will be acted on by the operators they give rise to. Since the excitations we are considering are at most doubles, the Pauli strings will have length four at most.

4.2.2.3 Minimal Pools

Finally, the article that introduced [Qubit-ADAPT-VQE](#) proposed a way of creating minimal pools containing operators capable of transforming any real state into any real state. These minimal pools were proven to be complete and allow convergence when used on random real Hamiltonians. An example of one such pool is the pool G , presented in the original article and reproduced here in 4.8.

$$\begin{aligned} G_1 &= iZ \otimes Y \otimes I^{\otimes(N-2)}, & G_2 &= iI^{\otimes 1} \otimes Z \otimes Y \otimes I^{\otimes(N-3)}, \\ \dots, & & G_{N-2} &= iI^{\otimes(N-3)} \otimes Z \otimes Y \otimes I^{\otimes 1}, & G_{N-1} &= iI^{\otimes(N-2)} \otimes Z \otimes Y, \\ G_N &= iY \otimes I^{\otimes(N-1)}, & G_{N+1} &= iI^{\otimes 1} \otimes Y \otimes I^{\otimes(N-2)}, \\ \dots, & & G_{2N-3} &= iI^{\otimes N-3} \otimes Y \otimes I^{\otimes 2}, & G_{2N-2} &= iI^{\otimes N-2} \otimes Y \otimes I^{\otimes 1} \end{aligned} \tag{4.8}$$

The minimal pools proposed scale like $\mathcal{O}(N)$ on the number of qubits, with a total of only $2N - 2$ operators. This is a remarkable reduction from the quartic scaling from the previous pools: the number of gradients that must be evaluated per iteration is reduced from $\mathcal{O}(N^4)$ to $\mathcal{O}(N)$.

The minimal pools allow for a significant decrease in the circuit depth per operator, and can have a very local structure (as does G , that only ever acts on two followed qubits).

Unfortunately, while they work for generic Hamiltonians (such as random real Hamiltonians), they are not adequate for the specific structure of molecular Hamiltonians that is imposed by fermionic symmetries. In particular, we can see that none of the operators in the G pool conserves the spin or particle number in the Hartree-Fock state, which has severe consequences.

The Hartree-Fock state can be written using OpenFermion's [2] orbital convention as $|000\dots111\rangle$. Acting on this state with a parameterized Y rotation acting on qubit k that is in a computational basis state $|x\rangle$, we get:

$$\begin{aligned} |\psi\rangle &= e^{jY_k\theta} |000\dots\rangle \otimes |x\rangle_k \otimes |\dots111\rangle \\ &= \cos\theta |000\dots\rangle \otimes |x\rangle_k \otimes |\dots111\rangle + (-1)^{1-x} \sin\theta |000\dots\rangle \otimes |1-x\rangle_k \otimes |\dots111\rangle \\ &= \cos\theta |n\rangle + (-1)^{1-x} \sin\theta |m\rangle \end{aligned}$$

Clearly, $|n\rangle$ and $|m\rangle$ represent computational basis states with a different number of electrons (different number of 1 states). Considering that the Hamiltonian matrix is real and Hermitian, so that $H_{m,n} = H_{n,m}$, the energy and gradient evaluated in this state are

$$\begin{aligned} E_{|\psi\rangle} &= \langle\psi| H |\psi\rangle = \cos^2\theta H_{n,n} + \sin^2\theta H_{m,m} + 2(-1)^{1-x} \sin\theta \cos\theta H_{m,n}, \\ \left. \frac{dE_{|\psi\rangle}}{d\theta} \right|_{\theta=0} &= \left(2\cos\theta \sin\theta (H_{m,m} - H_{n,n}) + 2(-1)^{1-x} \cos(2\theta) H_{m,n} \right) \Big|_{\theta=0} \\ &= 2(-1)^{1-x} H_{m,n} \end{aligned}$$

Since $|n\rangle$, $|m\rangle$ correspond to Slater determinants different numbers of electrons, $H_{m,n}$ is zero and the gradient vanishes. Since $|n\rangle$ represents the Hartree-Fock state, the lowest-energy Slater determinant, its energy corresponds to the lowest value in the diagonal of the Hamiltonian. Thus, it is evident that a single parameterized Y rotation (acting on the Hartree-Fock state) has null gradient and does not lower the energy beyond the Hartree-Fock energy.

All operators in pool G are either Y rotations or conditional Y rotations, so that the the same applies. Since all gradients will be zero, the [ADAPT-VQE](#) algorithm will be prevented from starting in the first place, because the selection method will fail. This is a consequence of the minimal pool operators not respecting particle number and spin symmetries.

For the [ADAPT-VQE](#) algorithm to be capable of starting (by adding the first operator to the ansatz), the pool must include operators that preserve symmetries in the Hartree-Fock state. However, this still does not guarantee proper convergence; for the ground state to be reached, the pool operators must obey extra restrictions. An explicit approach for creating symmetry-adapted minimal pools was introduced in 2021 in [75]. Because of the added restrictions, the symmetry-adapted version of the minimal pools can

actually have less operators than the minimal complete pools presented before. The size of the pool is still $\mathcal{O}(N)$, N the number of qubits/spin-orbitals. This means that the measurements per iteration scale as $\mathcal{O}(N^5)$ (against $\mathcal{O}(N^8)$ in fermionic-ADAPT-VQE or qubit-ADAPT-VQE with the original pools), a more modest increase against the $\mathcal{O}(N^4)$ scaling of static ansatz VQE.

4.3 Application to LiH

In order to analyse the fermionic- and qubit-ADAPT-VQE performance, and compare ADAPT-VQE with UCCSD-VQE, the algorithm was applied to multiple molecules.

LiH is a molecule represented by 12 qubits in a minimal basis. Its Hartree-Fock ground energy is significantly distant from the FCI value, and the Hartree-Fock approximation can be significantly improved upon in this case. Unlike H_2 , a smaller molecule, more than a single operator is required to reach chemical accuracy in the case of *LiH*, resulting in a more interesting evolution of the ansatz. While the simulations are heavier computationally, ADAPT-VQE with a threshold ϵ of 0.01 is still viable, and this suffices to reach an interesting accuracy. Because of this, *LiH* was chosen to test the algorithm in a fully noise-free scenario.

4.3.1 Evolution of the Optimization

Figure 23 shows the evolution of the ADAPT-VQE algorithm for *LiH* at an interatomic distance of 1.45Å.

It can be seen that the energy reaches chemical accuracy after four iterations; at this point, the ansatz has 4 spin-adapted operators and the corresponding 4 variational parameters. For reference, the UCCSD ansatz obtained from OpenFermion [2] has for the same molecule 44 variational parameters and 64 (simple) excitation operators. Evidently, this also corresponds to a greater accuracy than the 4-operator ansatz of ADAPT-VQE; but in the plot we see that as we let the algorithm evolve further, it manages to steadily improve the energy estimate. In the original article [35], through simulations including several molecules, it was shown that ADAPT-VQE is capable of outperforming UCCSD-VQE.

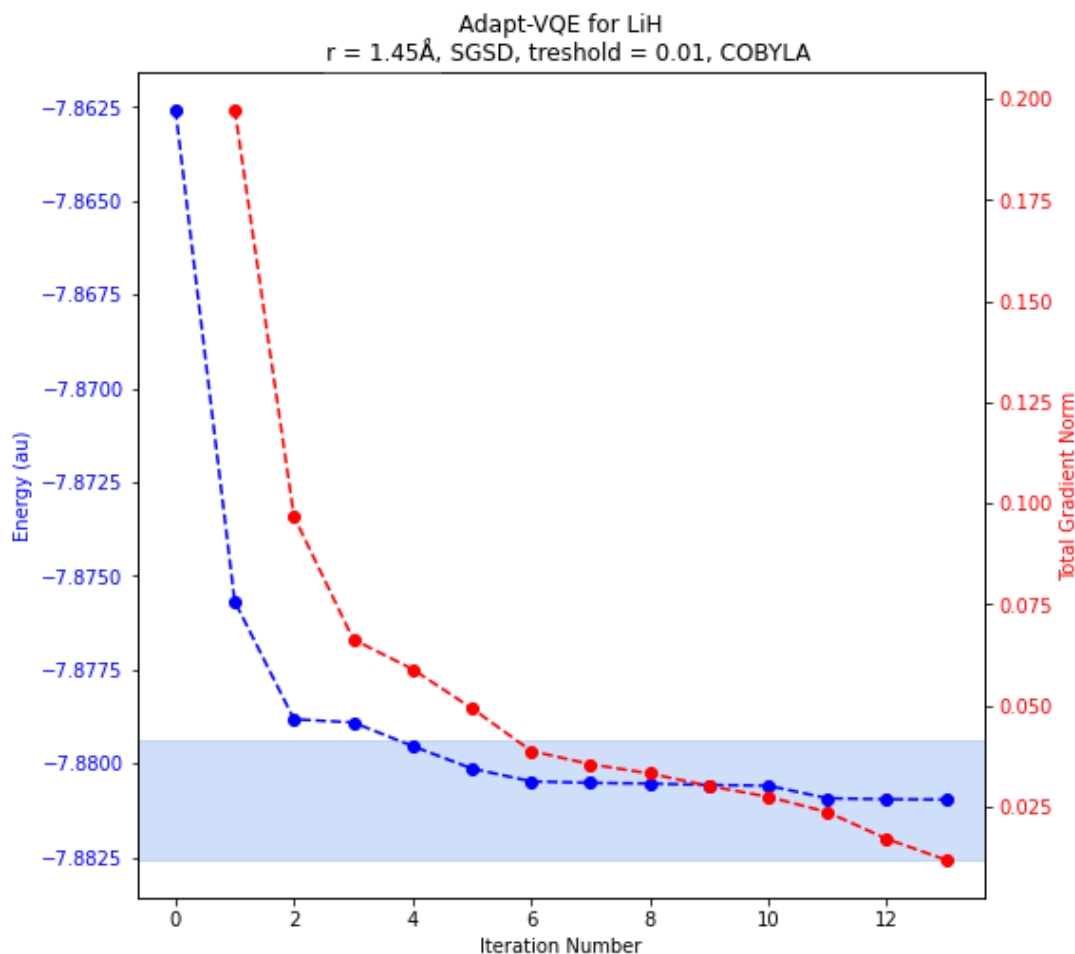


Figure 23: Evolution of the fermionic-ADAPT-VQE energy and gradient norm along the iterations, for the lithium hydride (LiH) molecule at an interatomic distance of 1.45\AA . The convergence criterion was set to a gradient threshold of 0.01. The SGSD pool and the COBYLA optimizer were used. The shaded blue area represents the region of chemical accuracy (error of less than 1kcal/mol). The simulation was done via matrix algebra, and does not include any type of noise.

4.4 Application to Molecular Hydrogen

Like UCCSD-VQE, the Qubit-ADAPT-VQE (with the Pauli string pool cleared of the Z string) was applied to the hydrogen molecule. The small size of this molecule allows it to be represented by only 4 qubits, and results in shallower ansätze. Because of this, it was chosen to run the algorithm on real quantum computers and analyse the effect of noise. Since it is enough to reach chemical accuracy in this case, a single iteration of ADAPT-VQE was used to obtain the ground state. At this point, there is a single variational parameter in the ansatz, which can be implemented with a CNOT count and depth of 6 assuming all-to-all connectivity.

In this section, ADAPT-VQE results will be presented and compared against UCCSD-VQE, with a focus on noise-resilience and (to a lesser extent) measurement costs.

4.4.1 Bond Dissociation Graph

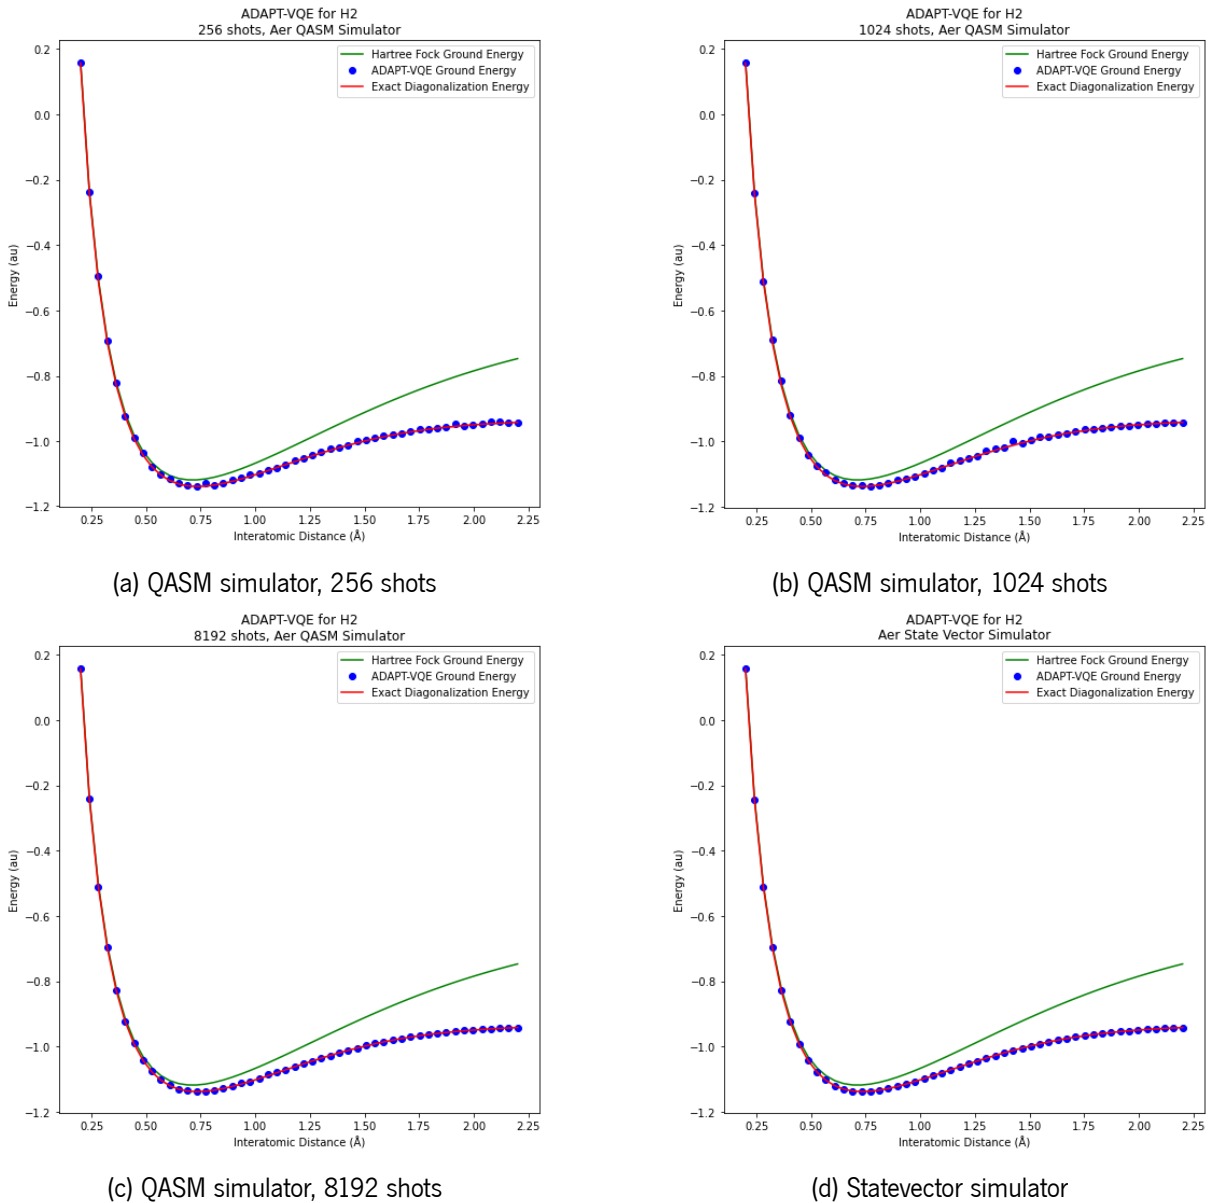


Figure 24: Single-run, single-iteration [ADAPT-VQE](#) ground energy along the bond dissociation curve of the hydrogen molecule, obtained using IBMQ [42] simulators. The results plotted in figures 24a-24c, ordered by increasing shot count, were obtained using the QASM simulator. The results plotted in figure 24d were obtained using the state vector simulator (equivalent to infinite shots). The red and green curves represent the [FCI](#) energy and the Hartree-Fock energy, respectively.

Following the fully noise-free simulation, the bond dissociation graphs of the molecule with sampling noise will now be analysed. The plots in 24 show the [ADAPT-VQE](#) energies along the curve, for different shot counts; they should be compared with figure 15, that presents the same plots for [UCCSD-VQE](#).

Like before, there is a visible improvement on the performance of [ADAPT-VQE](#) with the increase of the number of shots from figure 24a to figure 24c. In the limit of infinite shots (figure 24d), the error is of the

order of 10^{-8} Hartree along the whole curve.

As did **UCCSD-VQE**, **ADAPT-VQE** successfully brings the energy into the **FCI** solution from the Hartree-Fock approximation (which is the reference state). The extra variational freedom (in this case, a single variational parameter) seems to be enough to account for electronic correlation.

What is interesting is that **ADAPT-VQE** clearly suffers less from sampling noise. With only 256 shots (figure 24a), **ADAPT-VQE** is visibly closer to the **FCI** curve than **UCCSD-VQE** with 1024 shots (figure 15b). This is not a circuit depth dependent noise source; what causes the **UCCSD-VQE** performance to be worse is the extra variational freedom, that results in a more arduous optimization.

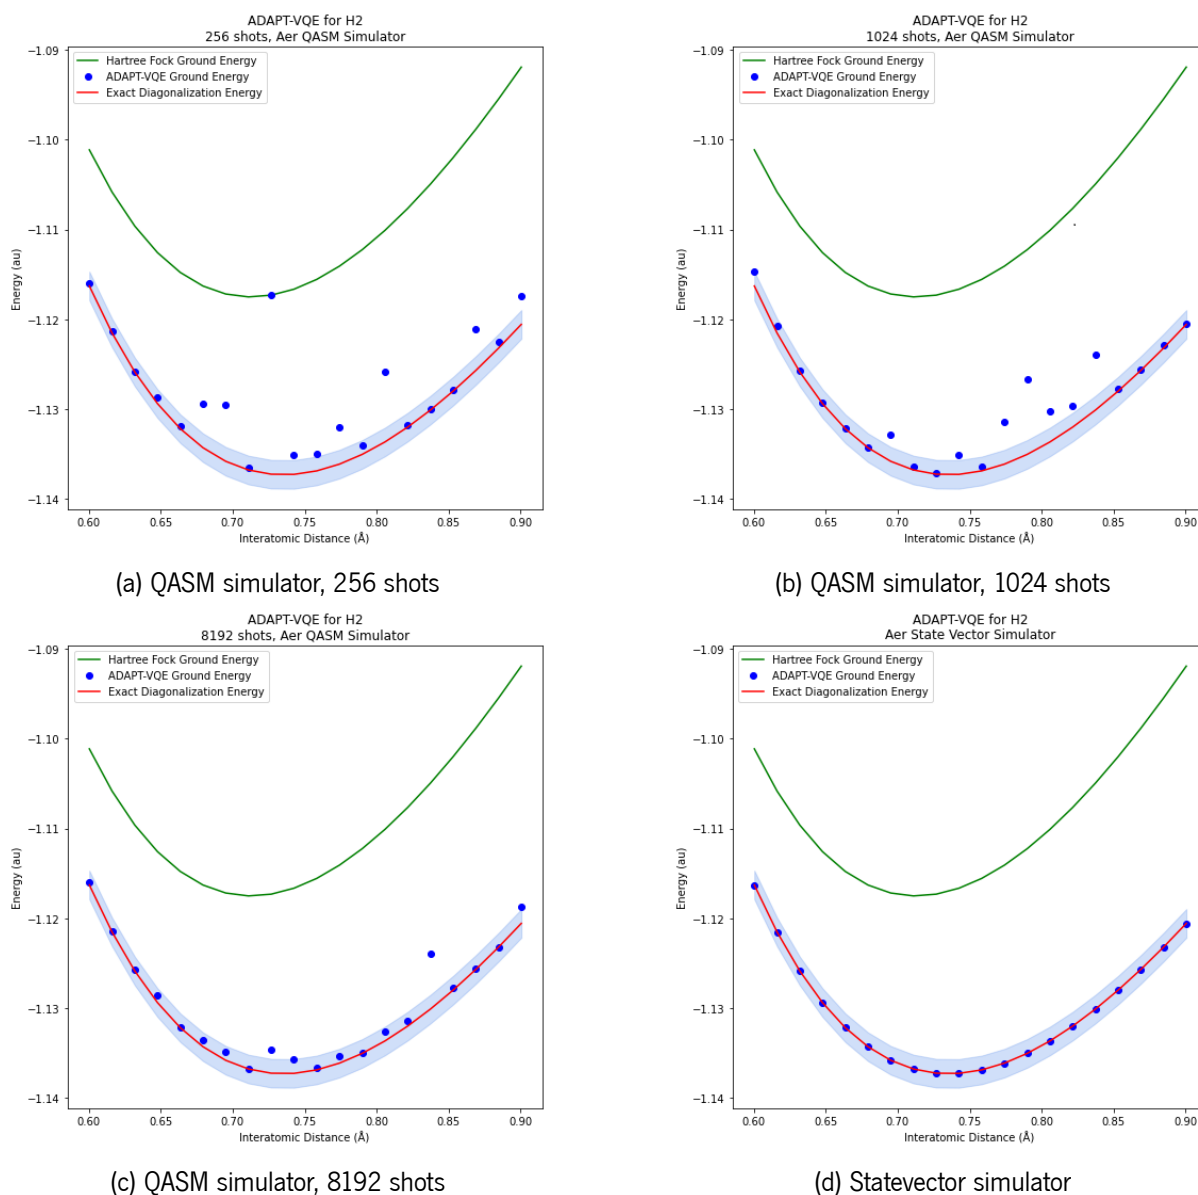


Figure 25: Bond dissociation curves of figure 24, zoomed in on the region nearing the equilibrium bond length. The blue area marks the region of chemical accuracy (error of less than 1kcal/mol).

Figure 25 shows the same plots of figure 24, magnified around the equilibrium bond length. This can

be compared with figure 16, that shows these plots for UCCSD-VQE.

The only scenario in which UCCSD-VQE beats 1-iteration qubit-ADAPT-VQE is the fully noise-free one: even though it is not visible in the graph, simulations showed that under those circumstances UCCSD-VQE could achieve a precision of the order of 10^{-9} Hartree against 10^{-8} from ADAPT-VQE. However, this precision is exceedingly high: not only it is not necessary for most applications, it also can hardly be ever met, considering not only sampling but also other sources of noise prevalent in NISQ devices.

With sampling, regardless of the shot count, single-iteration ADAPT-VQE manages to outperform UCCSD-VQE. With 256 shots (figures 25a and 16a) and 1024 shots (figures 25b and 16b), the UCCSD-VQE result does not fall within chemical accuracy in any of the points, while ADAPT-VQE falls within chemical accuracy in 11 and 13 points respectively. With 8192 shots (figures 25c and 16c), ADAPT-VQE only falls outside of chemical accuracy in 3 out of 20 points, against 12 out of 20 fails of UCCSD-VQE.

4.4.2 Running the Algorithm on Cloud Quantum Computers

After considering the impact of sampling noise, we can now observe the impact of other types of noise by analysing the performance of ADAPT-VQE on a real quantum computer. Figure 26 shows the evolution of the first iteration of ADAPT-VQE in three scenarios: no noise (figure 26a), sampling noise only / ideal quantum processor (figure 26b), real quantum processor (figure 26c). These can be compared with those from figure 17, that show the evolution of the UCCSD-VQE optimization for the same three backends.

In the setting with sampling noise only, ADAPT-VQE already outperforms UCCSD-VQE. As it was mentioned previously, this is mainly due to the extra strain on the optimization caused by the larger number of variational parameters and greater expressibility. While ADAPT-VQE stabilizes the energy inside the region of chemical accuracy after as few as five iterations, UCCSD-VQE is still in and out of this region by the time convergence is reached, which happens after more than double the iterations of ADAPT-VQE.

When a real quantum computer is used, and the algorithms are exposed to all types of noise, the difference becomes even sharper. The ADAPT-VQE and UCCSD-VQE optimizations in the Belem processor will be compared at a later subsection both graphically and numerically.

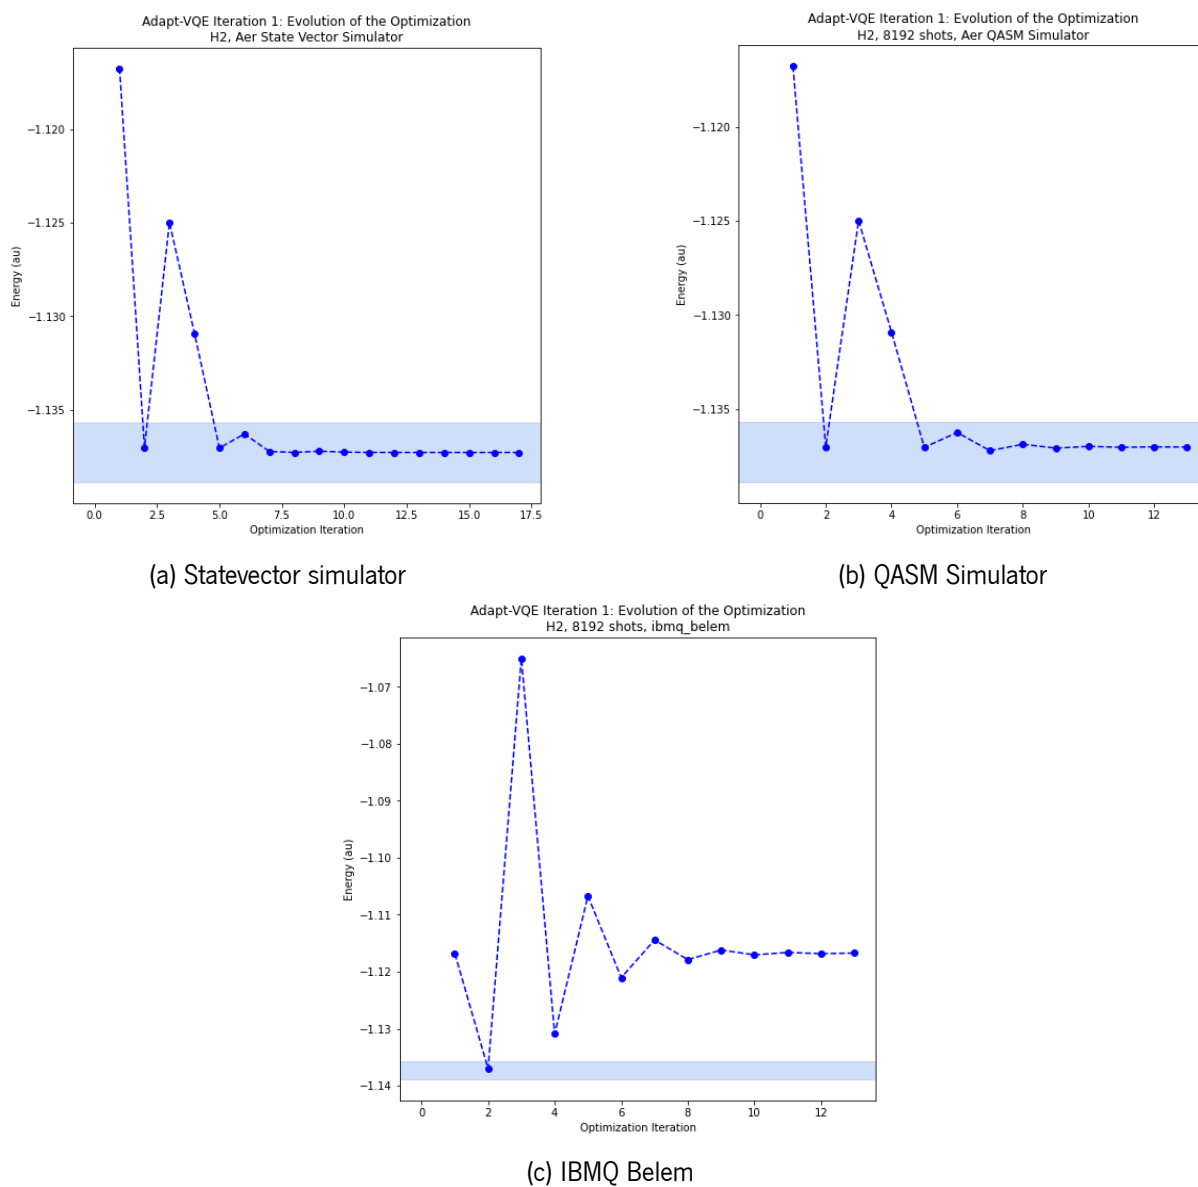


Figure 26: Evolution of first the [ADAPT-VQE](#) optimization for H_2 at an interatomic distance of 0.74 \AA , using three different IBM Quantum [42] backends: the statevector simulator (with no noise of any kind), the QASM simulator (with sampling noise) and the Belem device (a 5-qubit Falcon processor). In the last two cases, 8192 shots were used. The energy is plotted in blue; the pale blue region marks the area of chemical accuracy. The same starting point, the Hartree-Fock reference state, was used for all backends. The chosen optimizer was COBYLA.

4.4.3 Impact of Thermal Relaxation, SPAM errors, and Sampling

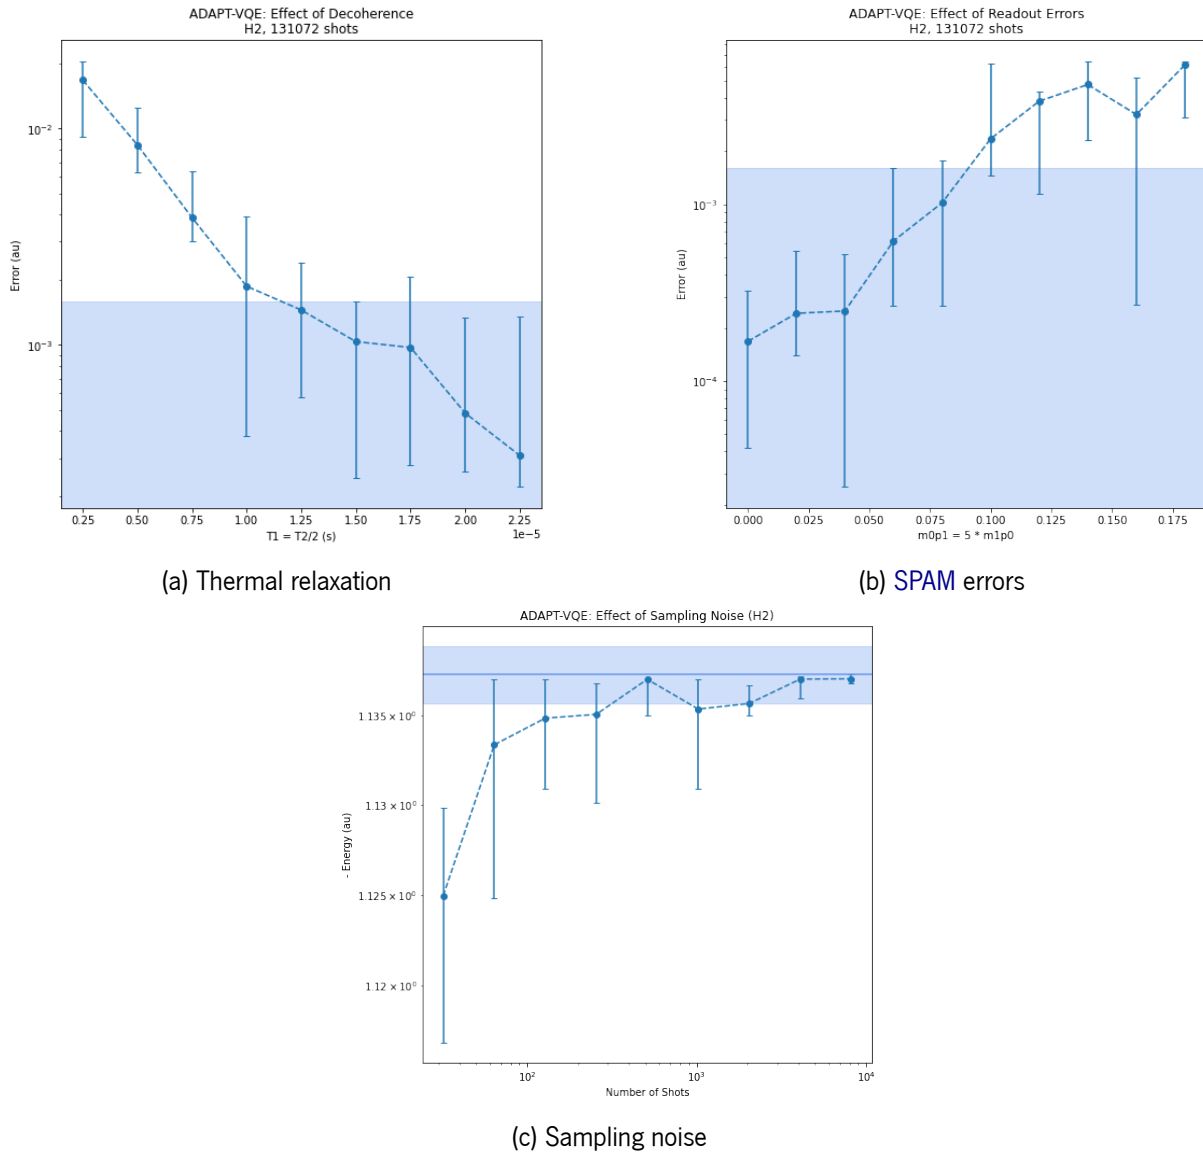


Figure 27: Plots showing the impact of different types of noise on the performance of ADAPT-VQE for H_2 at an interatomic distance of 0.74\AA . For thermal relaxation, a simple model assuming average gate times and a constant ratio between the coherence times $T1$ and $T2$ was used. Figure 27a shows the evolution of the error in the energy as a function of these times. For SPAM noise, the ratio between the probability of measuring 0 after preparing 1 and the probability of measuring 1 after preparing 0 was fixed (and greater than one, for the model to be realistic). Figure 27b shows the evolution of the error as a function of these probabilities. Figure 27c illustrates the impact of sampling noise by plotting the error as a function of the number of shots. Twenty runs were used for each point; the median was used as the estimator (as it filters out aberrant runs), and the error bars show the interquartile ranges. All noise models were created using Qiskit [3]. The chosen optimizer was COBYLA.

Figure 27 shows the impact of thermal relaxation, SPAM errors, and sampling noise in ADAPT-VQE. The circumstances were chosen to meet exactly those from figure 21, that show the same plots from

UCCSD-VQE. The same models of thermal relaxation, **SPAM** noise, and sampling noise, created in Qiskit [3], were used.

From 27a, we can see that the majority of the runs fall inside of chemical accuracy when T1 is greater than 2×10^{-5} seconds (with T2 being fixed at 4×10^{-5} s). This condition is actually met by IBMQ's devices, that typically offer coherence times of the order of 10^{-4} seconds.

The impact of **SPAM** errors can be seen on figure 27b. For the runs within the interquartile range to fall inside of chemical accuracy, the probability of measuring $|0\rangle$ after preparing $|1\rangle$ should be 0.06 (with the probability of measuring $|1\rangle$ after preparing $|0\rangle$, made proportional, at 0.012). The Belem backend used to run **ADAPT-VQE** before nearly meets the requirement, with average error probabilities around 0.04 and 0.009, respectively.

Finally, figure 27c shows that **ADAPT-VQE** requires no more than 4000 shots to reach chemical accuracy in the middle 50% runs in a sampling noise only scenario. This shot count is actually under the maximum allowed in IBMQ's cloud quantum computers.

4.4.4 Comparison with UCCSD-VQE

While the previous **ADAPT-VQE** results were already compared against those from **UCCSD-VQE** presented on subsection 3.2.1.4, it will be interesting to compare them side-by-side.




Figure 28: Evolution of the optimization of the first iteration of **ADAPT-VQE** (figure 28a) and of the **UCCSD-VQE** algorithm, for H_2 at an interatomic distance of 0.74 Å. IBMQ's Belem processor was used in both cases.

Figure 28 compares the **ADAPT-VQE** (figure 28a) and **UCCSD-VQE** (figure 28b) optimizations, using the same scale.

Neither of the algorithms manages to reach chemical accuracy with this backend and shot count, at least in this particular iteration. Regardless, [ADAPT-VQE](#) secures the energy close to the Hartree-Fock energy, with a final error of the order of centesimals of Hartree. In contrast, the [UCCSD-VQE](#) outputs an energy with an error greater than unity. The appearance of the [UCCSD-VQE](#) optimization plot is a symptom of a *noise-induced barren plateau*: due to the excessive circuit depth, noise washes away the quantum information and erases the characteristics of the cost function, causing the optimizer to vainly search through a flat optimization landscape.

When transpiled onto the Belem backend, the single-iteration [ADAPT-VQE](#) ansatz has 12 [CNOTs](#), against 130 from [UCCSD-VQE](#). This evidently impacts the purity of the state by the end of the circuit.

As was done in the previous chapter, an approximate density matrix representing the state at the end of the noisy [ADAPT-VQE](#) ansatz was obtained by simulating the algorithm on the QASM simulator with a noise model aiming to mimic the behaviour of the Belem processor.

The fidelity of the final state with the ideal state at the end of the [Adapt](#) ansatz was found to be 0.92 (against only 0.081 in [UCCSD-VQE](#)), while the fidelity with the fully mixed state was 0.087 (against 0.40). The final purity of the state was found to be around 0.84 (against 0.16).

Figure 29 compares the noise resilience of [ADAPT-VQE](#) and [UCCSD-VQE](#) against different sources (thermal relaxation, [SPAM](#) errors, and sampling). The same noise sources were already explored individually in 21 and 27; however, now the scales were adjusted so that both curves start when [ADAPT-VQE](#) is still outside of chemical accuracy and finish when [UCCSD-VQE](#) is already inside.

To reach chemical accuracy in more than 50% of the runs, [UCCSD-VQE](#) required roughly 850 times greater coherence times than [ADAPT-VQE](#). The difference in performance is visible in the plot 29a, with [ADAPT-VQE](#) showing around one order of magnitude greater accuracy for any given value of T1 and T2. This is a consequence of the greater circuit depth required by the [UCCSD](#) ansatz.

Figure 29c shows the impact of sampling noise, and figure 29b the impact of [SPAM](#) noise. In order to reach chemical accuracy in at least half of the runs in the presence of sampling noise exclusively, [ADAPT-VQE](#) required only 5% of the number of shots required by [UCCSD-VQE](#). [ADAPT-VQE](#) also tolerated roughly 150 times higher error probability in state preparation and measurements. It is interesting to see how both of these noise sources have greater impact in the performance of [UCCSD-VQE](#), regardless of being independent of circuit depth. Once again it becomes clear that more variational parameters and more variational flexibility also imply an added difficulty in the optimization, especially in the presence of noise.

The noise resilience of [ADAPT-VQE](#) against [UCCSD-VQE](#) has become evident, but it must be noted that this is not the only advantage of [ADAPT-VQE](#) against [UCCSD-VQE](#): in the original article [35], it was shown that while the latter often fails to reach chemical accuracy for strongly correlated molecules, [ADAPT-VQE](#) manages to reach it as long as the convergence criterion is sufficiently ambitious. These simulations were not replicated due to the heavy computational demands involved.

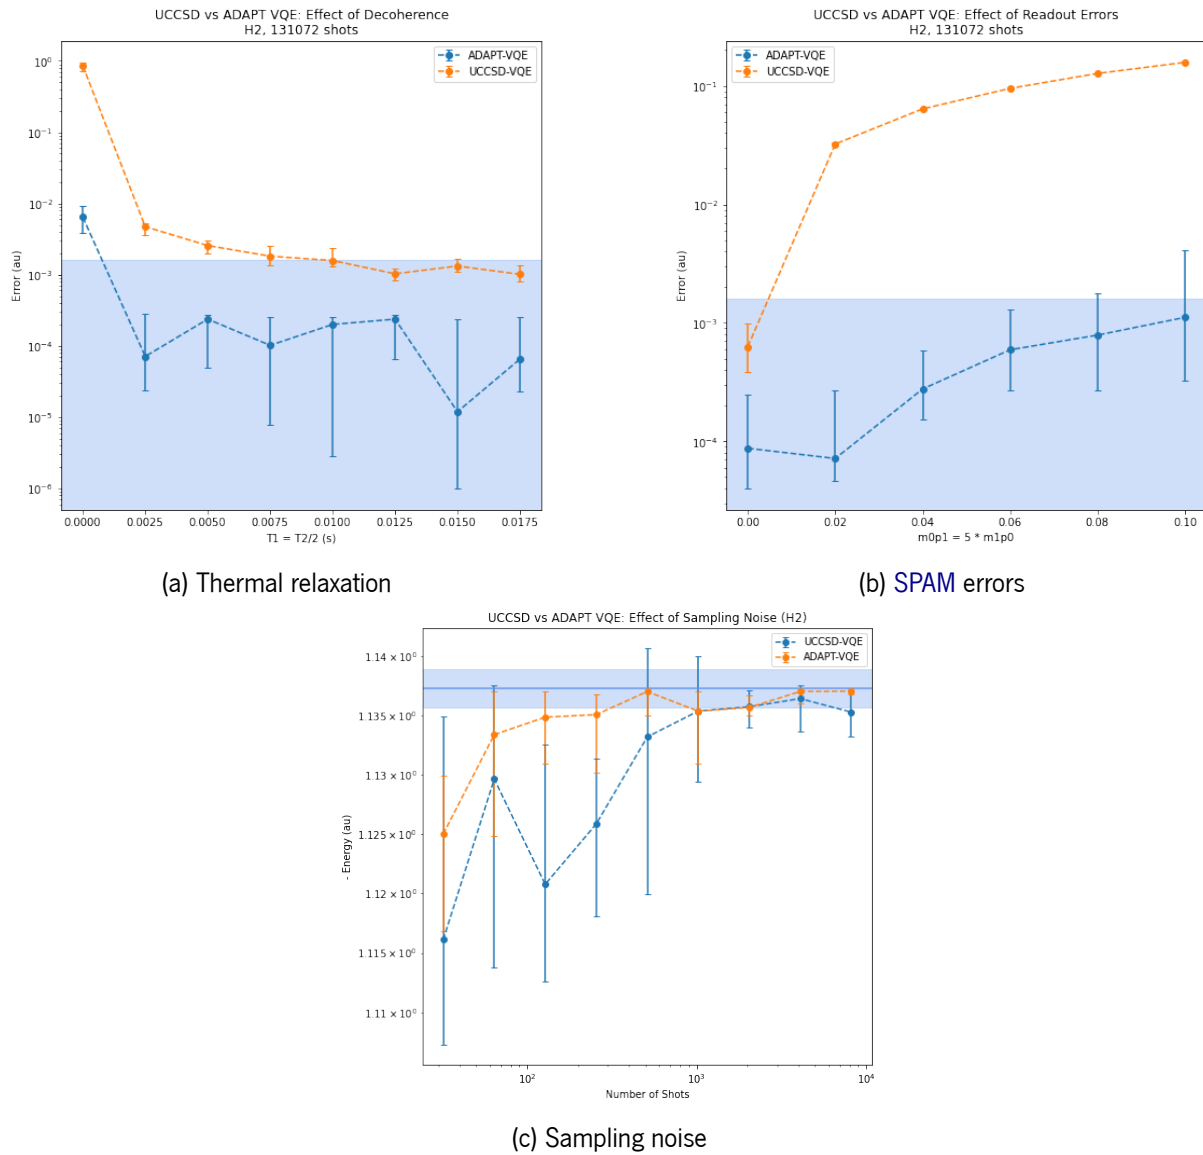


Figure 29: Comparison of the effect of different types of noise on **ADAPT-VQE** and **UCCSD-VQE** for H_2 at an interatomic distance of 0.74\AA . The same noise models described in figures 21 and 27 were used.

While the results regarding accuracy and noise-resilience favour the latter, comparing the performance of [UCCSD-VQE](#) and [ADAPT-VQE](#) is a delicate matter that goes beyond the impact of noise or the final error in the energy. There are many different costs at play, and many links between them. [ADAPT-VQE](#) requires evaluating the expectation value of a significantly larger number of observables, because all gradients must be evaluated each iteration. Additionally, because each iteration requires re-optimizing all the parameters, there will be a significant accumulated number of optimizations by the end of the algorithm, unlike in [UCCSD-VQE](#) (which requires a single optimization). This makes it seem likely that [ADAPT-VQE](#) will imply a larger number of measurements in total.


However, attention must be paid to the fact that since [ADAPT-VQE](#) creates the ansatz from scratch, most of the optimizations will have a significantly lower number of variational parameters. Because they're lower dimensional, these optimizations are likely to be easier, and to require a lower number of function evaluations (i.e. calls to the quantum computer). Further, in the common case that the final [ADAPT-VQE](#) ansatz includes less variational parameters than the [UCCSD](#) ansatz, [ADAPT-VQE](#) will tolerate sampling noise better, and thus require a lower number of shots per term for the same accuracy. Additionally, the measurement overhead from the gradient measurements may be softened by the existence of common factors among the Pauli strings that occur in the $[\hat{H}, \hat{A}_i]$ observables (and in the Hamiltonian itself). This allows reusing measurements.

Figure 30 aims to carefully analyse all of the different factors that weigh in on the costs. The number of shots for each algorithm was chosen so that the precision was matched; this implied a significantly larger number of shots for [UCCSD-VQE](#) (32 times larger, in this case). As explained before, this is due to the fact that the optimization is higher-dimensional, thus more difficult, and more sensitive to noise.


Figure 30b compares the number of energy evaluations throughout the optimizations. This number is larger for [UCCSD-VQE](#); once more, we can attribute this to the larger number of variational parameters in the optimization. However, this plot does not consider the fact that the number of *observables* measured in [ADAPT-VQE](#) is larger, because the gradients of the pool operators will also be measured in each iteration.

The additional measurements required for the evaluation of the gradients in [ADAPT-VQE](#) are factored in in figure 30c. Here, the number of observables includes *all* observables measured along a complete run of either algorithm. For [ADAPT-VQE](#), this includes all energy and gradient measurements; for [UCCSD-VQE](#), it matches the number of energy evaluations. In both cases, the utilized optimizer was gradient-free, so that all calls to the quantum computer throughout the optimization consisted of energy evaluations. As expected, the total number of measured observables was larger for [ADAPT-VQE](#). In this case, the pool had 20 operators, implying an additional 20 observables to be measured each iteration (the $[\hat{H}, \hat{A}_i]$).

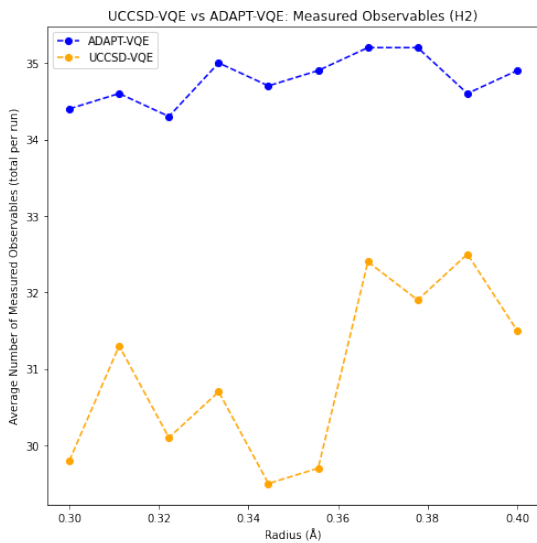
Finally, figure 30d compares a very important metric: the total number of shots in a full run. This number includes all the shots used for the evaluation of *all* Pauli strings in *all* observables in *all* times they were measured. This plot thus reflects a multitude of factors: it weighs in the fact that [UCCSD-VQE](#) requires more shots per string and more energy evaluations per optimization, as well as the fact that [ADAPT-VQE](#) requires additional measurements per iteration to obtain the gradients. The total number of shots is the total number of circuits, repeated or not, that have to be executed on the quantum computer.




(a) Error in the energy



(b) Number of energy evaluations



(c) Number of measured observables



(d) Total number of shots

Figure 30: Plots comparing the costs associated with UCCSD-VQE and ADAPT-VQE for a similar final error, in a setting with no noise other than sampling noise. For each interatomic distance, 10 runs of the algorithms were performed. The median of the errors is plotted in figure 30a. Figures 30b, 30c, and 30d present respectively the average number of energy evaluations, measured observables, and total shots per run. The number of shots per Pauli string was set to 2^{13} and 2^{18} for ADAPT-VQE and UCCSD-VQE respectively. These values were chosen so that the error was roughly matched between the two algorithms (the average error was 0.00038 a.u. and 0.00039 a.u. by the same order).

Interestingly, [ADAPT-VQE](#) required a *smaller* number of shots in total than [UCCSD-VQE](#) for the same final error. In addition to implying shallower circuits and a lower number of variational parameters, [ADAPT-VQE](#) does not in this case imply an overhead in measurements as compared to [UCCSD-VQE](#) - on the contrary. It required only a fraction of the number of shots in total (1.7%).

In this example, the fact that sampling noise is better tolerated by [ADAPT-VQE](#) allows for a reduction of the number of shots *per Pauli string* that compensates for the fact that there will be a larger number of observables being measured. Additionally, it must be mentioned that as was conjectured before, there was a significant amount of repeated Pauli strings among these observables. For the molecule in question (H_2), the total number of Pauli strings in the $[\hat{H}, \hat{A}_i]$ observables was 152. However, the set of Pauli strings that were *unique* among these and that didn't occur in the Hamiltonian, i.e. the Pauli strings that should actually be measured, was found to consist of only 24 elements. Evidently, acknowledging this fact allows for a reduction of the total number of measured Pauli strings, which in turn reduces the total number of shots. This was already considered in obtaining the results plotted in figure 30d.

Of course, these results should be interpreted with caution: we're only considering a particular molecule, and a very small one. As the size of the system grows, so will the number of operators in the pool, increasing the measurement overhead of [ADAPT-VQE](#). The total number of optimizations required for the same accuracy will grow as well. On the other hand, the difficulty that is brought on by the extra variational parameters in [UCCSD-VQE](#) will also increase with the size of the system. This algorithm will then become even more sensitive to sampling noise, and thus require an even larger precision on the energy evaluations (which will already require a larger number of shots for the *same* precision given that the system is bigger). Additionally, it has been established that [ADAPT-VQE](#) is better suited for real quantum processors, as the resulting circuits are shallower. The most evident benefit of this is that it improves the viability of the algorithm in near-term devices. But given that additional sources of noise also imply an additional difficulty in the optimizations, their presence is also likely to benefit the relative performance of [ADAPT-VQE](#) against [UCCSD-VQE](#) in what comes to measurement costs. With more qubits and deeper circuits, errors will accumulate faster, so this is another factor that is affected by the size of the system.

Taking all of this into consideration, it is not certain how the two algorithms will compare in terms of the total number of required shots for larger molecules and in more realistic settings. Such simulations were not attempted on account of the prohibitive computational complexity.

4.5 Application to H_4

The H_4 molecule is represented by 8 qubits in a minimal basis. While it is already too large a system to allow for meaningful results to be recovered in real quantum computers, it is still significantly easier to simulate classically than LiH . Because it allows for more iterations to be simulated in useful time, H_4 was used to analyse the effect of removing the Jordan-Wigner string from the Qubit Pool operators, and to compare fermionic-[ADAPT-VQE](#) against Qubit-[ADAPT-VQE](#).

4.5.1 Qubit Pool: Effect of the Jordan Wigner String




Figure 31: Impact of the Jordan-Wigner string on the convergence of the qubit-ADAPT-VQE algorithm. The graphs compare the evolution of the error in the energy along the 20 first iterations of the algorithm, with and without the Z-strings in the pool operators. The molecule in study is the H_4 molecule at an interatomic distance of 1.45Å.


In one of the Qubit Pools introduced before, the chain of Pauli Z operators responsible for faithful translation of the anticommutation of fermions into the qubit states was eliminated from these Pauli strings. The effect of its removal on the convergence of the algorithm for the H_4 molecule is plotted on figure 31.

The plot shows that there is very little variation between qubit pools with and without the Jordan-Wigner string. This is a remarkable fact. Firstly, it is remarkable because it means that representing the fermionic anticommutation accurately through the ansatz operators brings no benefit. But most of all, it is remarkable because it greatly aids in reducing circuit depth, as was explained before. These Z strings are non-local and act on a number of qubits that scales on average linearly with the size of the system. After relinquishing them, only the qubits corresponding to the one-body wave functions directly involved in the excitations will be acted on by the operators. As a consequence, if the employed excitations are at most

doubles, the pool operators will act on at most four qubits regardless of the molecule in study. We saw that the circuit depth required per operator will be $\mathcal{O}(1)$ instead of $\mathcal{O}(N)$.

4.5.2 Fermionic vs Qubit Pools: Performance Comparison

The purpose of the qubit pool is to make the [ADAPT-VQE](#) algorithm more [NISQ-Friendly](#), in hopes of enabling an earlier practical implementation. As it was explained in the previous section, the operators in the qubit pool are in fact implemented by shallower circuits than those in the fermionic pool. However, because they don't preserve fermionic symmetries and consist each of a single Pauli string, they are bound to make convergence slower as a function of the number of iterations. An N -operator ansatz from fermionic-[ADAPT-VQE](#) will typically imply a much greater accuracy than an N -operator ansatz from qubit-[ADAPT-VQE](#). However, it will also imply a much longer circuit. In a fair comparison, circuit depth, the number of variational parameters, and the measurement overhead should all be taken into consideration.



(a) Error as a function of the number of iterations of the algorithm.

(b) Error as a function of the number of CNOT gates in the ansatz circuit.

Figure 32: Evolution of the error in the [ADAPT-VQE](#) energy along the evolution of the algorithm for the [SCGSD](#) pool, the [SGSD](#) pool, and the qubit pool. The molecule in study is the H_4 molecule.

Figure 32 shows the evolution of the error in the [ADAPT-VQE](#) energy for three types of pools, two fermionic ones ([SCGSD](#) and [SGSD](#)) and the qubit pool.

The energy as a function of the iteration number, which corresponds to the number of operators in the ansatz, is plotted in figure 32a. As expected, we observe that qubit-[ADAPT-VQE](#) requires more iterations/operators to achieve the same accuracy as fermionic-[ADAPT-VQE](#).

In figure 32b, we can see that the tendency is reversed when we consider the error as a function of the number of CNOTs in the circuit, to which the circuit depth will be proportional. Qubit-[ADAPT-VQE](#) converges much faster than any of the fermionic pools in this case. Further, the molecule in study (H_4)

is rather small: its representation in a minimal basis set requires only eight spin-orbitals (eight qubits). It was mentioned before that there is a difference in the scaling of the circuit depth necessary to implement qubit pool operators (constant) versus fermionic pool operators (linear). As such, the difference is bound to become even more notable with the increase of the system size (number of molecular orbitals).

However, the error as a function of the number of CNOTs presented in 32b (and, in general, the circuit depth) does not represent the full picture, because measurement and variational parameter overhead are not considered. Because the qubit pool consists of fermionic operators broken into their constituent Pauli strings (and cleared of the anticommutation operator), it is significantly larger. For example, for H_4 , we already have 328 operators in the qubit pool, while SCGSD and SGS (fermionic pools) have respectively 111 and 66 operators. Since in each iteration the gradients of all the pool operators must be measured, this represents a sizeable measurement overhead. Further, the scaling of the number of variational parameters is as in 32a, which means that qubit-ADAPT-VQE will, for a given accuracy, imply a higher-dimensional (harder) optimization than fermionic-ADAPT-VQE. This implies more effort of the classical optimizer, which in turn will demand more calls to the quantum computer, resulting in even more measurements being required.

In essence, Qubit-ADAPT-VQE reduces the circuit depth necessary to reach a certain accuracy, at the expense of a greater burden on the classical optimizer and extra measurements on the quantum computer.

ADAPT-VQE: Exploring other Pool Options

In this chapter, [ADAPT-VQE](#) will be tested with several pools. The pools will be introduced here and differ both from the ones used in fermionic-[ADAPT-VQE](#) [35] and from the ones used in qubit-[ADAPT-VQE](#) [85].

The purpose is to explore the importance of the similitude between pool operators and fermionic excitations, and investigate the origin of such importance, analysing whether it stems from preserving quantities such as particle number and spin, from respecting antisymmetry, or from other factors.

The constitution of the [ADAPT-VQE](#) state in terms of Slater determinants, as well as the performance of the algorithm, will be analysed and compared between the different pools.

5.1 Motivation

The operators in the pools from fermionic-[ADAPT-VQE](#) are fermionic excitations that, when transformed into qubit operators, can consist of a superposition of up to 48 Pauli strings, the average length of each scaling linearly with the size of the system. These operators preserve particle number and spin, respect fermionic anticommutation, and properly represent fermionic excitations.

In qubit-[ADAPT-VQE](#), the operators are decomposed into the individual Pauli strings, and the Jordan-Wigner string is removed, causing their length to be independent of the system size. These operators do not preserve particle number or spin, do not respect fermionic anticommutation, and do not directly represent any proper fermionic excitation.

In the previous chapter, we saw that while fermionic-[ADAPT-VQE](#) converges faster as a function of the iteration number, qubit-[ADAPT-VQE](#) converges faster as a function of circuit depth. The qubit pools have more operators than the fermionic pools, which implies a higher overhead from measuring the gradients in each iteration; they will also typically require a larger number of optimizations and variational parameters for the same accuracy, demanding more effort from the classical optimizer and further increasing the number of calls to the quantum computer. In compensation for this, qubit pools reduce the circuit complexity of each operator, and result in shallower circuits for a given accuracy.

In the previous chapter, no options in-between the fermionic-inspired and qubit-inspired pools were tested. As well as spin-complemented and spin-adapted excitations, simple excitations can be used. These excitations can themselves be cleared of the Jordan-Wigner string while still preserving particle number and spin, and representing proper fermionic excitations up to the antisymmetry principle. And further, it is even possible to create operators that preserve (partially or in full) particle number and spin, but don't correspond directly to a specific fermionic excitation.

In order to assess the relative importance of each of these factors, several other pools are introduced and tested in this chapter. These pools find an intermediary position between those from fermionic-ADAPT-VQE and those from qubit-ADAPT-VQE. Because they have more operators than the former and less than the latter, their overhead due to gradient measurements is in-between these two algorithms; and because the number of Pauli strings in each operator is lower than in the former and higher than in the latter, the depth of their circuit implementation finds also an intermediary position.

We will begin by analysing the operators in the pool from qubit-ADAPT-VQE, as well as the impact that they have on the trial state. This is relevant for the definitions of the pools that ensue. After the pools have been introduced, the evolution of the state when each is used will be analysed. Finally, the performance associated with these pools will be compared.

5.2 Analysis of the Qubit Pool Operators

From the Jordan-Wigner transform in definition 2.57 and the fact that excitations are subtracted of their Hermitian conjugate to assure unitarity of the exponential operators, it is evident that any indices q, p will be represented by at most two different Pauli strings in the Qubit Pool (without Z strings):

$$i \cdot Y_q X_p, \quad i \cdot X_q Y_p. \quad (5.1)$$

And double excitations, q, p, s, r will be represented by at most eight different Pauli strings:

$$\begin{aligned} i \cdot Y_q X_p X_s X_r, & \quad i \cdot X_q Y_p X_s X_r, & \quad i \cdot X_q X_p Y_s X_r, & \quad i \cdot X_q X_p X_s Y_r, \\ i \cdot X_q Y_p Y_s Y_r, & \quad i \cdot Y_q X_p Y_s Y_r, & \quad i \cdot Y_q Y_p X_s Y_r, & \quad i \cdot Y_q Y_p Y_s X_r. \end{aligned} \quad (5.2)$$

Apart from the phase factor i , the coefficients do not matter because the final coefficients will be the variational parameters. These operators consist exclusively of Pauli X and Y operators, and the number of each is always odd. This is because the other terms are canceled out by the Hermitian conjugate, which is actually what leads to the unitarity of the exponentiated operators: the coefficient of these terms would be real.

It was mentioned that single excitations will lead to *at most* two Pauli strings, and double excitations to *at most* eight Pauli strings in the pool. This is because not all sets of spin-orbitals lead to valid excitations. For example, single excitations will only exist between spin-orbitals of the same type ($\alpha \rightarrow \alpha$ or $\beta \rightarrow \beta$). A similar restriction exists in double excitations; if in a set there is an odd number of β orbitals and α orbitals (e.g. p, q, r, s are of types $\alpha, \alpha, \alpha, \beta$), no double excitations involving these four orbitals will exist.

5.3 Analysis of the Qubit-ADAPT-VQE State

If we use the pool without Z strings, the i th operator in the Qubit-ADAPT-VQE pool consists of an exponentiated Pauli string ($\hat{P}_i = \bigotimes_N \sigma_k$, with $k \in \{0, z, x, y\}$) acting on N qubits (N will be two/four for singles/doubles respectively), multiplied by a real variational parameter θ_i and the imaginary unit i . We can expand this as

$$e^{i\theta_i \hat{P}_i} = 1 + \frac{i\hat{P}_i \theta_i}{1} + \frac{(i\hat{P}_i \theta_i)^2}{2!} + \frac{(i\hat{P}_i \theta_i)^3}{3!} + \frac{(i\hat{P}_i \theta_i)^4}{4!} + \frac{(i\hat{P}_i \theta_i)^5}{5!} + \dots$$

As do the individual Pauli operators, squared tensor products of them amount to the identity. As such, $(\hat{P}_i)^k$ is the identity if k is even, or simply \hat{P}_i if k is odd. As such, we can rewrite the expression into a simpler form.

$$\begin{aligned} e^{i\theta_i \hat{P}_i} &= \hat{I}^{\otimes 4} \left(1 - \frac{\theta_i^2}{2!} + \frac{\theta_i^4}{4!} - \dots \right) + i\hat{P}_i \left(\frac{\theta_i}{1!} - \frac{\theta_i^3}{3!} + \frac{\theta_i^5}{5!} - \dots \right) \\ &= \hat{I}^{\otimes N} \cos \theta_i + i\hat{P}_i \sin \theta_i \end{aligned}$$

Operators in the ansatz of Qubit-ADAPT-VQE then act on a computational basis state $|n\rangle$ in a very simple manner.

$$e^{i\theta_i \hat{P}_i} |n\rangle = \cos \theta_i |n\rangle + i \sin \theta_i \hat{P}_i |n\rangle \quad (5.3)$$

In the previous section we saw that each \hat{P}_i has an odd number of Y operators. Since $|n\rangle$ is a computational basis state, acting on it with such an operator will introduce a factor of $\pm i$. Other than that, the effect of \hat{P}_i is to flip the states of the four qubits it acts on (it must be remembered that there are no Z operators in \hat{P}_i). So we can rewrite 5.3 as

$$e^{i\theta_i \hat{P}_i} |n\rangle = \cos \theta_i |n\rangle \pm \sin \theta_i |m\rangle .$$

Here, $|m\rangle = i\hat{P}_i |n\rangle$ is a computational basis state that differs from the original (also computational basis) state $|n\rangle$ in and only in the states of the four qubits P_i acts on. $|m\rangle$, $|n\rangle$ represent Slater determinants with the same particle number if and only if the qubits acted on by P_i in $|n\rangle$ had an equal number of $|0\rangle$ and $|1\rangle$ states.

From this, it's obvious that all eight Pauli strings acting on the same two/four spatial orbitals (for singles/doubles respectively) listed in the previous subsection will have the same effect on a given computational basis state, except for a possible phase factor of -1 associated with $|m\rangle$. Operators acting on different indices may bring different Slater determinants (computational basis states) into the superposition among themselves; operators acting on the same indices will only differ in their impact on the wave function through interference.

For example, in the simple case that we have four spin-orbitals and two electrons, the Hartree-Fock state (using alternate orbital ordering) can be written $|0011\rangle$. The impact of each of the eight operators in the Qubit Pool associated with double excitations is

$$|0011\rangle \rightarrow \cos \theta |0011\rangle \pm \sin \theta |1100\rangle.$$

The sign is plus for $XXYX$, $XXXY$, $XYYY$, $YXYY$ and minus for $YXXX$, $XYXX$, $YYXY$, $YYYY$.

5.4 Alternative Pools

In order to test the relative importance of bringing new Slater determinants into the superposition *versus* manipulating interference effects, as well as the impact of particle number and Z spin preservation on convergence, [ADAPT-VQE](#) was tested with five other pools. The pools were purposefully defined to facilitate the analysis of the impact of these factors, separately, on the performance of the algorithm: they were chosen to have varying degrees of affinity with fermionic excitations. The introduction of each new pool follows.

5.4.1 Generalized Singles and Doubles

The [GSD](#) pool includes the simplest possible generalized excitations (that are not spin-adapted nor spin-complemented), up to order two. Of course, they're still forced to be anti-Hermitian.

Single excitations will consist of a linear combination of two two-body operators,

$$a_p^\dagger a_q - a_q^\dagger a_p, \quad (5.4)$$

and double excitations will consist of a linear combination of two four-body operators,

$$a_p^\dagger a_q^\dagger a_r a_s - a_r^\dagger a_s^\dagger a_p a_q. \quad (5.5)$$

5.4.2 Eight Pool

The Eight Pool consists of the same operators as the [GSD](#) pool, with the Jordan-Wigner string removed from all operators. It was designated this way because each operator will consist of a linear combination of at most eight Pauli strings.

As an example, the generalized double excitation defined in [5.5](#) gives rise to the Eight Pool operator in [5.6](#).

$$\begin{aligned} \hat{t} = & i(-X_q X_p X_s Y_r - X_q X_p Y_s X_r + X_q Y_p X_s X_r - X_q Y_p Y_s Y_r \\ & + Y_q X_p X_s X_r - Y_q X_p Y_s Y_r + Y_q Y_p X_s Y_r + Y_q Y_p Y_s X_r) \end{aligned} \quad (5.6)$$

We can factorize the operator as in 5.7.

$$\begin{aligned}
 \hat{\tau} &= -iX_qX_pX_sY_r(+1 + I_qI_pZ_sZ_r - I_qZ_pI_sZ_r - I_qZ_pZ_sI_r \\
 &\quad - Z_qI_pI_sZ_r - Z_qI_pZ_sI_r + Z_qZ_pI_sI_r + Z_qZ_pZ_sZ_r) \\
 &= -iX_qX_pX_sY_r(+1 + Z_qZ_pZ_sZ_r) \times \\
 &\quad (+1 + I_qI_pZ_sZ_r - I_qZ_pI_sZ_r - Z_qI_pI_sZ_r)
 \end{aligned} \tag{5.7}$$

From this, we can easily see the effect of $\hat{\tau}$ on computational basis states. It should be remembered that states $|0\rangle_k, |1\rangle_k$ indicate that the spin-orbital that qubit k represents is unoccupied and occupied, respectively (under the Jordan-Wigner mapping as defined in 2.57). It is also worth noting that since the qubits with indices p, q, s, r are the only ones acted on by the operator, we can ignore the state of the remaining qubits.

The factorized expression with Z operators acts as $8 \cdot \hat{I}_{qpsr}^{\otimes 4}$ for $|1100\rangle_{qpsr}$ and $|0011\rangle_{qpsr}$, and as 0 for the remaining computational basis states. The effect of it is then to block transitions that do not correspond to the proper excitation. Thus, $\hat{\tau}$ acts on computational basis states $|1100\rangle_{qpsr}$ and $|0011\rangle_{qpsr}$ to transform them into $+8|0011\rangle_{qpsr}$ and $-8|1100\rangle_{qpsr}$ respectively, and it acts as identity in the remaining computational basis states.

The original generalized double excitation in 5.5 acts to excite (or de-excite) electrons from spin-orbitals r, s to spin-orbitals p, q . The Jordan-Wigner transformed operator then has to be faithful to this action; it can't e.g. transform $|1010\rangle_{qpsr}$ into $|0101\rangle_{qpsr}$, as that would correspond to a different excitation ($q, s \rightarrow p, r$). Such an excitation might even not preserve S_Z . Other transitions might not preserve particle number: that is the case if $|0000\rangle_{qpsr}$ is transformed into $|1111\rangle_{qpsr}$, for example.

In a similar way to how we saw the (exponential) ansatz operators act on computational basis states when the pool consists of single Pauli strings with imaginary coefficients (equation 5.3), we can see what the effect of the ansatz operator corresponding to pool operator $\hat{\tau}$ (associated with the excitation $(r, s \rightarrow p, q)$ and coefficient θ is. On the computational basis states $|0011\rangle_{qpsr}$ and $|1100\rangle_{qpsr}$, it acts as

$$e^{\theta\hat{\tau}}|0011\rangle_{qpsr} = -\sin(8\theta)|1100\rangle_{qpsr} + \cos(8\theta)|0011\rangle_{qpsr},$$

$$e^{\theta\hat{\tau}}|1100\rangle_{qpsr} = +\sin(8\theta)|0011\rangle_{qpsr} + \cos(8\theta)|1100\rangle_{qpsr},$$

and it has trivial action on the remaining computational basis states.

The Eight Pool operators conserve particle number and S_z , and properly represent fermionic excitations, up to the anticommutation requirement. This was the pool used in the [Qubit-Excitation-Based \(QEB\)-ADAPT-VQE](#) introduced in [99], where they called such operators *qubit excitations* for obvious reasons (they are essentially excitations that don't respect fermionic antisymmetry).

5.4.3 One Pools

Pools of this type are a subset of the Qubit Pool (the one without the Jordan-Wigner strings). From any fermionic excitation, One Pools keep only one Pauli string (from the two in 5.1 or the eight in 5.2, in the case of single and double excitations respectively).

The One Pools will hereby be denoted after the strings that are kept from double excitations. For example, if from any four spin-orbital indices q, p, s, r only the string $i \cdot X_q X_p Y_s X_r$ is kept, the pool will be called **XXYX** pool. From single excitations, the string $Y_q X_p$ was the one kept always.

This type of pool does not conserve S_z or particle number, much like the original Qubit Pool. However, keeping only strings of a specific form greatly limits variational flexibility by narrowing the possibilities of control over quantum interference effects. As such, One Pools are likely to perform worse.

5.4.4 Two Pools

Two Pools are obtained from the previous One Pools by multiplying each operator arising from a single excitation by $(1 - Z_q Z_p)$, and each arising from a double excitation by $(1 + Z_q Z_p Z_s Z_r)$. The effect of this is to block some improper transitions between Slater determinants. When only two Pauli strings are used per operator, it is not possible to block them all.

Regarding double excitations, the effect is to avoid those transitions that attempt to act on a set of four spin-orbitals with an odd number of electrons: the action of Two Pool operators on the eight 4-qubit computational basis states with an odd number of 1 states is trivial. This is a natural option because it can be done with a simple parity check.

We can see that $(I_{qprs}^{\otimes 4} + Z_q Z_p Z_s Z_r)$ is one of the factors appearing in 5.7 that prevents improper excitations. Without the other factor, not all of them can be prevented. This one only prevents eight improper transitions between Slater determinants, and allows six improper ones in addition to the correct two.

In the case of single excitations, the effect is to block those transitions that attempt to act on a set of two spin-orbitals with an even number of electrons; this is enough to conserve particle number and S_z . In fact, the single excitation operators in Two Pools have exactly the same effect on the state as those in the Eight Pool.

However, aside from single excitations, these quantities are not in general conserved by the operators in this pool.

5.4.5 Four Pools

Four Pools are obtained from the Two Pools. The operators arising from single excitations are kept as before, and the ones arising from double excitations are multiplied by an operator of the form $(1 - Z_a Z_b)$. The indices a, b depend on the type of spin-orbitals (up or down) the operator acts on. The following

specifies what was the utilized factor for each type of spin-orbitals. Case $\alpha\beta\alpha\beta$, for example, is used to mean that the indices q, p, s, r correspond in this order to up, down, up, down spin.

Case A, $\alpha\beta\alpha\beta$ or $\beta\alpha\beta\alpha$:

$$I_{qprs}^{\otimes 4} - I_q Z_p I_s Z_r$$

Case B, $\alpha\alpha\beta\beta$ or $\beta\beta\alpha\alpha$:

$$I_{qprs}^{\otimes 4} - I_q I_p Z_s Z_r$$

Case C, $\alpha\beta\beta\alpha$ or $\beta\alpha\alpha\beta$:

$$I_{qprs}^{\otimes 4} - Z_q I_p I_s Z_r$$

Case D, $\alpha\alpha\alpha\alpha$ or $\beta\beta\beta\beta$:

$$I_{qprs}^{\otimes 4} - I_q I_p Z_s Z_r$$

$$I_{qprs}^{\otimes 4} - I_q Z_p I_s Z_r$$

$$I_{qprs}^{\otimes 4} - Z_q I_p I_s Z_r$$

This pool preserves S_z and particle number, without being completely faithful to a true fermionic excitation. To see this, we can analyse a specific case of spin-orbitals q, p, s, r - e.g. case **A**. If we had picked the $XXYX$ One Pool, the corresponding operator in the Four Pool will be

$$i \cdot X_q X_p Y_s X_r (I_{qprs}^{\otimes 4} + Z_q Z_p Z_s Z_r) (I_{qprs}^{\otimes 4} - I_q Z_p I_s Z_r).$$

It is easy to see that $(I_{qprs}^{\otimes 4} + Z_q Z_p Z_s Z_r) (I_{qprs}^{\otimes 4} - I_q Z_p I_s Z_r)$ is 0 when this operator acts on any but four computational basis states of the involved qubits (the state of the remaining qubits is irrelevant):

$$|0110\rangle_{qpsr}, \quad |1001\rangle_{qpsr}, \quad |1100\rangle_{qpsr}, \quad |0011\rangle_{qpsr}.$$

When we add the variational parameter and exponentiate the operator, it will be a conditional rotation $e^{i\theta X_q X_p Y_s X_r}$ applied only to these four states.

When acting on $|n\rangle$, $|n\rangle$ being one of these four computational basis states, this operator rotates it into a superposition $a|n\rangle + b|m\rangle$, where $|m\rangle$ is the same as $|n\rangle$, except qubits q, p, s, r have their states flipped. The new state appearing in the superposition has the same particle number and S_z as the original state, e.g.

$$|0110\rangle_{qpsr} \rightarrow a|0110\rangle_{qpsr} + b|1001\rangle_{qpsr}.$$

The action of the operator is trivial in the cases were this would not happen - e.g., the action of $e^{i\theta X_q X_p Y_s X_r}$ in $|1000\rangle_{qpsr}$ would be

$$|1000\rangle_{qpsr} \rightarrow a |1000\rangle_{qpsr} + b |0111\rangle_{qpsr} .$$

The particle number is altered by 2 and S_z is altered by 1 in the computational basis state $|0111\rangle_{qpsr}$. The two factors $(I_{qprs}^{\otimes 4} + Z_q Z_p Z_s Z_r)$, $(I_{qprs}^{\otimes 4} - I_q Z_p I_s Z_r)$ prevent all these improper rotations. However, they do not prevent transitions that do not correspond to the original excitation. This is the difference between the Four Pool and the Eight Pool.

For each set of four spin-orbitals q, p, s, r of type **A**, the Eight Pool will have two operators, each corresponding to specific excitations / de-excitations. The excitation $a_p^\dagger a_q^\dagger a_r a_s - a_r^\dagger a_s^\dagger a_p a_q$ (5.5) would give rise to an Eight Pool operator that acts non-trivially on two computational basis states:

$$\begin{aligned} |0011\rangle_{qpsr} &\rightarrow a |1100\rangle_{qpsr} , \\ |1100\rangle_{qpsr} &\rightarrow b |0011\rangle_{qpsr} . \end{aligned}$$

Corresponding to the excitation $r, s \rightarrow p, q$ and the respective de-excitation $p, q \rightarrow r, s$. The coefficients a, b are irrelevant because they will be multiplied by the variational parameter. Another excitation $a_p^\dagger a_s^\dagger a_r a_q - a_r^\dagger a_q^\dagger a_p a_s$ would correspond to an Eight Pool operator that acts non-trivially on two other computational basis states:

$$\begin{aligned} |1001\rangle_{qpsr} &\rightarrow c |0110\rangle_{qpsr} , \\ |0110\rangle_{qpsr} &\rightarrow d |1001\rangle_{qpsr} . \end{aligned}$$

Corresponding to the excitation $r, q \rightarrow p, s$ and the respective de-excitation $p, s \rightarrow r, q$. In contrast, there will be a single Four Pool operator acting on these four spin-orbitals. This is the operator we saw before, with non-trivial action on four spin-orbitals:

$$\begin{aligned} |0011\rangle_{qpsr} &\rightarrow e |1100\rangle_{qpsr} , \\ |1100\rangle_{qpsr} &\rightarrow f |0011\rangle_{qpsr} , \\ |1001\rangle_{qpsr} &\rightarrow g |0110\rangle_{qpsr} , \\ |0110\rangle_{qpsr} &\rightarrow h |1001\rangle_{qpsr} . \end{aligned}$$

This operator corresponds simultaneously to all excitations / de-excitations acting on this set of spin-orbitals that are possible (i.e. that preserve particle number and S_z).

In case **D**, the One Pool operator yields three operators in this pool. This is because, in the special case that all orbitals are of the same type, enforcing particle number preservation is enough to ensure S_z preservation. As a consequence, there are less restrictions in the excitations that are allowed to occur within these spin-orbitals.

5.4.5.1 Considerations

Because of the relations between Pauli matrices, the constituents of all of these pools are linear combinations of operators belonging to the original Qubit Pool. As compared to the latter, the number of operators in any of the One, Two, and Four Pools is decreased almost 8-fold, and almost 4-fold in the Eight Pool. The operators in these new pools consist of linear combinations of at most one, two, four and eight Pauli strings, in the same order; the increase in number implies an increase in affinity with proper fermionic excitations.

It's important to note that implementing the exponential of a linear combination of N Pauli strings does not necessarily correspond to an N -fold increase in circuit depth, as compared to the single Pauli strings in the Qubit Pool. Because they have a very specific form resembling controlled rotations, the exponentials of these newly introduced pools can be implemented more efficiently than those consisting of linear combinations of arbitrary Pauli strings. For example, reference [98] presented circuits for implementing the exponentiated Eight Pool operators (*qubit excitations*) with a CNOT count of 13 and depth of 11, against the 48 CNOT count and depth of the ladder-of-CNOTs [64] method. Figures 33 and 34 exemplify the two different circuit implementations for a specific operator (the operator in 5.8). The former illustrates the vanilla approach, while the latter shows the more efficient option.

$$\hat{\tau} = e^{-\frac{it}{8} (+X_0 Y_1 X_2 X_3 + Y_0 X_1 X_2 X_3 + Y_0 Y_1 Y_2 X_3 + Y_0 Y_1 X_2 Y_3 - X_0 X_1 Y_2 X_3 - X_0 X_1 X_2 Y_3 - Y_0 X_1 Y_2 Y_3 - X_0 Y_1 Y_2 Y_3)} \quad (5.8)$$

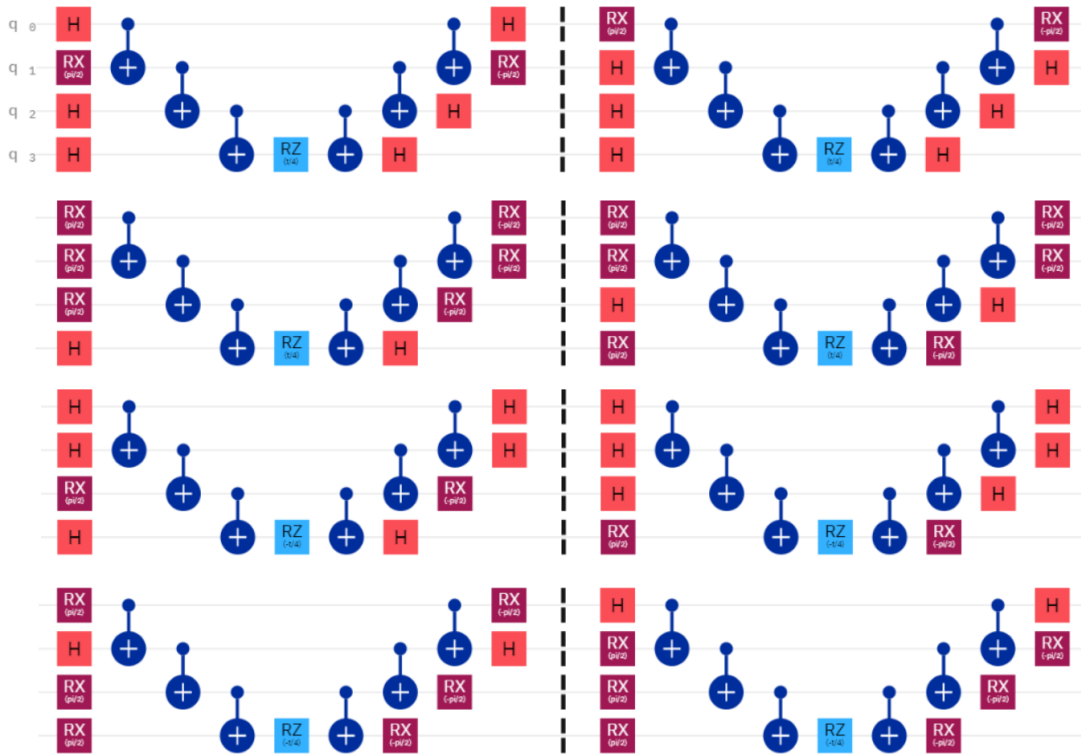


Figure 33: Circuit for implementing the operator in 5.8 using the ladder-of-CNOTs method explained in subsection 2.1.3. The circuit was drawn in the IBM Quantum Composer [41].

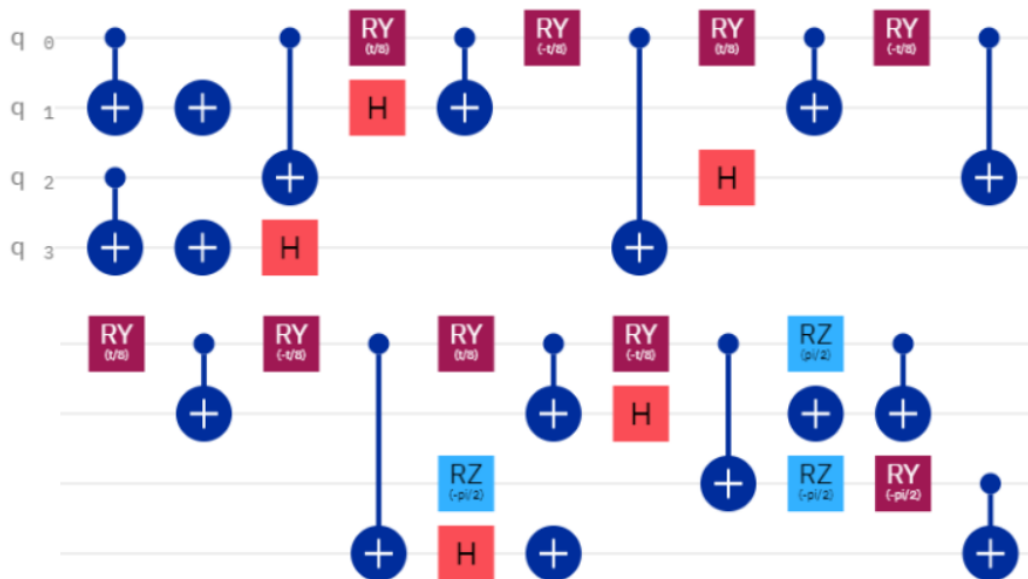


Figure 34: Circuit for implementing the operator in 5.8 using the more efficient method introduced in [98]. The circuit was drawn in the IBM Quantum Composer [41].

5.5 Graphical Analysis of the State Evolution

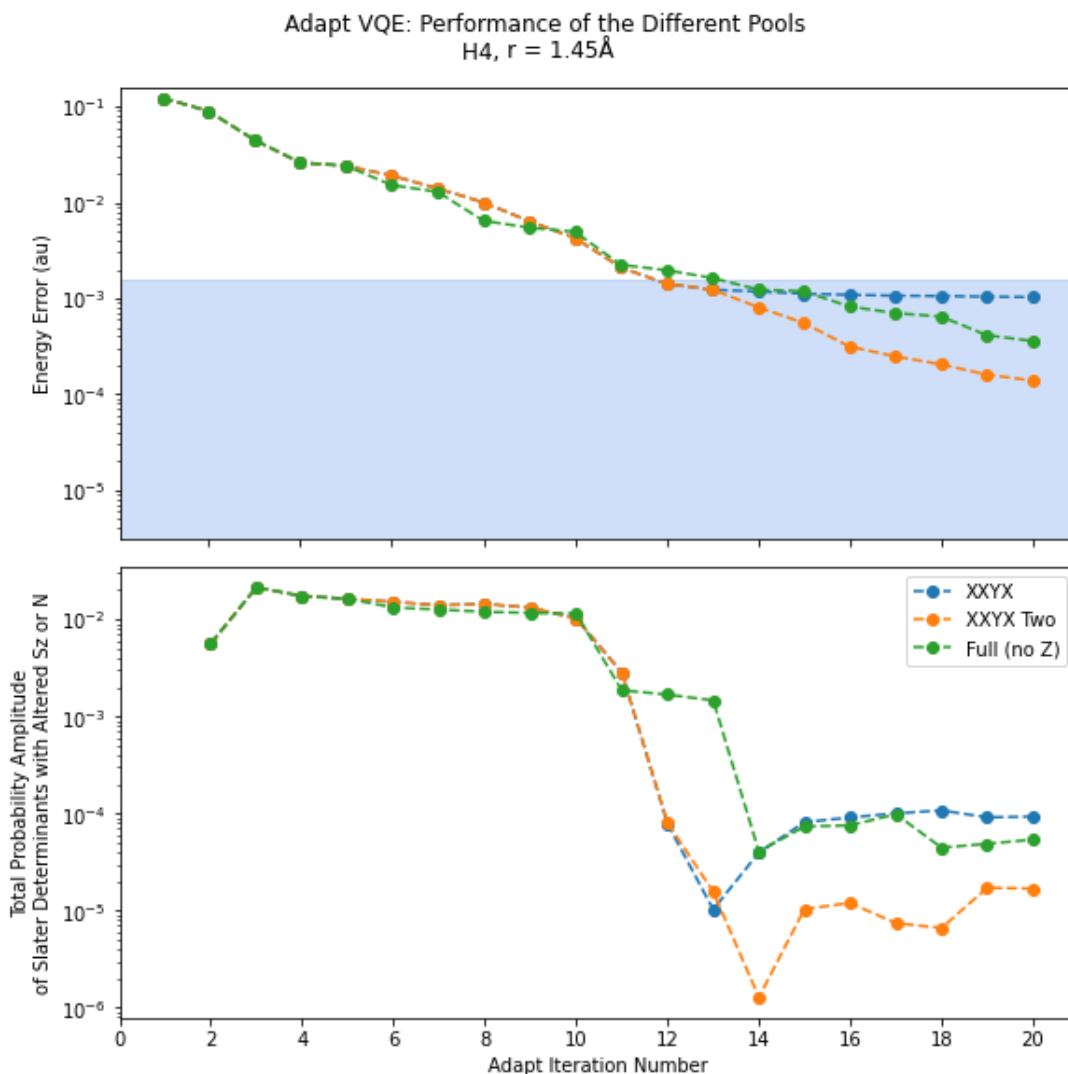


Figure 35: Evolution of the 20 first iterations of the [ADAPT-VQE](#) algorithm for three different pools: the XXYX One Pool, the XXYX Two Pool, and the original Qubit Pool (without the Jordan-Wigner string). The upper plot shows the evolution of the error in the energy. The lower plot shows the evolution of the total probability amplitude of computational basis states representing Slater determinants with an altered particle number or Z spin projection. The molecule in consideration is the H_4 molecule at an interatomic distance of 1.45Å.

Before comparing all of these pools, an illustration of the impact of ‘partial’ conservation of S_z and particle number on convergence is presented. Figure 35 shows a comparison of the performance of all of the pools that do not preserve in full any of these quantities. Other than the evolution of the error, the plot shows the significance of Slater determinants with altered Z spin projection and particle number in the wave function, for all three pools.

Even though the XXYX One Pool contains operators capable of bringing any Slater determinant into the wave function, it doesn’t seem to be capable of convergence, as the error stabilizes at around 10^{-3}

Hartree. The same does not happen for the original Qubit Pool: it becomes clear that the different Pauli string formats are important for indispensable interference effects. However, the efforts to partially preserve S_z and particle number in the XXYY Two Pool are clearly rewarded: this pool does not only match the performance of the Qubit Pool, it surpasses it. The upper and lower plots in figure 35 do indeed suggest a link between convergence and preservation of these quantities. It seems like wave functions that are more contaminated by Slater determinants with altered S_z and particle number imply a slower convergence.

It is interesting to observe that the curves in the lower plot of figure 35 are non-monotonic. All curves have a minimum at iteration 13 or 14. It seems as though as while in the beginning it is possible to get closer to the ground state while decreasing the probability amplitude of Slater determinants with altered S_z or N , after a point the energy can't be lowered further without increasing it. Regardless, this increase quickly comes to a halt, and does not seem to impede convergence.

Figure 36 shows the evolution of the ADAPT-VQE algorithm for four different pools defined in the previous section: the XXYY One Pool, the XXYY Two Pool, the XXYY Four Pool, and the Eight Pool. The evolution of the algorithm with the original Qubit Pool (without the Jordan-Wigner string) is also plotted for reference. Other than the evolution of error, the plots show the evolution of the number of Slater determinants (computational basis states) in the superposition state, split through three groups: those with altered S_z , those with altered particle number, and those with unaltered S_z and particle number. The purpose is to contrast the convergence speed against the composition of the state in terms of Slater determinants.

An interesting aspect all plots have in common is that by the end of these iterations, there are twenty fully correct Slater determinants in the state. In the case of the Eight Pool (figure 36d), this is enough to reach a very high accuracy (almost 10^{-7} Hartree, well inside chemical accuracy, of the order of 10^{-3} Hartree). Further, this number stabilizes after a certain number of iterations that ranges from 11 (in the Four Pool, figure 36c) to 15 (in the Qubit Pool, figure 36e); this stabilization does not seem to imply a stabilization of the error in general. This seems to confirm that the relevant limitation after this point is *not* the lack of computational basis states in the superposition that correspond to Slater determinants with relevance for the ground state, rather the lack of variational freedom with respect to their coefficients.

The error only does seem to stabilize after the 20 correct Slater determinants are reached when the XXYY One Pool is used (figure 36a). The reason was explained before: while this pool allows for enough degrees of freedom to bring all necessary Slater determinants into the superposition, it does not allow for enough degrees of freedom that play with interference effects. Restricting the format of Pauli strings to XXYY forbids the use of different operators in the same four spin-orbitals that would lead to different local phases (in this case, the only options are ± 1 , because the state is real). This should be contrasted with the results obtained using the original Qubit Pool (figure 36e). The composition of state in terms of correct and incorrect Slater determinants is remarkably similar between these two pools. However, the Qubit Pool continues to decrease the error after the number of Slater determinants has stabilized at 20. This is because the existence of different formats of Pauli strings for the same spin-orbitals (the eight in

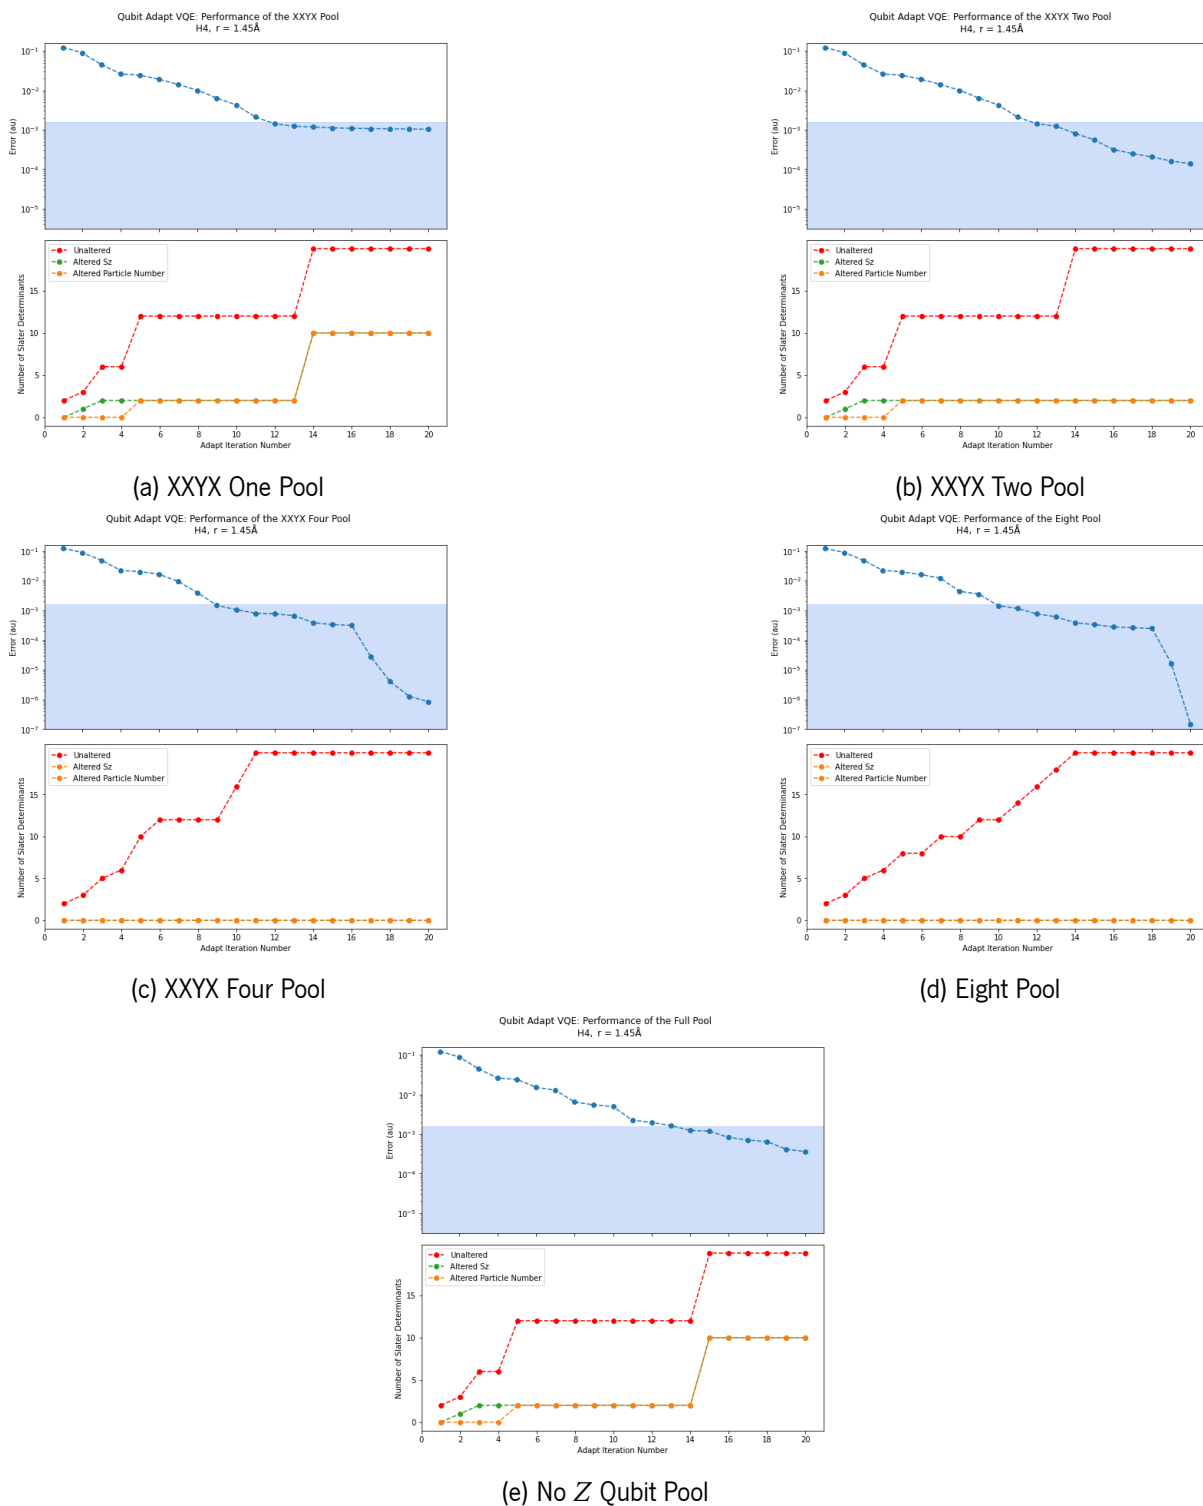


Figure 36: Evolution of the 20 first iterations of the ADAPT-VQE algorithm for different pools. The upper plot shows the evolution of the error in the energy. The lower plot shows the number of Slater determinants in the superposition state with altered particle number, S_z , or neither. In 36c and 36d, altered S_z and particle number correspond to the same colour because the curves coincide. When the green curve is not visible, it lies under the orange one. The molecule in consideration is H_4 at an interatomic distance of 1.45\AA .

5.2) empowers ADAPT-VQE with the ability to better control interference effects, and use them to converge: namely, it can leverage them to decrease the probability amplitude of wrong Slater determinants.

When we compare the XXYY One Pool against the corresponding Two Pool (figure 36b), we can find a reason for the better performance of the latter: while it does not fully preserve particle number and S_z , it decreases the amount of wrong Slater determinants in the state. While the One Pool finishes with ten Slater determinants with wrong S_z and particle number in the state, the Two Pool finishes with only two of each. The error decreases steadily up until the last iteration. Unlike in the Qubit Pool, the improved convergence is not due to different operators in the pool implying different local phases; the Two Pool operators associated with double excitations consist of a linear combination of Qubit Pool operators, but there is still only one operator for each set of four spin-orbitals (the size of the Two Pool is almost an eighth of the size of the Qubit Pool). The effect of multiplying the individual Pauli strings from the One Pool by $(1 + Z_q Z_p Z_s Z_r)$ is to preserve particle number and S_z in most computational basis states. While this does not provide ADAPT-VQE with more control over interference effects, it reduces the presence of incorrect Slater determinants in the state. Presumably, this decreases the need for the extra options of variational freedom, leading to improved convergence. In fact, the Two Pool even outperforms the Qubit Pool.

Finally, we can contrast the composition of the state associated with the Four Pool (figure 36c) against that of the Eight Pool (figure 36d). Their ease of convergence will be compared in the next section.

Both the Four Pool and the Eight Pool preserve S_z and particle number; the difference is that the operators in the latter are in direct correspondence with proper fermionic excitations. The former is actually faster to reach the 20 correct Slater determinants in the state, even though it has less and smaller operators in the pool (with an almost 2-fold reduction in both the size of pool and the number of Pauli strings per operator). The reason why this happens is clear from the previous section: the operators in the Four Pool and in the Eight Pool have non-trivial action on four and two Slater determinants, respectively. This is also the reason why the number of Slater determinants in the Eight Pool increases at most by two upon the addition of a single operator.

5.6 General Comparison of Pool Performance

Up to now, we have been analysing the relation between convergence and composition of the state in terms of Slater determinants. This section will abstract from such details and directly compare the performance of different pools, in light of the acquired knowledge.

Firstly, figure 37 compares multiple One Pools and multiple Two Pools. The purpose is to verify that the evolution of the error is not overly dependent on the chosen Pauli string format, and so the results presented previously for the One, Two, and Four pools constructed from XXYY strings hold in general. In fact, while there is evidently some variation between pools (because the operators do imply different local phases), the tendency seems identical between all plotted One Pools (YXXX, XYXX, XXYY, XXXY) and between all corresponding Two Pools.

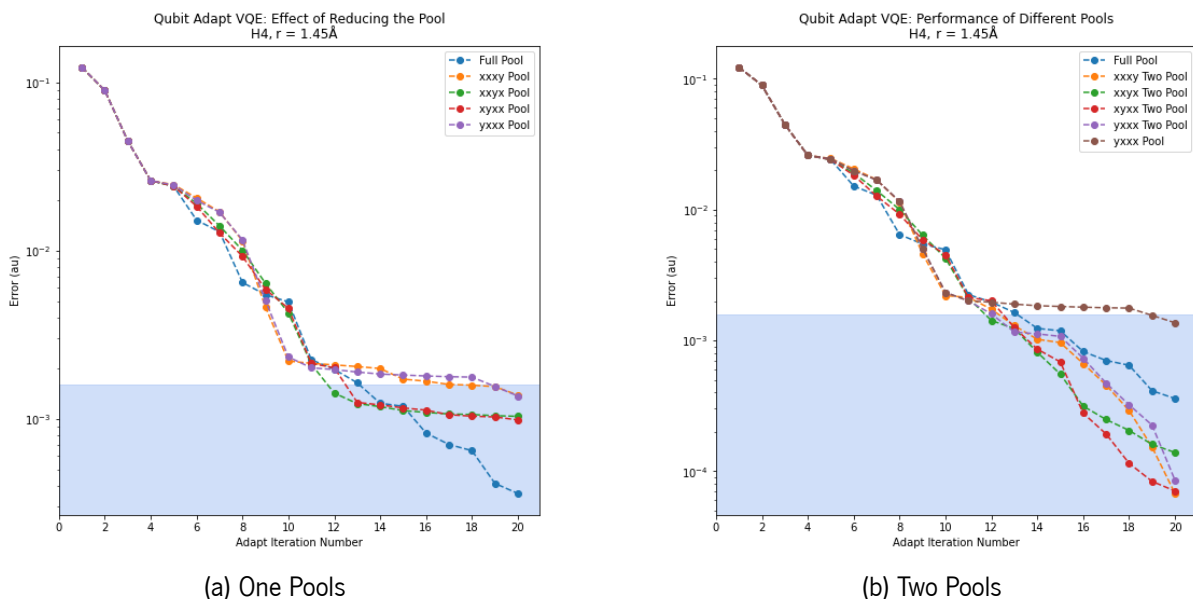


Figure 37: Comparison of the convergence of the ADAPT-VQE algorithm using several different One Pools (figure 37a) and Two Pools (figure 37b) against the full Qubit Pool (without Jordan-Wigner strings). In figure 37b, a curve corresponding to a One Pool was also plotted for reference. The error in the energy is plotted as a function of the iteration number. The molecule in consideration is H_4 at an interatomic distance of 1.45\AA .

In figure 37a, we can see that all these four One Pools perform worse than the original, full Qubit Pool. In fact, all of them seem to nearly stop decreasing the error after around iteration 10, while the Qubit Pool is steadily improving the approximation to the ground state. This seems to confirm that the different formats of Pauli string are necessary for interference effects to allow convergence, when a single string is used per operator.

In contrast with the performance of the One Pools, all the corresponding Two Pools (figure 37b) converge faster than the original Qubit Pool. This does indeed suggest that the preservation of S_Z and particle number in most computational basis states accelerates convergence in general, regardless of the string format used to construct the pool.

In figure 38, we can compare the performance of the ADAPT-VQE algorithm with the GSD Pool and with the Eight Pool. It should be remembered that the difference between the respective operators is that in the latter, all Jordan-Wigner strings responsible for the anticommutation of fermions are removed. As such, this figure is comparable with figure 31, that compared the performance of ADAPT-VQE using the full Qubit Pool with the Jordan-Wigner strings and without them. Once again, removing such strings barely seems to have an effect on the algorithm.

As happened with the individual Pauli strings from the Qubit Pool, removing the Jordan-Wigner strings here represents a relevant saving in the circuit depth. Both using the ladder-of-CNOTs method [64] and using the more efficient method of [98], the scaling of the depth necessary to implement the Eight Pool and the GSD Pool operators is respectively $\mathcal{O}(1)$ and $\mathcal{O}(N)$, N the number of spin-orbitals.




Figure 38: Comparison of the convergence of the **ADAPT-VQE** algorithm using the **GSD Pool** and the **Eight Pool**. The error in the energy is plotted as a function of the iteration number. The molecule in consideration is H_4 at an interatomic distance of 1.45\AA .

Finally, figure 39 compares the performance of all four newly introduced pools. The twenty first iterations of **ADAPT-VQE** for H_4 are plotted in 39a and the ten first for LiH are plotted in 39b. As expected, the observed tendency is an improvement of convergence with increased similarity with fermionic excitations. However, it is interesting to see how the **XXYX Four Pool** performs similarly to the **Eight Pool**, even though the operators consist of half the Pauli strings and the pool has nearly half the operators.

In the case of LiH (figure 39b), the curve corresponding to the **XXYX Four Pool** is not visible because it lies under the one corresponding to the **Eight Pool**. In the case of H_4 (figure 39a), the curves only overlap in some iterations. In most iterations in which this doesn't happen, the **XXYX Four Pool** has a visibly better performance (iterations 7-11 and 17-19). Only in iteration 20 does the **Eight Pool** ansatz prepare a better approximation to the ground state.

As there is no clear trend, tests on more molecules would be required to draw conclusions. It seems possible that as the energy plunges well into chemical accuracy, the **Eight Pool** allows fine-tuning of the wave function by allowing selection of specific double excitations, unlike the **XXYX Four Pool**, which only guarantees preservation of S_Z and particle number. However, since it allows action on more computational basis states at once, there is also reason to believe that the **XXYX Four Pool** may hasten convergence until points of extreme accuracy. For H_4 , this pool outperforms the larger **Eight Pool** until the error is lowered


(a) H_4 molecule at an interatomic distance of 1.45Å.(b) LiH molecule at an interatomic distance of 1.5Å.

Figure 39: Impact of the choice of pool (between One, Two, Four, and Eight Pools) in the **ADAPT-VQE** algorithm, for two different molecules. The error in the energy is plotted against the iteration number. In figure 39b, the XXXY Four Pool and the Eight Pool correspond to the same color because the curves coincide. These curves also overlap in some iterations in figure 39a.

to under 10^{-6} , which is roughly 0.07% of what is actually necessary to be inside chemical accuracy.

ADAPT-VQE: Further Ansatz Manipulation

This chapter aims to explore the possibility of introducing extra flexibility and deliberation into the evolution of the [ADAPT-VQE](#) ansatz, instead of having it defined by a single criterion of selection.

Two options will be outlined, having in common the intent of keeping the preparation circuit shallower than [ADAPT-VQE](#) (while attempting to avoid a prohibitive number of additional optimizations). The first one focuses on removing operators from the ansatz along the evolution of the algorithm, with the purpose of removing the ones that performed worse than expected. The second one attempts to be more conservative in the addition of new operators, in order to avoid getting bad performing operators in the ansatz to begin with. The approaches will be compared against each other, as well as against the original [ADAPT-VQE](#) algorithm, for multiple molecules and multiple pools. This comparison will be done by using the different approaches to obtain ansätze with the same number of operators, and analysing the error resulting from each in light of the number of optimizations accumulated along the execution (which is strongly related to the measurement cost).

6.1 Removing Operators

6.1.1 Motivation

The [ADAPT-VQE](#) algorithm selects the operators based on a single quantity: the magnitude of the derivative of the energy with respect to the operator coefficients θ_i , at point zero. This is a convenient selection method, since the gradient can be measured in a quantum computer. The circuits used for these measurements at a given iteration are barely deeper than the ansatz itself at that same iteration, thus guaranteeing that the selection method itself isn't the bottleneck of the algorithm. What is more, there being no further information, how much an operator is impacting the energy locally, at $\theta_i = 0$, is the best possible indicator of how much it will have lowered it once it has been added to the ansatz and had its coefficient optimized.

However, this is not a faultless selection method. While it seems likely that the best performing operator will be among those with the largest gradients, it is certainly possible that its gradient is not the largest one of all. Conversely, the operator with the largest gradient is not guaranteed to have a great impact on the energy. It might even be that the derivative of the selected operator was large at $\theta_i = 0$, but upon optimizing the coefficient and exploring the optimization landscape further, the optimizer will find that the change in energy quickly comes to a halt, as a local minimum is reached, and return a state with little energy difference as compared to the previous one. Of course, there is no way to know in advance, and this criterion allows for a much more informed choice than random selection.

But in spite of their actual impact on the energy, once they are chosen by their gradient, [ADAPT-VQE](#) simply accepts the operators. The selected operators remain in the ansatz until the algorithm terminates, regardless of how much they actually lowered the energy. As a consequence, it is possible that some operators in the ansatz will bring little advantage in converging to the ground state, but nevertheless contribute to deepen the circuit (placing more demands on the quantum computer) and increase the number of variational parameters (placing more demands on the classical optimizer, which in turn will require more function evaluations until convergence, leading to more calls to the quantum computer).

The problem of finding the optimal ansatz is combinatorial in nature. The best solution could be found by exhaustive search: if we tested *all* n -operator ansatze that could be formed from the operators in our pool, we could find the optimal one. Of course, this approach is not viable due to its computational complexity - hence, [ADAPT-VQE](#) settles for a procedure of growing the ansatz that, while not perfect, is at least tractable. However, there is a lot of space between [ADAPT-VQE](#) and a brute-force solution. It is possible to keep the algorithm tractable, but nevertheless be stricter in allowing the operators to actually remain in the ansatz, instead of simply keeping them once they are selected.

This subsection builds on the idea of removing bad performing operators from the ansatz along the execution of the algorithm. An attempted solution to the problem of finding a removal procedure that reaches a lower error while keeping the optimization overhead to a minimum will be developed.

6.1.2 Heuristics

The heuristics behind removing terms has to address two major points: when to remove operators, and when to allow them to be added back. If not done carefully, the process can actually harm the evolution of the algorithm by requiring unnecessarily many additional iterations, and consequently also optimizations and measurements.

The purpose of the following exposition is to develop a set of rules to remove operators which satisfies two conditions: resorting exclusively to data that is available during the process of running [ADAPT-VQE](#), and resulting in a modest overhead in optimization/measurement costs.

6.1.2.1 Removal Criterion

The first question that should be asked is:

When should an operator be susceptible to being removed?

A sensible first answer is *when it performs badly*, in the sense that it doesn't lower the energy as much as it could. However, defining *bad performance* is in itself cumbersome.

In a realistic run of the algorithm, we never know how far the current energy lies from the [FCI](#) ground energy. As such, we are ignorant of what consists a reasonable change in energy. What is more, *reasonable* might depend on the particular choice of pool: naturally, large operators as one encounters in the [SGSD](#) pool will have a greater capacity of changing the energy as compared to the length four Pauli strings in the qubit pool.

A natural option is to characterize performance *relative* to the other operators. Since the operators used in a particular run all belong to the same pool, the latter problem is solved. However, attention must be paid to which operators we are using as standard for the comparison.

As the state grows closer to the true ground state, and the energy grows closer to the true ground energy, there will be a decreasing trend in the impact of the newly added operators on the energy. It becomes clear that it is not fair to compare an operator to its predecessors in the ansatz: having acted on a state farther away from the ground state, they were in a better position to lower the energy.

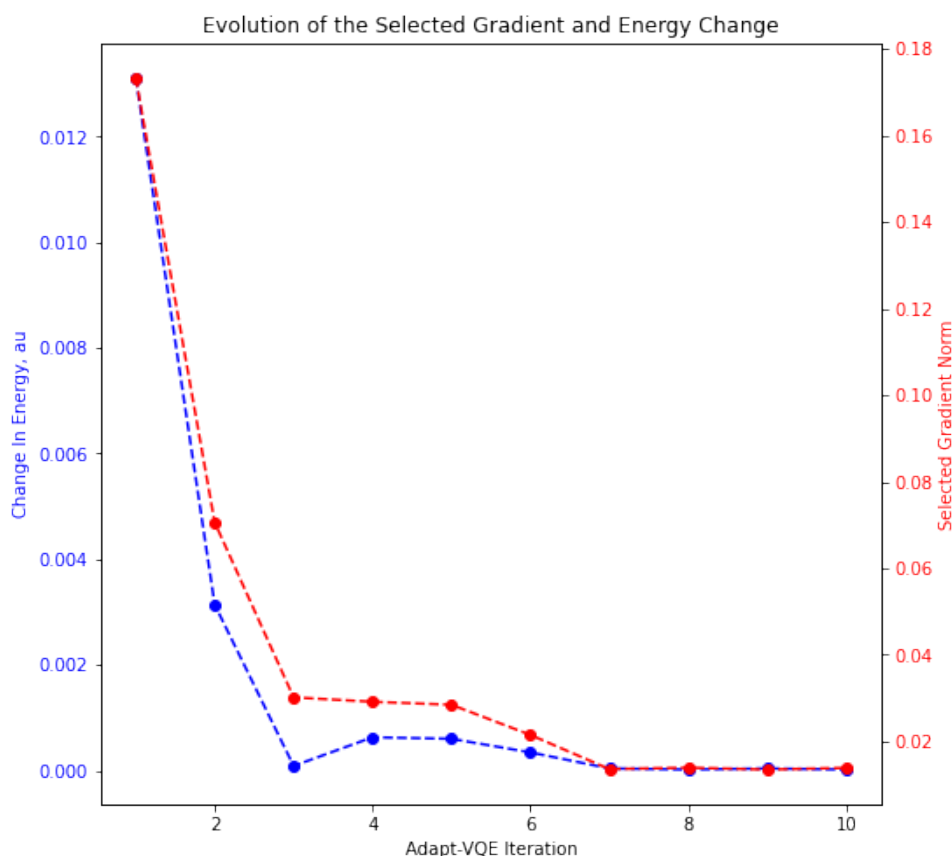


Figure 40: Norm of the selected operator and change in energy in the first ten iterations of [ADAPT-VQE](#). The molecule in consideration is the *LiH* molecule at an interatomic distance of 1.45Å. The [SGSD](#) pool was used.

This is better illustrated by the plot in figure 40, that shows the change in energy in the first ten iterations of the [ADAPT-VQE](#) algorithm for *LiH*. The change in energy is taken as compared to the previous iteration, with the energy at iteration 0 taken to be that of the reference state (the Hartree Fock state).

As expected, the tendency is for the impact that the operators have on the energy to decrease significantly along the evolution of the algorithm. As such, not lowering the energy as much as its predecessor seems like a highly inadequate criterion for removing an operator: taking this case, for example, it would imply removing almost all the operators in the ansatz.

Another, more viable possibility would be comparing two operators by their energy change *relative*

to their gradient. After all, by setting it as the selection criterion, we are using the gradient as the sole indicator of how much we expect a given operator to change the energy. The absolute value of the ratio between the change in energy caused by an operator and its the gradient seems like an interesting metric for characterizing the performance of an operator against expected. In addition, the selected gradient norm has, as does the energy change, a tendency to decrease as the algorithm evolves. This can be confirmed from figure 40, and understood intuitively: as the possible impact on the energy by the operators decreases, so does the absolute value of the maximum gradient, as a measure of how much the operator is affecting the energy at a certain point. Even though, as it was mentioned, the norm of the gradient of an operator isn't a rigorous indicator of how much it will impact the energy, this statement certainly holds in a broad sense, and the tendency is verified in practice.

The norm of the ratio between the energy change produced by an operator and its gradient at the time it was selected will hereby be denoted *performance ratio* for simplicity (formula 6.1).

$$\text{performance ratio} = \left| \frac{\text{energy change}}{\text{gradient}} \right| \quad (6.1)$$

For all that was exposed, one could imagine that the performance ratio would facilitate comparing an operator to its predecessor. However, numerical simulations show that this ratio also has a tendency to decrease along the evolution of the algorithm. Figure 41 shows its evolution along an ADAPT-VQE run. Once again, if one took the decrease in this ratio as compared to the previous iteration as the criterion for removing the operator that was added last, a very large amount of operators would be removed. This would warrant, for example, removing the operator added at iteration 2. However, this is the second best performing operator in all ten iterations. The operator added at iteration 7 would also be removed, as there is a sharp drop in the performance ratio at that iteration; however, the ratio does not go back to previous values in iterations 8, 9 or 10. Looking at the big picture, it does not seem like operator 7 performed particularly bad for its position on the ansatz.

It seems therefore difficult to remove operators by comparing them with their predecessors. However, when we add a new operator to the ansatz, the only information we have available is precisely that regarding the operators previously added: as for the remainder of the operators in the pool, they were not added to the ansatz, which means that there is no way to know how much they're capable of lowering the energy.

The possibility that is left is to compare an operator against its successors. If an operator added to the ansatz later on has a larger impact on the energy, there is a reason to believe that, up until that point, it was reasonable to expect changes in energy of at least that order of magnitude.

Of course, this is a simplified view of things: it is also possible that the operators added previously were necessary for the energy change caused by the new operator to be as large as it was. It is important to remember that the operators are selected based on their energy gradient taken on the state *at the point at which they are added*. So it is even possible that, had the previous operator not been there, the choice of operator would have been different altogether.




Figure 41: Absolute value of the performance ratio in the first ten iterations of [ADAPT-VQE](#), for the same pool and molecule as [figure 40](#). The performance ratio at a given iteration corresponds to the operator added then.

As a specific example, an operator might add to the state Slater determinants that did not appear in the superposition before. The term added afterwards can have relied on its impact on those determinants to change the energy as much as it did. In that case, even if the previous operator performed badly, removing it from the ansatz would prove to be a bad decision, as the energy would increase significantly due to the change that this operator produced not on the energy, but on the state.

However, there is an easy way to discover whether removing an older operator is a bad decision: removing it, re-optimizing the energy, and comparing the increase in energy from removing the operator with the decrease that had been caused by adding it. If the increase is significantly larger, then the operator is more important to the ansatz than one could have imagined from its impact on the energy.

It is then interesting to contemplate the possibility of removing an operator at a later iteration, once a *successor* has a greater impact on the energy, rather than immediately in the iteration at which it was added. Looking once again at [figure 40](#), it can be seen that the operator added at iteration 3 has quite lower an impact on the energy than its immediate successor, added in iteration 4. As much as we don't

have a view on how the future operators will impact the energy as the algorithm proceeds, at iteration 4 we can certainly look back at the previous iterations and notice that the third one produced a remarkably low difference in energy as compared to the fourth one. Looking at the numerical data, one finds that the energy change of iteration 4 is 739% that of iteration 3.

Considering that removing terms also slows down the evolution of the algorithm, and implies an added cost in measurements and optimizations, it is important that the criterion for removal is chosen to be strict enough that reasonably performing operators will not be removed. If operator A is followed by operator B that produces 101% its energy change, it doesn't seem fruitful to remove operator A: that would demand an extra optimization step after removing it, possibly only to add it back to the ansatz soon (how and when this should happen will be discussed shortly) and re-optimize all the coefficients once more. This would imply two extra optimizations for little to no gain.

With all of this in consideration, the following rules were chosen as the procedure for removing operators:

- Upon adding an operator to the ansatz at iteration j , compare the resulting energy change ΔE_j with that of all previous iterations $i < j$.
- If an operator added at iteration i caused an energy change ΔE_i such that $\Delta E_i > r\Delta E_j$ ¹, attempt to remove this operator and re-optimize the coefficients of the ansatz without it. Calculate the energy change produced by this action, $\Delta E'_i$.
- Compare the increase in energy caused by removing the operator with the decrease that had been caused by adding it². If $-\Delta E'_i > t\Delta E_i$, remove the operator.

The parameter r , $0 < r < 1$, determines what energy change is considered acceptable for previous iterations when a new operator is added and a new energy change is known of. As mentioned previously, a value of r too close to 1 would often cause the removal process not to compensate its measurement/optimization overhead. In the following results, $r = 0.5$ was used. This means that when an operator is added to the ansatz and produces a certain change in energy, all its predecessors that lowered the energy by under 50% of this are liable to be removed.

Once an operator is considered suitable for removal, the parameter $t > 1$ determines the tolerance for the difference in impact that there might be as a consequence of this operator not being the last one in the ansatz, and as such affecting the energy also indirectly through its successors. In the simulations, $t = 1.5$ was used. This means that, once the removal has been attempted, an operator is effectively removed if the increase in energy resulting from removing it is no greater than 150% the decrease in energy that had been caused by adding it. As explained before, not using this limit would make the process of removing terms blind to the possibility of operators with little impact in the energy actually being important for the wave

¹It should be noted that the $\Delta E_{i,j}$ are expected to be negative.

²Being caused by the removal of an operator, $\Delta E'_i$ will usually be positive.

function that was built after them. One could also imagine different alternatives: for example, comparing the change in energy from removing the operator i with the change in energy from adding operator j instead (removing operator i if $-\Delta E'_i > t\Delta E_j$), or the sum of both ($-\Delta E'_i > t(\Delta E_j + \Delta E_i)$), adjusting the parameter t accordingly. In essence, what we want to do is guarantee that the operator is only removed if there is something to be gained from it; that can be assured in several different ways.

6.1.2.2 Blocking Operators

Once the question of when to add operators has been dealt with, it is important to decide when and where they will be allowed to be added back to the ansatz.

At the point we've removed an operator, we know that even though its gradient was high, it didn't impact the energy as much as we could have hoped. Thus, it seems unwise not to make sure it doesn't get picked again too soon. This is likely to happen if we use the gradient as the sole selection method.

It is possible that the removed operator doesn't have the highest gradient anymore, since another operator has been added to the ansatz and so the state in which we are calculating the gradient is different; however, in practice, that is seldom the case. Usually, the change that a single operator produces on the wave function is not enough to cause significant changes to the gradients of the other operators in the pool. If nothing is done, a removed operator will almost always be added back to the ansatz in the next iteration, and once again have little impact on the energy. This behaviour is not desirable, as it prevents the removal of the term from being advantageous at all: once the term is removed, we should be giving the opportunity of being in the ansatz to other operators, that despite having lower gradients might perform better.

One solution could be blocking removed operators, taking them off the pool altogether and not allowing them to be added to the ansatz ever again. But this is another extreme that might jeopardize the performance of the algorithm: as was discussed, as the ansatz evolves, the change in energy produced by each operator will have a tendency to decrease. What consists a bad performance early on might be a great one later.

Taking figure 40 once again, we see that while the energy change at iteration 3 is the smallest in the first 6 iterations, it seems comparable to the change observed in iterations 7 to 10. In fact, the numerical data shows that all these iterations have associated energy changes of the order of 10^{-5} , with the change in 3 being the greatest one of all. What is more, it was verified that if the operator added in iteration 3 is removed and added back after a few iterations, the energy change it produces is very similar to the one it had produced back in position 3 - meaning that even after being removed, it seems like the best possible addition to the wave function at any of iterations 7 to 10. As an example, removing this operator and adding it back at iteration 8 would produce an energy change corresponding to 97% the change it had produced back at iteration 3. Importantly, at this point the removed operator is capable of producing more than double the change in energy of the operator added to the ansatz in case this one is not allowed to reenter it at all. This is only due to worsen as the ansatz evolves: the operator that was removed might be

crucial to convergence after a certain point.

In dealing with terms that have been removed, there are thus two main requirements: they can't be allowed to be added back too soon, and they can't be forbidden from being added back forever.

An approach that seems reasonable is blocking the operators that are removed until an operator is added that performs similarly. Since operators are removed when they have little impact on the energy as compared to operators added after them, it seems natural that they would be allowed back once a newcomer has an impact comparable to the one they had.

There are, however, a few problems with this method, that are better illustrated by an example. The following is a depiction of the evolution of an actual simulation of ADAPT-VQE with term removal (again for the LiH molecule at an interatomic distance of 1.45Å, using the SGSD pool). Aside from showing the drawbacks of blocking operators, this will serve to demonstrate how the term removal procedure flows. The criterion for removing operators is as defined before, and a removed operator i is allowed to be selected again once an operator k added to the ansatz later produces an energy change that is no more than the double the one operator i had produced ($\Delta E_i < 0.5\Delta E_k$).

	Ansatz			
Position	0	1	2	
Index	193	187	25	
Gradient (x 10⁻³)	173	71	30	
ΔE (x 10⁻⁵)	-1311	-314	-9	\emptyset

Figure 42: Relevant data structures at iteration 3 of the ADAPT-VQE algorithm, with an additional term removal feature based on blocking and unblocking operators. The molecule in consideration is *LiH*, and the pool the SGSD pool. The units of energy, omitted to avoid overcrowding of the figures, are Hartree throughout the whole chapter. Because the variational parameters are dimensionless, the gradient shares the energy units.

Figure 42 represents the relevant data structures at iteration three of the algorithm. The most important one is the ansatz, that specifies the current state preparation circuit.

As the name suggests, *position* represents the position of each operator in the ansatz: the column labeled with 0 concerns the first operator in the ansatz, the one labeled with 1 concerns the second one, and so on. The first operator is both the one that was first added and the one that comes first in the circuit.

The index identifies which operator in the pool was selected; what fermionic operator this corresponds to is irrelevant, as the purpose is just attributing an identifier to each pool operator. For a matter of simplicity, this identifier is a number rather than a formula.

The first three iterations are uneventful: the criterion for attempting to remove a term is never met. As can be seen in figure 42, the energy change always decreases from an iteration to the following one. As

such, no operator is suitable to be removed, and the algorithm proceeds just as the original ADAPT-VQE. The set of blocked operators is empty.

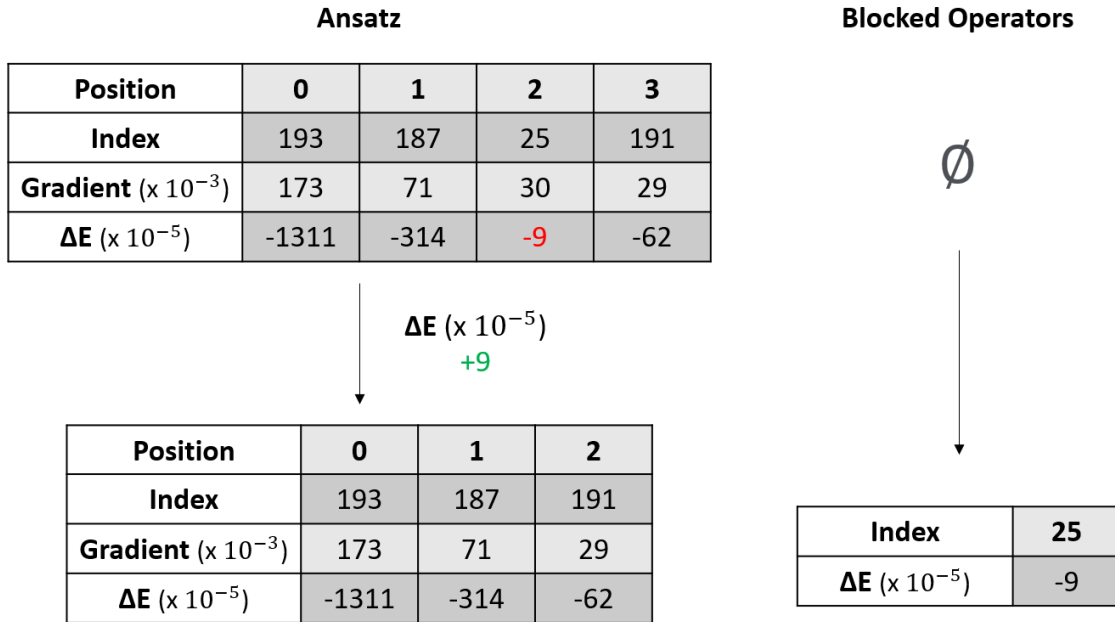


Figure 43: Relevant data structures at iteration 4.

Finally, at iteration 4 (figure 43), the condition necessary to attempt to remove an operator is met. The condition reads

$$\Delta E_i > r\Delta E_j \Leftrightarrow -9 > 0.5 \times (-62),$$

which does indeed hold: the operator 191 (j), added in this iteration, produces a change in energy significantly higher than that of operator 25 (i), added to the ansatz in the previous iteration.

As such, operator 25 is removed from the ansatz, and the coefficients of the remaining operators are re-optimized. From this, it is possible to calculate the increase in energy caused by removing this operator, which turns out to be 9×10^{-5} a.u.. The condition for effectively removing the operator reads

$$-\Delta E'_i > t\Delta E_i \Leftrightarrow -9 > 1.5 \times (-9),$$

which once again holds. It can be concluded that the impact of this operator in the energy isn't much changed by the posterior addition of operator 191 into the ansatz. This will not necessarily always be the case: sometimes, one might find that removing an operator from the middle of the ansatz will cause a drastic increase in energy, even if adding that operator in the first place had barely lowered it. This is a consequence of the presence of the operator in the ansatz being crucial for the impact on the energy of the next ones, and possibly decisive for their selection.

By the end of iteration 4 operator 25 is blocked, so as to seek operators that might have a greater impact on the energy, even if their gradient is lower.

Ansatz							Blocked Operators	
Position	0	1	2	3	4	5	Index	25
Index	193	187	191	188	184	64	ΔE ($\times 10^{-5}$)	-9
Gradient ($\times 10^{-3}$)	173	71	29	28	22	14		
ΔE ($\times 10^{-5}$)	-1311	-314	-62	-60	-35	-4		




Figure 44: Relevant data structures at iteration 7.

The algorithm proceeds without noteworthy developments until iteration 7. At iterations 5 and 6, operators 188 and 184 are added. They lower the energy roughly 7 and 4 times more than 25 (the blocked operator) respectively. This is a sign that the removal was worthwhile. Finally, at iteration 7, an operator is added that causes a significantly smaller change in energy. The condition for unblocking operator 25 reads

$$\Delta E_i < 0.5\Delta E_k \Leftrightarrow -9 < 0.5 \times (-4),$$

so that operator 25 is released and allowed to be selected again in future iterations. The set of blocked operators is once again found empty.

Unsurprisingly, 25 is the operator selected in the next iteration (iteration 8). It is interesting to note how the gradient and energy change of this operator are now, at iteration 8 (figure 45), very similar to those back at iteration 3 (figure 42), when the operator was first added. Despite the position in the ansatz being changed (from 2 before, to 6 now), the effect of the operator on the ansatz seems to be about the same.

Since the previous operator (64) had produced an energy change of -4×10^{-5} a.u., roughly half of that seen in iteration 8, that operator is removed from the ansatz. After re-optimizing the coefficients, the energy has only increased by 4×10^{-5} a.u., so the algorithm goes forward with the removal. At the end of iteration 8, operator 64 is blocked.

By this iteration, one can already notice undesirable behaviour. Operator 25 had been removed and blocked, but hadn't it been for that, it would precede operator 64. The removal criterion is based on subsequently added operators, but if they are only added later because they had been removed from an earlier position in the ansatz, this seems rather inadequate.

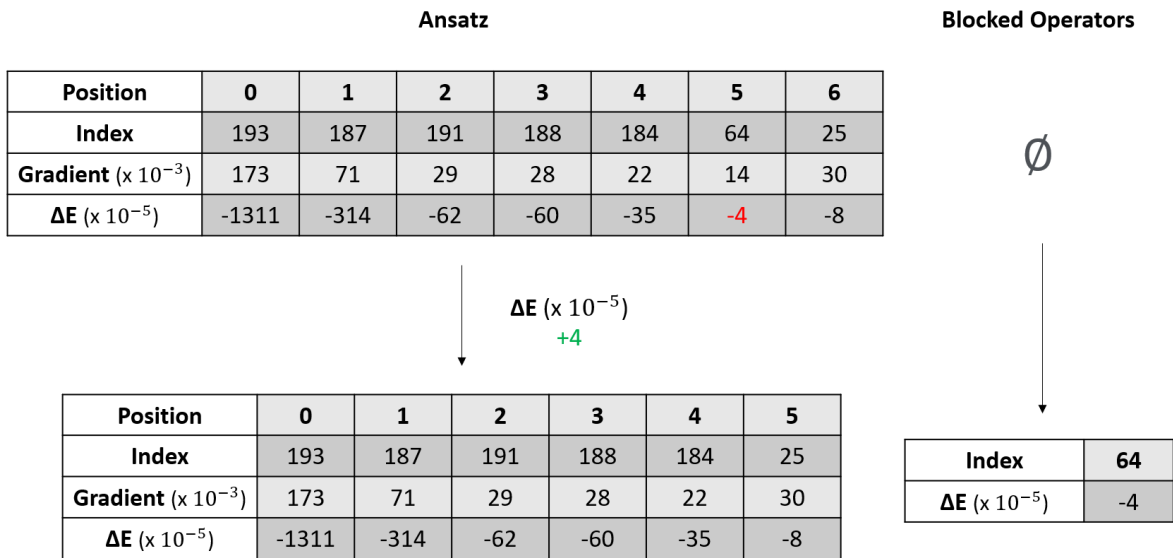


Figure 45: Relevant data structures at iteration 8.

An easy solution for this would be not allowing operators to be removed upon the addition of a previously blocked operator. However, this is not the only problem with the approach, as will become clear soon.

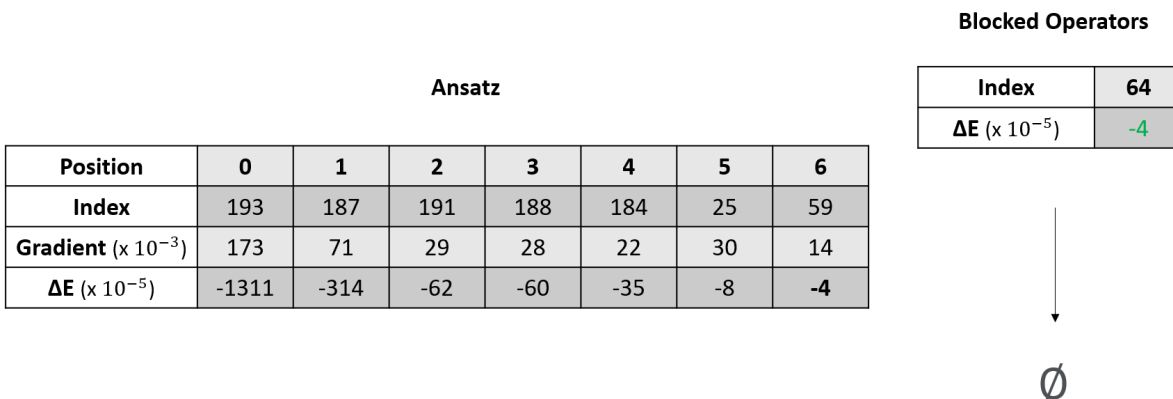


Figure 46: Relevant data structures at iteration 9.

By iteration 9 (figure 46), operator 59 is added. This operator produces an energy change similar to the one operator 64, added in iteration 7 (figure 44) and blocked in the ensuing iteration (figure 45). As such, the operator is unblocked.

Nothing remarkable happens at iteration 10: operator 29 is added, and no operator is suitable to be removed upon this addition. That changes by iteration 11 (figure 47), when the previously removed operator 64 is added back to the ansatz. At this point, operator 29 is removed and blocked.

Once again, it should be pointed out that this is really inefficient. Operator 64 had been removed from being outperformed by an operator that belonged to an earlier point of the ansatz, and this proved

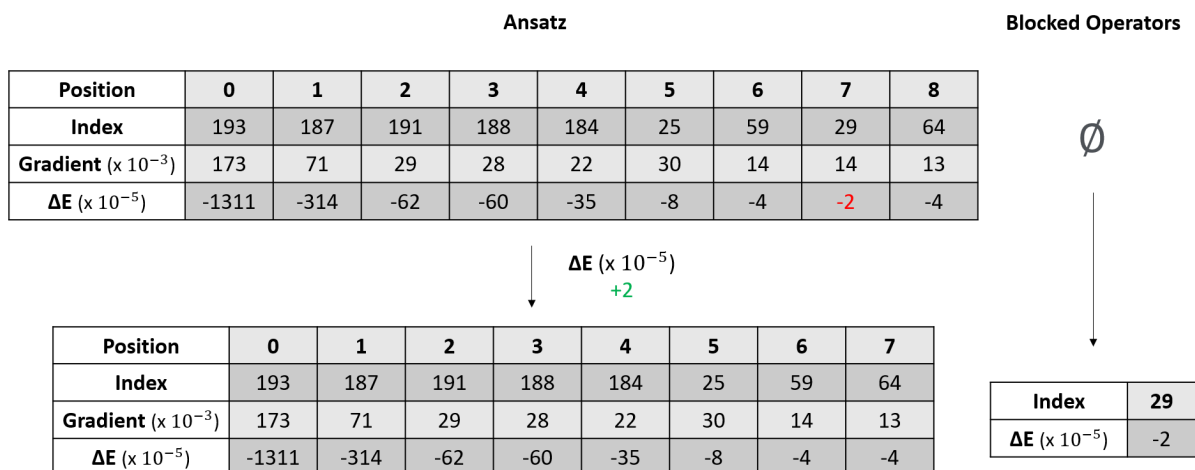


Figure 47: Relevant data structures at iteration 11.

to be a bad choice: it was added shortly after, with no benefits and a cost of two extra optimization steps. Now, this very operator causes the same problem: it removes operator 29 from the ansatz, even though it would have come before it if we hadn't removed it.




Figure 48: Relevant data structures at iteration 12.

In fact, this operator (29) is unblocked right in the following iteration (iteration 12, figure 48), once another operator is added (32) that causes about the same impact in the energy. And it will be added again immediately, by iteration 13. Once again, removing the operator costed us two extra optimization steps, and had no effect other than moving it a couple of positions forward in the ansatz (with no benefit in the energy).

At iteration 14 (figure 49), something very interesting happens: the selected operator (35) causes an energy change of -34×10^{-5} a.u., a change so large that it is only surpassed by that of iteration 6. Despite having a lower gradient than any other operator so far, it largely outperforms any of the five last operators,

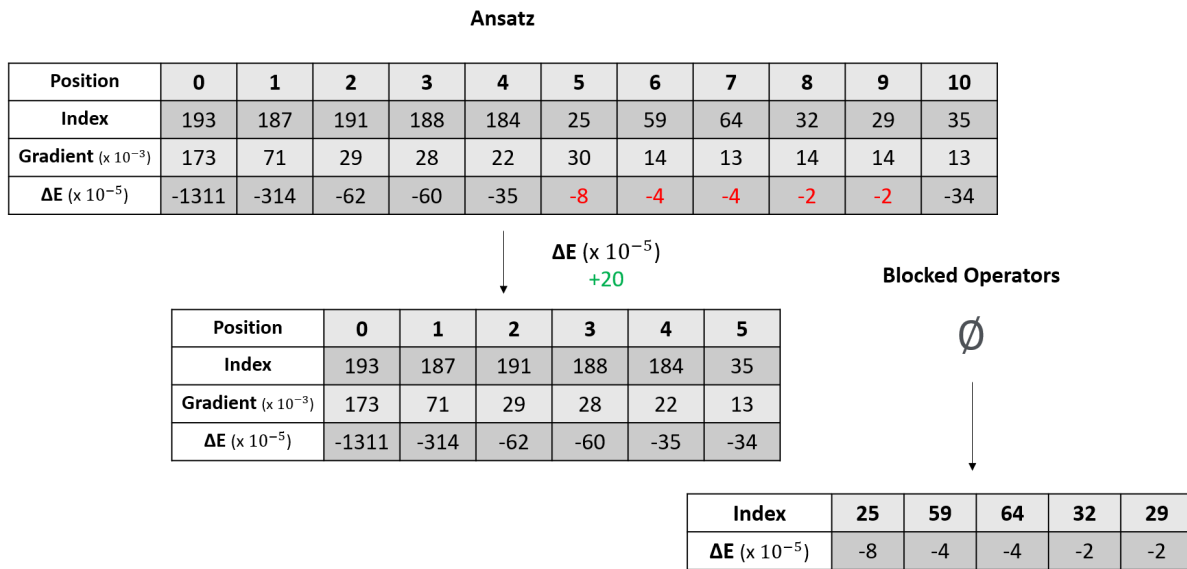


Figure 49: Relevant data structures at iteration 14.

with an energy change of up to almost 20 times that of its predecessors.

Thus, unprecedentedly, there are five operators that are liable to be removed. They are removed one by one, the coefficients being re-optimized each time; and each time, it is verified that the operator can be effectively removed. This is remarkable: removing these five operators caused an increase in energy roughly equal to the sum of the decreases that they had individually caused. This means that operator 35, added only in this iteration, was by far a better candidate to hold position 5 of the ansatz than any of those in positions 5-9 up to now. However, it was not picked until as late as iteration 14, because its gradient was not large enough.

Looking at figure 49, and focusing especially on the data regarding the ansatz at the beginning of the iteration, one can notice that operators in positions 6-10 all had very similar gradients upon being selected. However, there are relevant differences in the energy change that they produced: and in fact, it was the operator with the lowest gradient among these (35) that caused the greatest change in energy. Moreover, selected with a gradient of more than double and holding position 5 in the ansatz, operator 25 produced only a fraction of this energy change. This further highlights the downsides of using the gradient as the selection method, and proceeding in disregard of the energy change that the selected operator actually produced.

At this point, we have an ansatz with only 6 operators that prepares a state with lower energy than any before. It is a significantly closer approximation to the ground state than the ansatz with 10 operators we had back at iteration 13, despite having a 40% decrease in the number of variational parameters and (approximately) circuit depth.

But in spite of the the ansatz looking favorable at this point, the way it was reached was certainly not the best. By iteration 14, three operators (25, 29 and 64) have been removed from the ansatz twice.

Removing an operator and adding it back represents an overhead in optimization costs (coefficients must be re-optimized upon removal and upon re-addition), so it is certainly beyond optimal to have them removed twice.

There could be no greater sign that removed operators are being allowed back into the ansatz too soon than seeing them removed a second time. In conclusion, it seems that the strategy of blocking them until a new operator produces a change in energy similar to the one they had produced is not the most appropriate. After examining this example run, this became exceedingly evident.

This is not quite so unpredictable. After we remove an operator, we know that it didn't perform as well as could be expected given its gradient. In a first analysis, blocking it until there is a similar energy change seems adequate. However, unblocking operators based on the energy change produced by a single new operator is dangerous: there is the possibility that the new operator also impacted the energy significantly less than another one with the same gradient could. What is more, once unblocked, the operator will go back to competing with the remaining operators in the pool on the same standing. The selection method is still the gradient, so it is likely that it will be picked very soon after it's unblocked (as did happen multiple times in the example). The information that an operator had little impact in energy is being underused: we are only taking it into account to remove and block the operator. Once it is unblocked, we allow it to be selected again by its gradient, even though we now know that it is an inadequate indicator for that operator especially.

In the previous exposition, we can see that operator 25, the first being removed (iteration 4, figure 43), is allowed back into the ansatz by the addition of operator 64 (iteration 7, figure 44). However, operator 64 is later removed, signaling that it is not a good reference for unblocking others. What is more, this last operator is itself unblocked by the addition of 59 (iteration 9, figure 46), later removed (iteration 14, figure 49). By iteration 14 all the aforementioned operators are removed once again, because an operator with a smaller gradient turns out to be a far better choice than any of them. However, this operator did not have a chance to be added to the ansatz before all the removed operators were added again, because they did still have higher gradients.

This is particularly bad considering how close the gradient of this operator (35) was to that of the bad performing operators. For example, at iteration 11 (figure 47), previously removed operator 64 is added back to the ansatz. In this iteration, the gradient of operator 64 is only 107% that of operator 35, the above average performing one. And even though we have no way of knowing that this operator will perform so well, we do know that operator 64 did not perform as well as it could be expected from its gradient. As such, it would definitely be interesting to give an opportunity to others with close, but slightly smaller gradients.

The goal of removing terms was precisely to explore operators with lower gradients, and this goal is not being met: blocking removed operators was proved to be an inadequate option. As soon as a bad performing operator is added to the ansatz, all of the previously blocked ones are unblocked and likely to be selected.

After this discussion, the solution that seems the most adequate is to somehow change the selection

criterion so as to incorporate information regarding previously removed operators: ideally, an operator that has been removed would be less likely to be selected than others. Something as binary as blocking and unblocking operators doesn't seem to take advantage of the information that we have to the full extent.

6.1.2.3 Performance Penalty

The best way to incorporate extra factors into the selection criterion without drastically changing it would be to impose a penalty on the gradient of the removed operators. The selected operator would be the one having the highest 'effective' gradient, with this quantity reflecting the below average performance of those operators that had been previously removed. Instead of just having the operator blocked and later unblocked, we'd have its effective gradient lower than the actual gradient, in accordance with how poorly it performed. Evidently, it would be interesting to have some quantification of just how bad the energy change was, so as to apply a proportional penalty.

The purpose would be signaling to the algorithm that, even if a removed operator has a high gradient, we know that it doesn't change the energy as much as one could expect. As such, we want to attempt to add other operators to the ansatz, even if they have lower gradients, in hope of finding some that have a greater impact on the energy. The goal is using the knowledge that we now have (regarding the removed operators) to make a more informed selection: this penalty would be a way of correcting, in part, the flaws of the gradient as an indicator of the potential energy change. In those cases in which the gradient was a particularly poor indicator, to the point we removed the operator, the effective gradient would be lowered accordingly.

The aforementioned performance ratio seems interesting for this purpose. After we select a given operator and re-optimize the ansatz, we can calculate the associated performance ratio from its gradient and the resulting energy change (as compared to the previous iteration) from formula 6.1.

The energy change is what we actually want from an operator: that it brings us closer to the ground energy. As the selection method, the gradient quantifies in a way how much we expect it to change the energy: we select the operator with the highest gradient hoping that it is also the one with the greatest impact on the energy. So intuitively, this ratio can be interpreted as something like

$$\text{performance ratio} = \frac{\text{actual performance}}{\text{expected performance}}.$$

While the performance ratio seems useful and has the nice property of being dimensionless, it doesn't allow us in itself to establish a penalty: we don't know what is a good or bad performance ratio.

Once again, the easiest solution seems to be assessing performance relative to that of the remaining operators. If we take the average over all known performance ratios (meaning, over the ratios of all of the operators ever added to the ansatz), we get a very interesting quantity: the average (absolute) energy change per unit gradient. At last, we know what energy change it is reasonable to expect from an operator with a given gradient.

One can then define a dimensionless penalty for removed operators as

$$\text{penalty} = \frac{\text{performance ratio}}{\text{standard performance ratio}}, \quad (6.2)$$

where the standard performance ratio is our reference - it can just be the average taken over all known performance ratios. This metric tells us just how below average the performance of the operator was. If the energy change per unit gradient of the removed operator was half of the average, the effective gradient will be half of the actual gradient. This 'sanction' will prevent it from being selected back into the ansatz as soon.

As long as there are other operators in the pool that, considering their gradient, will have a greater impact on the energy *assuming that their energy change per unit gradient is average*, the removed operator will not be selected again. Only when, from the gradients of the other operators in the pool and under the same assumption, we don't expect that any will impact the energy more, will the removed one be selected.

As it was exposed before in figure 41, the performance ratio has a tendency to decrease along the iterations; that is why it wasn't a suitable metric to decide which operators to remove. Taking the average will already soften this effect: the value will be constantly updated as more performance ratios are known, and as such will follow their tendency to decrease. To further avoid over-penalizing the removed operators, it will be interesting to use the moving average, rather than the average, as the standard performance ratio. This will prevent data from too old iterations to affect the calculation of the penalty.

The window over which we calculate the average shouldn't fall in the other extreme and be too small. If we only consider the ratio of two or three of the last added operators, and one of them performs particularly bad, it will have a significant impact on the average - thus causing the penalty to be too soft, and the removed operator to be added back too soon. It is better to be cautious, because the resulting optimization overhead can significantly worsen how fast the algorithm is to converge, by wasting iterations adding and removing the same operator multiple times. In the following examples, a window of 10 was used.

The final version of the procedure of operator removal can now be stated. The basic algorithm is the same as the original [ADAPT-VQE](#), except at every iteration j , when a new operator is added, the following must be done:

- Calculate the performance ratio of the new operator (equation 6.1). Store this value (associated with the operator) and update the standard performance ratio.
- Go through the previous iterations and compare the energy changes produced in them with that produced in this one. If an operator added at iteration i caused an energy change ΔE_i such that $\Delta E_i > r\Delta E_j$, attempt to remove this operator and re-optimize the coefficients of the ansatz without it. Calculate the energy change produced by this action, $\Delta E'_i$.
- Compare the increase in energy caused by removing the operator with the decrease that had been caused by adding it. If $-\Delta E'_i > t\Delta E_i$, remove the operator. In the following iterations, multiply

its gradient by a penalty calculated as in equation 6.2. The penalty should be recalculated each iteration, to use the most recent value of the standard performance ratio.

As it was mentioned before, one should choose $0 < r < 1$ and $t > 1$. The examples to follow used $r = 0.5$ and $t = 1.5$. The 10-iteration moving average of the performance ratio was used as the standard.

In order to compare this approach with the previous one (blocking / unblocking operators), the execution of the algorithm will be exemplified for the same test case as before, albeit more succinctly.

	Ansatz			Penalized Operators
Position	0	1	2	
Index	193	187	25	
Gradient ($\times 10^{-3}$)	173	71	30	
ΔE ($\times 10^{-5}$)	-1311	-314	-9	
Performance Ratio ($\times 10^{-3}$)	76	44	3	\emptyset

Standard Performance Ratio ($\times 10^{-3}$): 41
(10-iteration moving average)

Figure 50: Relevant data structures at iteration 3 of the ADAPT-VQE algorithm, with an additional term removal feature based on a performance penalty. The molecule in consideration is *LiH*, and the pool the SGSD pool.

At iteration 3 (figure 50), the ansatz is exactly the same as previously (figure 42). Until this point, as was pointed out before, no operators have met the removal criterion.

The only difference is that now we're keeping track of the performance ratio of the operators, easily calculated from their gradient upon being selected and the energy change that they caused (equation 6.1). This ratio has decreased from iteration to iteration up to now. It was already known that the produced energy change and the gradient of the selected operator were strictly decreasing until iteration 3, but it is also interesting to note that the absolute value of the change in energy per unit gradient is also decreasing.

At this point, the standard performance ratio sits at 41×10^{-3} . Evidently, it has also been on a monotonous decrease.

As before, the monotony is broken at iteration 4 by an operator that impacts the energy significantly more than its predecessor. We attempt to remove it and do so successfully, as the impact in the energy is still low after the addition of the last operator.

But now we don't 'block' operator 25: we just mark its bad performance, and keep a note that it should be penalized in the following iterations.

The developments of iterations 5 to 13, with a focus on the penalized operator, are represented in figure 52.

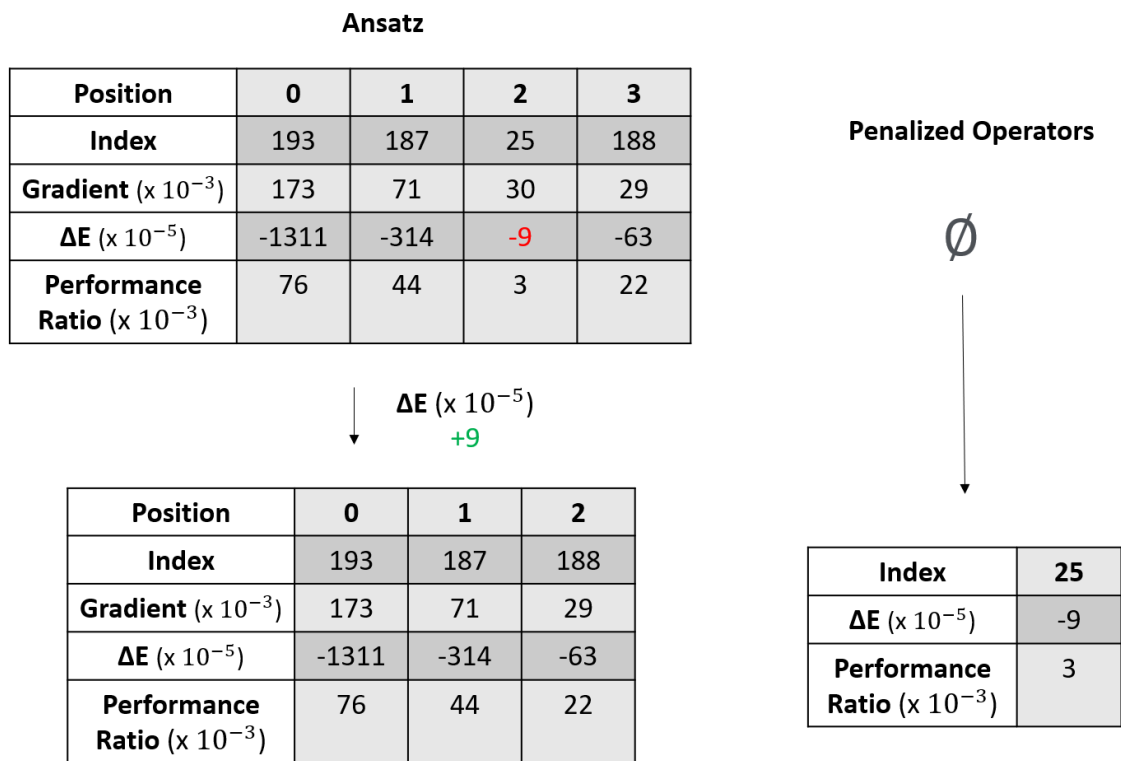


Figure 51: Relevant data structures at iteration 4.

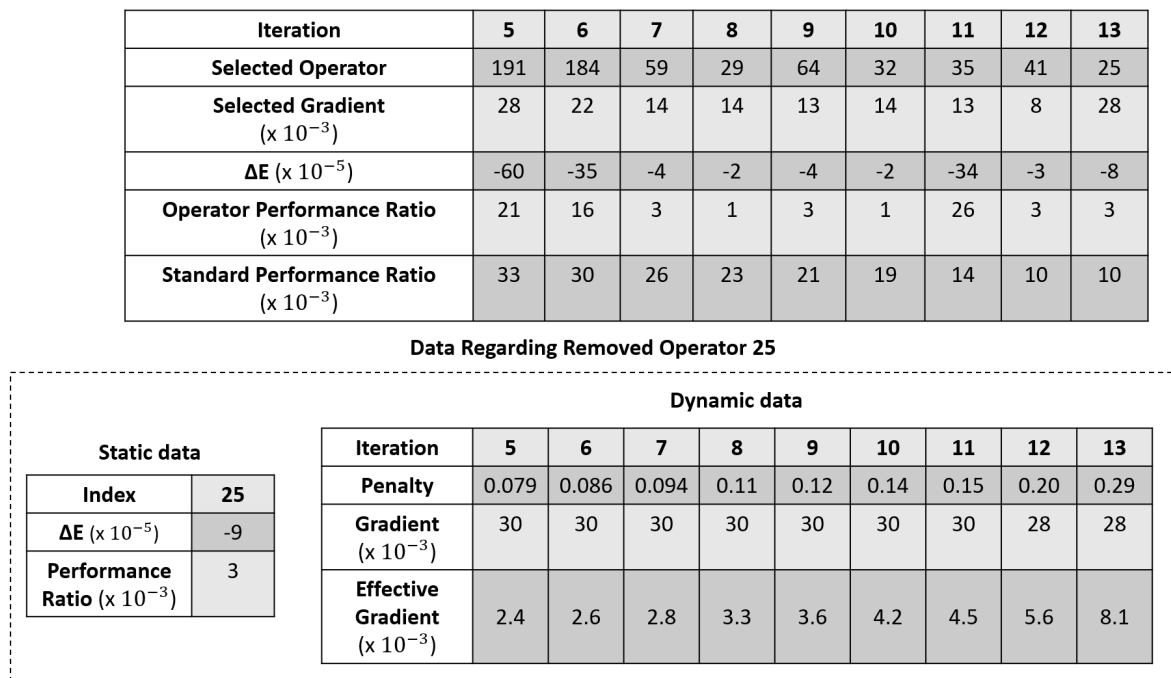


Figure 52: Evolution of the relevant data, iteration 5 through 13. On top, we have data on the evolution of the algorithm and on the selected operators. On the bottom, we have data specifically regarding operator 25, that was removed at iteration 4.

The relevant static data regarding this operator are simply the energy change it produced and its performance ratio. The energy change is only presented for its relevance here, in the context of analysing and comparing the changes produced by different operators; it is not directly used by the algorithm. Only the performance ratio is used in calculating the penalty.

The rest of the data concerning operator 25 changes from iteration to iteration. In each of them, equation 6.2 is used to calculate the penalty to apply to its gradient. The gradient is then multiplied by this penalty to obtain the *effective gradient*. This is, as should, lower than the original value: we know that this operator produced an energy change per unit gradient under average, and want to use that information to make a more informed selection. As such, we lower its gradient by exactly how under average the energy change by unit gradient was.

As can be seen from the figure, the gradient itself might suffer slight changes, but these changes are not very relevant. What changes significantly is the penalty, as the standard performance ratio decreases with the addition of new operators. In the iterations that immediately follow its removal, this penalty greatly damages the effective gradient of the operator, causing it to be as low as 2.4 (7.9% of the actual gradient) in iteration 5. This is because all operators added so far impacted the energy much more in proportion to the gradient than operator 25. As such, it is interesting to attempt adding other operators rather than this one into the ansatz, despite them having lower gradients: they might produce a greater energy change per unit gradient. This is, as discussed, the purpose of the penalty. Even though we didn't block operator 25, the lowered effective gradient postponed it being selected to the ansatz a second time.

At iteration 5, the selected operator is 191, with a gradient of 28×10^{-3} , 93% of the gradient of the removed operator. Adding this operator with a slightly lower gradient proved profitable: it produced an energy change almost seven times bigger than that of the sanctioned operator. Operator 184 follows in iteration 6, with a gradient of 73% that of the removed operator. Once again, it outperforms it, impacting the energy roughly four times more.

From iteration 7 to 10, that changes: the selected operators not only have lower gradients, but also lower impacts on the energy than the removed operator. Accordingly, their performance ratio is below average. One can see that, as a consequence, the standard performance ratio (as the 10-iteration moving average) is significantly decreasing: it goes from 30 at the beginning of iteration 7 to 19 at the end of iteration 10.

The decrease softens the penalty on operator 25, but it's not enough to have it selected again yet. Even though there are 4 iterations in a row with less beneficial operators, the 10-iteration moving average of the performance ratio is still quite over that of the removed operator (because of the previous operators). That causes the algorithm to still 'have hope' that a better operator with a lower gradient will be found.

That is actually the case. By iteration 11, operator 35, with a gradient of 13 (only 43% that of operator 25), is selected and lowers the energy by -34×10^{-5} a.u. - an almost fourfold improvement as compared to operator 25.

Obviously, there was no way to know beforehand which operators were going to be a better addition to the ansatz: all operators added in iteration 7 through 11 had very similar gradients, and before adding

them, that is all the information we had about them. However, there was a great disparity in the energy changes they produced, as we had seen before.

We now see that adding poorly performing operators 59, 29, 64 and 32 (iterations 7-10) was a ‘necessary evil’: there was no other way to find operator 35 in iteration 11 than to test them all.

This time, it was in fact avoided to add operator 25 to the ansatz too soon. We managed to find better operators not only in the iterations immediately after (5 and 6) but also in iteration 11, still before adding operator 25 once again. This didn’t happen with the blocking / unblocking heuristic, in which the addition of a single operator with a poor impact on the energy was enough to allow adding the removed operator back to the ansatz. With the performance penalty, it was indeed accomplished that more operators were given a chance to before the re-addition of 25.

Finally, at iteration 13, the removed operator 25 is added back to the ansatz. This is allowed by the eased penalty (from the decrease of the performance ratio), and the fact that the gradients of the not-yet-selected operators in the pool are lower now.

It should be noted that, having the purpose of illustrating the *addition* of operators, figure 52 glosses over the fact that some of the operators added in iterations 5 to 13 are later removed from the ansatz. By the time operator 25 is added, operators 29, 32, 59 and 64 have been removed and have a sanction associated with their gradient. Because their gradient and performance ratio are lower than those of operator 25, their effective gradient is significantly lower. Even if they all belong to the set of penalized operators, they are penalized differently according to how they performed, which is a way of putting to use all the information we have.

Once operator 35 is added and outperforms many of its predecessors (iteration 11), the ansatz and the set of penalized operators is as represented in figure 53.

Interestingly, this is a very similar ansatz to that obtained in the end of iteration 14 using the blocking / unblocking heuristic. However, it took us three less iterations to get to it (11 now against 14 then) to get there, because we didn’t waste time removing and adding operators multiple times. Previously, three iterations had been spent adding back operators that had been removed recently and would be removed again soon. What is more, there is an extra optimization per removed operator, meaning that we saved a total of six optimizations by improving how and when the operators are added back to the ansatz.

There is additional interesting information in figure 53. For example, it can be seen that the average performance ratio of the operators in the ansatz is 179% the average taken over *all* the operators we have information about (in the ansatz or penalized). In particular, it is surprising how under average the performance ratio of the penalized operators is: each of them has a ratio *at least* five times lower than any of the operators in the ansatz. These operators change the energy very little in proportion to their gradient as compared to the other ones. Because their gradient is misleadingly high, the gradient as a selection method recognizes them as good options; but for our purpose (lowering the energy) they are not. Actually, *all* five operators that were removed and penalized produce together an energy change of -21×10^{-5} a.u.. This is only 62% of the energy change of the *worst* operator in the ansatz *alone*. Even though they represent about five times the cost in terms of the number of variational parameters and circuit depth,

Ansatz						
Position	0	1	2	3	4	5
Index	193	187	188	191	184	35
Gradient ($\times 10^{-3}$)	173	71	29	28	22	13
ΔE ($\times 10^{-5}$)	-1311	-314	-63	-60	-35	-34
Performance Ratio ($\times 10^{-3}$)	76	44	22	21	16	27

Ansatz Performance Ratio ($\times 10^{-3}$): 34

Average Performance Ratio ($\times 10^{-3}$): 19

Standard Performance Ratio ($\times 10^{-3}$): 14
(10-iteration moving average)

Penalized Operators

Index	25	29	32	59	64
ΔE ($\times 10^{-5}$)	-9	-2	-2	-4	-4
Performance Ratio ($\times 10^{-3}$)	3	1	1	3	3

Figure 53: Relevant data structures at the end of iteration 11.

they are together quite less important for lowering the energy than any single operator still in the ansatz by iteration 11.

All of this is a good indicator that we have picked interesting heuristics for removing operators and handling them afterwards. However, the exemplification of a few iterations is not enough to say that the procedure is sound.

It should be asked whether the advantage of removing terms is limited to certain molecules and/or operator pools.

The simple Qubit Pool that results from using all tensor products of Pauli operators occurring in the fermionic excitations (with or without the anti-commutation string) is redundant, and therefore not a suitable choice here. We would often find that we removed an operator only to add one with almost exactly the same effect on the state and energy. Regardless, this pool is also excessive and was proven to be significantly larger than necessary for convergence. As for all the other pools, there seems to be no reason to believe that they would not allow for removing terms.

With a reasonable choice of pool, there are still two possibilities for the procedure not being advantageous.

The most obvious one is the energy change in any given iteration not being significantly higher than that of any of the previous iterations. In case the removal condition is never verified, never will it be attempted to remove an operator. In this case, there is no benefit to the version of [ADAPT-VQE](#) with term removal, but there is also no extra cost in optimizations or measurements. If this happens, the algorithm perform exactly the same with and without term removal. However, it seems highly unlikely: even in the small example that was used before multiple terms were suitable to be removed. In a minimal basis, the representation of

LiH has only 12 spin orbitals; bigger molecules will have more spin orbitals, and as a consequence more operators in the pools to choose from. It is also likely that there will be more operators with comparable gradients, in which case the selection using the gradient seems even less fitting. Additionally, only 10 iterations were considered, which is a fairly small number; in larger molecules, more iterations will be necessary to reach a low error, comprising more possibilities for removing terms.

There is, however, another possibility: having the condition for operator removal verified, but upon removal, seeing the energy increase to the point that it doesn't compensate to go through with the removal. It was a part of the method presented in this chapter that, upon removing an operator, it should be verified that the removal did not produce an increase in energy over 150% the decrease that the operator had caused in the first place - thus avoiding to remove an operator that did not impact the energy much, but was an important part of the variational wave function once other operators were added.

Up to this point, that was never the case. In the example presented before, with lithium hydride, every time the removal condition was met the removal came to fruition. Each time, removing the operator produced a very similar change in energy to that observed when adding it. However, this shouldn't be taken for granted, as there is no reason to believe it should always be the case.

If we attempt to remove an operator and optimize the parameters of the wave function without it, only to conclude that it should not be removed on account of the energy having increased by too much, then we have absolutely no gain to compensate for the extra optimization step. Of course, this is a necessary evil because we can't know in advance what will happen. In this scenario, operator removal brings only extra cost to the algorithm, but this extra cost does not increase maximum requirements in what comes to circuit depth or number of variational parameters.

In order to better assess the performance of [ADAPT-VQE](#) with term removal, the next few subsections aim to present graphical results for four different molecules, and different pools.

6.1.3 Application to LiH

The first example will be the one used in the previous subsection to illustrate the result of the different heuristics: LiH at an interatomic distance of 1.45\AA , with the [SGSD](#) pool.

In figure [54](#), we can compare the performance of the original [ADAPT-VQE](#) with the version allowing term removal according to the method outlined before.

It is difficult to compare the two approaches, for in each iteration they may differ both in the number of operators in the ansatz and the number of optimizations performed so far. Since currently the bottleneck of these hybrid algorithms concerns the quantum computers, it was privileged to contrast clearly the maximum number of operators in the ansatz, which is closely related to the maximum circuit depth. The number of operators also corresponds to the number of variational parameters, and thus sets the dimensionality of the optimization.

As such, the plot contemplates 10 iterations of the original [ADAPT-VQE](#) procedure, and as many iterations of the version with term removal as occur *before an iteration finishes with 11 operators*. This




Figure 54: Plot comparing the original [ADAPT-VQE](#) and the version with term removal, for LiH at an interatomic distance of 1.45\AA and using the [SGSD](#) pool. The error in the energy (top) and the number of operators in the ansatz (bottom) are plotted against the iteration number. The shaded blue area on the upper plot marks the error values within chemical accuracy (less than 1kcal/mol).

means that none of the iterations included for the term removing version ended with an ansatz having more elements than the ansatz of the original [ADAPT-VQE](#) protocol did by iteration 10. This can be confirmed in the bottom plot, that represents the number of operators in the ansatz.

Evidently, the original [ADAPT-VQE](#) outperforms the version with term removal in any given iteration: the latter ‘wastes’ iterations removing terms. However, the iteration count is not in itself an interesting cost metric.


In what concerns maximum circuit depth, we can expect a similar cost from the 10 iterations of the original protocol as from the 18 of the version with term removal. Of course, there might be small variations arising from the difference between the circuits that implement the different excitations. What is more, the particular architecture of the quantum computer will have an influence. Qubit connectivity will play a big role: operators that represent excitations between spin-orbitals corresponding to qubits that are directly connected will naturally be implemented by shallower circuits. However, there is no reason to believe that those variations would benefit either procedure. As such, assuming an average circuit depth per operator seems reasonable.

The maximum number of variational parameters is also matched between the 10 iterations of the original [ADAPT-VQE](#) and the 18 iterations of the version with term removal. This represents a similar maximum level of effort required of the classical optimizer, which in turn is related to the number of calls to the quantum computer (an optimization with more variational parameters will in principle require more energy evaluations).

It is interesting to notice how removing terms helps reaching a higher accuracy than that obtained by iteration 10 of the original [ADAPT-VQE](#), even though the maximum ansatz size and number of variational parameters are matched. Numerically, we find that by the time the ansatz has 10 operators, the original [ADAPT-VQE](#) has reached a state with an energy approximately 3.9×10^{-4} a.u. over the ground state energy. For contrast, the version with term removal has an error of only 6.1×10^{-5} a.u. - only 16%, all the while keeping the maximum requirements on the quantum computer and the classical optimizer. The extra cost comes in the number of measurements and optimizations, which does not affect how viable the implementation is.

In figure 55, the energy and error plotted against the number of operators are shown. In the case of the term removing version, the ansatz in consideration for each operator number is the last occurring ansatz of that size. As before, to allow a fair comparison, iterations were only considered up until (and excluding) the first one ending with 11 operators.

These graphs do not have present the extra cost in measurements and optimizations brought on by term removal; however, it is interesting to note how it is possible to reach a lower error with more compact ansätze. The version of [ADAPT-VQE](#) with term removal manages to find better ansätze for any number of operators larger than two. What is more, as it was mentioned before, the maximum demands on the quantum computer and the classical optimizer in obtaining either curve are identical.



(a) Energy plotted against the number of operators in the ansatz.

(b) Error plotted against the number of operators in the ansatz.

Figure 55: Plots comparing the energy and energy error in the original [ADAPT-VQE](#) and the version with term removal, for LiH at an interatomic distance of 1.45\AA and using the [SGSD](#) pool. The shaded blue areas on the plots mark the values within chemical accuracy both for the energy (less than 1kcal/mol away from the [FCI](#) energy) and the error (less than 1kcal/mol).

6.1.4 Application to OH-

The second example is OH^- at an interatomic distance of 1.45\AA . At first, the same pool will be used: the [SGSD](#) pool. In the STO-3G minimal basis set, the pool for this molecule contains exactly the same excitations (in terms of operator representation) as the one for LiH. This is because they have the same number of spin orbitals and we are considering all generalized excitations, so that the particle number is irrelevant in creating the pool. Evidently, what differs is the spatial part of the orbitals each qubit represents; and because the Hamiltonian is different, the cost function is different altogether.

Using the same pool for a different molecule allows obtaining some more insight as to how general the success of the method for removing terms is. After this, a different pool will be tested, to further test generality of the procedure.

As before, the plot in figure [56](#) illustrates the evolution of the two versions of the algorithm (with and without term removal). Once more, both were allowed to grow the ansatz until an iteration finished with 10 operators in it.

Again, we see that by the time the size of the ansatz is matched between the two, the version with term removal has reached a lower error. It takes 15 iterations to get there, meaning that there were 5 operators removed in total. At this point, all operators which removal was attempted were in fact removed, because the condition to effectively remove a term was never violated. Removing an operator always resulted in about the same energy change as it had caused back in the iteration in which it was added.

The original [ADAPT-VQE](#) finishes with a 10-operator ansatz that prepares a state with an energy $9.7 \times$





Figure 56: Plot comparing the original [ADAPT-VQE](#) and the version with term removal, for OH⁻ at an interatomic distance of 1.45Å and using the [SGSD](#) pool. The error in the energy (top) and the number of operators in the ansatz (bottom) are plotted against the iteration number. The shaded blue area on the upper plot marks the error values within chemical accuracy (less than 1kcal/mol).

10^{-7} a.u. away from the FCI ground energy. As for the version with term removal, its final 10-operator ansatz is associated with an error of only 2.7×10^{-8} a.u. - only 2.8%.



(a) Energy plotted against the number of operators in the ansatz.

(b) Error plotted against the number of operators in the ansatz.

Figure 57: Plots comparing the energy and energy error in the original ADAPT-VQE and the version with term removal, for OH^- at an interatomic distance of 1.45 \AA and using the SGSD pool. The shaded blue areas on the plots mark the values within chemical accuracy both for the energy (less than 1 kcal/mol away from the FCI energy) and the error (less than 1 kcal/mol).

The energy plotted as a function of the number of operators in the ansatz is in figure 57a (the ansatz considered for each size was once again the last ansatz of that size). Interestingly, it seems like those iterations in which the energy does not change much as compared to the surrounding ones (iterations 4 and 5) are entirely avoidable. One could have imagined that the slowdown was inevitable, and that no operator would be capable of changing the energy more than the one chosen. However, this is not the case. Better choices of operators can help convergence be significantly faster as a function of the ansatz size.

In figure 57b, we can remark on how the version with term removal succeeds in finding a better ansatz for many numbers of operators. Interestingly, in contrast with what happened for LiH, one can see that the 7 and 8-operator ansätze are very comparable between the two versions of the algorithms. This is a sign that, for those numbers of operators, the simple gradient-based selection is nearly optimal.

In the following, results will be presented for the same molecule, but using the Eight Pool introduced in the previous chapter. The purpose is verifying whether the success of removing terms might be dependent on the pool; once the fermionic anti-commutation strings are removed and simpler operators are used, results might differ.

In figure 58, we can see that significant advantage can be achieved with this pool as well. By iteration 16 of ADAPT-VQE with term removal, we have a 10-operator ansatz that prepares a state with an energy





Figure 58: Plot comparing the original [ADAPT-VQE](#) and the version with term removal, for OH⁻ at an interatomic distance of 1.45Å and using the Eight Pool. The error in the energy (top) and the number of operators in the ansatz (bottom) are plotted against the iteration number. The shaded blue area on the upper plot marks the error values within chemical accuracy (less than 1kcal/mol).

7.4×10^{-6} a.u. away from the ground energy - only 4.6% of the error of the 10-operator ansatz that the original ADAPT-VQE algorithm has reached by iteration 10 (1.6×10^{-4} a.u.).



(a) Energy plotted against the number of operators in the ansatz.

(b) Error plotted against the number of operators in the ansatz.

Figure 59: Plots comparing the energy and energy error in the original ADAPT-VQE and the version with term removal, for OH⁻ at an interatomic distance of 1.45 Å and using the Eight Pool. The shaded blue areas on the plots mark the values within chemical accuracy both for the energy (less than 1 kcal/mol away from the FCI energy) and the error (less than 1 kcal/mol).

In figure 59, we can see that removing terms allows reaching a better ansatz for any number of operators after six. The difference in energy seems small on 59a, which is a consequence of the linear scale being unfit for showing the small but important changes that the energy suffers by the last iterations; once the error is plotted on a logarithmic scale (figure 59b), the difference becomes evident.

6.2 Conservative Ansatz Growth

An alternative approach to removing operators is to attempt to be more cautious in adding them to the ansatz in the first place. Once again, the idea is not to fully trust the gradient norm as an indicator of the possible impact on the energy. However, here the correction is done preemptively, rather than *after* the operators have been added to the ansatz.

Instead of blindly accepting the operator with the highest gradient into the ansatz in each iteration, this approach tests several operators before settling for one. At each iteration, the N operators with the highest gradients are independently added to the ansatz, and their coefficients are optimized. In the end, only the one that lowered the energy the most is effectively added. In the simulations to follow, $N = 5$ was used. Evidently, this implies a 5-fold increase in the number of optimizations that must be performed.

This method was employed in QEB-ADAPT-VQE [99].

6.2.1 Comparison with Operator Removal

6.2.1.1 Application to LiH

In order to compare these three approaches (original; term removal; conservative ansatz growth), simulations involving three different molecules were done.




Figure 60: Plot showing the minimum error reached for each ansatz size, for multiple ansatz growth strategies. The molecule in case is LiH at an interatomic distance of 1.45Å. The SGSD pool was used. In the points where the green curve, corresponding to conservative ansatz growth, is not visible, it lies directly under the term removal curve.

Figure 60 shows the familiar example of *LiH*, now complete with all three ansatz growing approaches. It is remarkable how similarly the term removal and conservative ansatz growth strategies perform. The latter is significantly behind at iteration 6, but recovers immediately after; further, even this difference would be effaced if one allowed a greater N in the strategy. In the other iterations, only one of the curves is visible because they overlap.

The figure does not show the extra overhead in the optimizations. As for the conservative ansatz growth strategy, it is straightforward that it implies a 5-fold increase in the total number of optimizations. The term removal version is not as easy to analyze, because the number of optimizations depends on the number of terms that were removed. In this case, term removal implied 16 extra optimizations to achieve a 10-operator ansatz. So the total number of optimizations for the original ADAPT-VQE, the version with term removal, and the version with conservative ansatz growth was respectively 10, 26, and 50. The maximum circuit depth, as well as the maximum number of variational parameters, are roughly the same for all. The term removal and conservative ansatz growth strategies achieve a greater accuracy for comparable circuit depth.

6.2.1.2 Application to OH-

Figure 61 presents the results for a different molecule (OH^-) and a different pool (the Eight Pool). The term removal curve coincides with the conservative ansatz growth curve. Here the original ADAPT-VQE, the version with term removal, and the version with conservative ansatz growth required respectively 10, 22, and 50 optimizations. The maximum circuit depth, as well as the maximum number of variational parameters, are again roughly the same for all. By the end, for the same size of ansatz, the term removal and conservative ansatz growth strategies reach an error of only 4.6% that of the original ADAPT-VQE algorithm.

6.2.1.3 Application to H4

Figure 62 presents the results for yet another molecule, H_4 . Interestingly, the results are very different here - the term removal version is not visible, but now because it lies directly under the *original* ADAPT-VQE curve. There are two terms that fit the criterion for being removed, but on the last step of the removal procedure, it is verified that removing them causes a significantly larger increase in the energy than the prior decrease. Remarkably, by the end, conservative ansatz growth results in an error of only 0.3% that of the other strategies (original ADAPT-VQE and term removal).

Here, the optimizations required were 10 for ADAPT-VQE, 12 for the term removal version, and 50 for the conservative ansatz growth. As much as attempting to remove terms does not improve accuracy, there is barely an optimization overhead. Additionally, the tendency in this graph it is not as clear as before in what comes to determining the success of the conservative ansatz growth strategy: through the first 10 iterations, there are many in which it actually corresponds to a greater error than the original ADAPT-VQE, despite the sizeable optimization overhead. There is a sharp drop in the error by iterations 9 and 10, seeming to hint that the procedure might become more advantageous as the algorithm evolves; but unfortunately, the simulations become increasingly expensive computationally as more terms are added. This behaviour was also not encountered in simulating any other molecule, which complicates the task of determining whether conservative growth is indeed advantageous when term removal fails.




Figure 61: Plot showing the minimum error reached for each ansatz size, for multiple ansatz growth strategies. The molecule in case is OH^- at an interatomic distance of 1.45\AA . The Eight Pool was used. The term removal and conservative ansatz growth strategies correspond to the same color because the curves coincide.




Figure 62: Plot showing the minimum error reached for each ansatz size, for multiple ansatz growth strategies. The molecule in case is H_4 at an interatomic distance of 1.45\AA . The original ADAPT-VQE and the version with term removal correspond to the same color because the curves coincide.

Table 1 presents a succinct comparison of the three ansatz growth strategies, for the case of the three systems analysed in this subsection. The final error is used as the figure of merit for characterizing performance, and the number of optimizations for characterizing measurement and optimization costs.

Method	System					
	LiH		OH^-		H_4	
	$ \Delta E $ ($\times 10^{-5}$)	# Opts	$ \Delta E $ ($\times 10^{-6}$)	# Opts	$ \Delta E $ ($\times 10^{-7}$)	# Opts
Original	39.0	10	160.2	10	2166.5	10
Term Removal	6.1	26	7.4	22	2166.5	13
Conservative	6.1	50	7.4	50	5.8	50

Table 1: Summary of the performance of the three methods for manipulating the ansatz (original ADAPT-VQE, term removal, conservative ansatz growth). The table presents the absolute error in the energy and the accumulated number of optimizations at the point a 10-operator ansatz is reached. The data concerns the three systems explored in this subsection (LiH , OH^- , and H_4) with the considered pools (SGSD, Eight Pool, and SGSD, respectively). The units of energy are Hartree, as always.

6.3 Discussion

Conservative ansatz growth seems like a solid procedure for creating compact ansätze. Unlike term removal, it has guarantees in what comes to maximum circuit size: at any level of accuracy, the maximum size of the ansatz is the current one. By being cautious with growing the ansatz, it is possible to burden the classical optimizer with the compensation of decreasing circuit depth. Further, growing the ansatz conservatively is viable even when the term removal procedure fails at the third step.

On the other hand, term removal has its advantages. By only removing terms *after* they are outperformed by their successors, the procedure avoids an overhead in optimizations at every single iteration. Extra optimizations only occur when there is a good chance that they are in fact benefiting for the ansatz, which results in a more modest overhead in total optimizations. Further, as the algorithm evolves, this saving becomes more and more rewarding; in the results presented before, even when limiting the *maximum* number of operators in the ansatz to 10, the term removal version matched or surpassed the accuracy of the version with conservative growth (with 5 optimizations per step) for OH^- and LiH , for ansätze with up to only 10 operators. This suggests that the conditions for removing a term are met sufficiently often that the size of the ansatz necessary for a given accuracy can be kept as modest with term removal as with conservative ansatz growth.

Further, it is expected that as the size of the system grows, testing a relevant portion of the operators in the pool per iteration will become harder and harder: larger molecules imply larger pools. In the cases presented here, testing the five operators with largest gradient in each iteration ($N = 5$) corresponded to testing up to 7.6% of the operators in the pool, which is a significant portion. If we increase N to always account for the same portion of the pool, we will have an added cost in optimizations that increases with the system size. In contrast, term removal is size-agnostic, because it will attempt to remove under-performing operators, rather than find out in advance which they are. Removing operators after they're outperformed (instead of testing several *a priori*) can keep the ansatz compact as other operators are explored, while saving on the extra optimizations.

What is more, because removed operators are penalized, they will not be added too soon after they're removed. In contrast, when the conservative ansatz growth procedure finds an operator that performs very poorly, it will likely still attempt to add it in the next iteration, instead of giving the opportunity to other operators with lower gradients. This approach does not attempt to leverage any information regarding the performance of the operators that was obtained in the previous iterations, which may prevent it from exploring more advantageous candidates and 'clog' several of the N attempts per iteration with operators that were already shown to perform badly.

Of course, removing terms is always more fallible than growing the ansatz conservatively. If, as happened with H_4 , term removal consistently fails on step 3, there is absolutely no added benefit to the procedure, only extra cost.

Conclusions and Future Work

The purpose of this dissertation was to explore several proposals of ansätze for the Variational Quantum Eigensolver in chemistry applications, with a focus on the noise-resilience that is paramount in improving short-term viability of the algorithm.

Ansätze can be split into three broad classes: predetermined and problem agnostic; predetermined and problem tailored; dynamic and problem-tailored. An ansatz from each class was implemented, and each was tested both in simulators and in real quantum computers made accessible on the cloud by IBMQ [42]. The effect of sampling noise, decoherence, SPAM errors and generic noise was analysed via simulations on different backends and creation of noise models. With this, we could compare the noise-resilience of different type of ansätze. In particular, the effect of noise on the predetermined UCCSD ansatz (considered a ‘gold standard’ in quantum chemistry) and the dynamic ansatz built in the ADAPT-VQE algorithm (a more recent proposal) was compared.

The benefits of ADAPT-VQE became evident in our simulations. For the H_2 molecule, a 1-operator qubit-ADAPT-VQE ansatz was only surpassed by UCCSD-VQE in a fully noise-free scenario, when even sampling noise was removed from the simulations. Such scenario would be no less than impossible to actually create, even if we had an ideal quantum processor available (which is far from reality). Two unavoidable factors, quantum projection noise and a finite number of shots, are enough to deteriorate the performance of UCCSD-VQE significantly. In the presence of generic noise, the difference between UCCSD-VQE and ADAPT-VQE became even more pronounced.

Problem-tailored (but predetermined) ansätze already help avoid some problems that arise in the context of problem-agnostic ansätze, *barren plateaus* being one example. This notwithstanding, *dynamic* ansätze are one step ahead: in addition to being problem-tailored, they are *system*-tailored, which brings along an array of benefits. Ansätze grown from scratch, like the one from ADAPT-VQE, can provide high-accuracy results with shallower circuits and less variational parameters than any predetermined alternative. Both of these factors increase the noise-resilience of VQE, resulting in it being an even more NISQ-Friendly algorithm when dynamic ansätze are used.

The choice of operator pool in ADAPT-VQE is the part of the algorithm that implies the biggest human involvement. In order to assess its importance, ADAPT-VQE was simulated resorting to a variety of pools in addition to those originally proposed. The constitution of the superposition state in terms of Slater determinants was analysed along with ease of convergence, and the relative importance of fermionic antisymmetry, preservation of quantities such as particle number and spin, and similarity with fermionic excitations for the success of the pool was explored.

This analysis allows us to conclude that there is a delicate balance between choosing operators that

respect fermionic symmetries, and choosing operators that can be implemented by shallow and convenient circuits. The former helps decrease the number of iterations and variational parameters, while the latter can reduce the circuit depth; but the two are tightly interwoven, and the problem at hand is to find the optimal compromise. Regardless, there seem to be aspects in which being faithful to the proper representation of fermionic operators brings no advantage, and we find that relinquishing them lowers costs free of negative repercussions. This is the case of anticommutation: relinquishing such principle eases, from linear to constant (on the size of the system), the scaling of the average circuit depth required to implement an operator, with no observable negative effect on the success of the algorithm.

Independently of the pool choice, there are also possibilities of improvement through additional manipulation of the ansatz. No selection criterion is unerring; but adjustments and corrections can aid in preventing the misuse of circuit depth on portions of the ansatz that bring little benefit. Two approaches were analysed to address this. Heuristics for removing operators from the ansatz along the evolution of the algorithm were developed. The main goal was keeping the ansatz shallow without hindering the efficiency of the algorithm by incurring in unnecessary measurements and optimizations, a consequence we showed could be brought about by an unpolished strategy. The second attempted strategy was simpler, consisting of testing the parallel addition of distinct operators to the ansatz per iteration, only to choose and add the best performing one.

The first strategy (term removal) was compared against the second one (conservative ansatz growth), with the original [ADAPT-VQE](#) as a reference. Tests encompassed simulations with several molecules and pools. Both strategies were shown to be able to bring a significant reduction to the circuit depth required to achieve a given accuracy. Term removal seems to have a more modest optimization overhead, since it only entails extra optimizations when it can sense that they may be advantageous, unlike conservative ansatz growth, which performs a constant number of optimizations per iteration irrespective of the evolution of the algorithm. Despite the reduced number of additional optimizations, removing terms proved to be often capable of maintaining the maximum ansatz depth as shallow as conservative ansatz growth. Further, the term removal procedure seems in principle to be unaffected by the size of the pool, unlike conservative ansatz growth. The portion of the pool we can test with a constant number of allowed optimizations per iteration will become smaller and smaller as the size of the system increases, leading to a degradation of the performance of the conservative ansatz growth strategy. However, this method is safer than term removal, as it does not rely on the evolution of the algorithm or on the characteristics of the commutator between operators to keep the ansatz compact. In one of the tested molecules, H_4 , the term removal procedure was incapable of changing the evolution of the [ADAPT-VQE](#) algorithm for the better. For this molecule, the success of conservative ansatz growth was also less clear; in spite of this, for the maximum simulated ansatz size (twenty operators), it did achieve a lower error.

Variational quantum algorithms are a recent class, created specifically in the context of near-term quantum computing. [VQE](#) was the first algorithm of this class, proposed in 2014 [67]; [ADAPT-VQE](#) introduced the possibility of using dynamic ansätze for chemistry applications in 2019 [35], with further developments in 2020 [85] and September 2021 [75] bringing remarkable decreases in both the size of

the pool and the circuit depth per operator.

The dates of these proposals speak already of the novelty and intensity of the research around variational quantum algorithms in general and dynamic ansätze in particular. For the former case, noise-resilient optimization strategies, improved ways of measuring relevant observables, and strategies for avoiding barren plateaus are doubtlessly of great interest. For the latter, and more relevant in the context of this dissertation, methods for developing and manipulating ansätze on the fly can certainly be expanded and improved upon.

The proposals of dynamic ansätze explored in this project were mainly directed at chemistry applications; it would be interesting to generalize the algorithm to other areas, such as nuclear and condensed matter physics, or even mathematics and machine learning. Since in [ADAPT-VQE](#) the specifics of the problem are tied to the content of the operator pool, the algorithm must be formulated differently for other problems. Further, recent results showed that the size of the pool can be significantly decreased by leveraging the characteristics of the specific molecule. We can speculate that, in general, leveraging problem-specific knowledge in the creation of the set of operators that serve as 'pieces' in the construction of the ansatz may allow for improving results and decreasing the measurement costs in the creation of adaptive wave functions.

Even within chemistry problems, [ADAPT-VQE](#) could be further explored. For example, pools could be chosen taking into account the native interactions and the connectivity of the device, with the purpose of obtaining a more hardware-friendly ansatz. Mappings other than Jordan-Wigner could be employed, changing the formulation of the problem, affecting the gradient measurements, the operator pool, and the qubit requirements, among other aspects.

And finally, as it was explored here, manipulating the ansatz beyond the steady addition of operators of [ADAPT-VQE](#) can improve the approximation to the solution while keeping the ansatz compact, by keeping in it only the most important operators. Removing terms or testing several per iteration are merely examples of what can be done. It would be interesting to test and compare these approaches on more molecules and pools, and for longer runs.

Explored or unexplored, there are many options and possibilities in designing ansätze for the variational quantum eigensolver - even more when we expand the scope to variational quantum algorithms in general. Finding increasingly more compact and convenient state preparation circuits is a fundamental piece in the pursuit of a chance at quantum advantage with variational quantum algorithms: if they are naturally [NISQ](#)-friendly, they can become even more so through a refinement of the ansatz. A proper choice can decrease the depth of the required circuits, facilitate the task of the classical optimizer, increase the chances of convergence, and reduce the number of calls to the quantum computer. All efforts done to improve the ansatz pave the way for the success of general variational approaches to quantum algorithms.

Bibliography

- [1] D. Aharonov and M. Ben-Or. “Fault-Tolerant Quantum Computation with Constant Error Rate”. In: *SIAM Journal on Computing* 38.4 (Jan. 2008), pp. 1207–1282. doi: [10.1137/s0097539799359385](https://doi.org/10.1137/s0097539799359385). url: <https://doi.org/10.1137%5C%2Fs0097539799359385>.
- [2] J. R. M. et al. “OpenFermion: the electronic structure package for quantum computers”. In: *Quantum Science and Technology* 5.3 (June 2020), p. 034014. doi: [10.1088/2058-9565/ab8ebc](https://doi.org/10.1088/2058-9565/ab8ebc). url: <https://doi.org/10.1088%5C%2F2058-9565%5C%2Fab8ebc>.
- [3] M. S. A. et al. *Qiskit: An Open-source Framework for Quantum Computing*. <https://qiskit.org/>. 2021. doi: [10.5281/zenodo.2573505](https://doi.org/10.5281/zenodo.2573505).
- [4] N. J. M. et al. *Adapt-VQE*. <https://github.com/mayhallgroup/adapt-vqe>. 2019.
- [5] Q. S. et al. “PySCF: the Python-based simulations of chemistry framework”. In: *WIREs Computational Molecular Science* 8.1 (2018), e1340. doi: <https://doi.org/10.1002/wcms.1340>. eprint: <https://wires.onlinelibrary.wiley.com/doi/pdf/10.1002/wcms.1340>. url: <https://wires.onlinelibrary.wiley.com/doi/abs/10.1002/wcms.1340>.
- [6] V. et al. “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python”. In: *Nature Methods* 17 (2020), pp. 261–272. doi: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2).
- [7] P. Aliferis, D. Gottesman, and J. Preskill. “Quantum accuracy threshold for concatenated distance-3 code”. In: *Quantum Information and Computation* 6.2 (Mar. 2006), pp. 97–165. doi: [10.26421/qic6.2-1](https://doi.org/10.26421/qic6.2-1). url: <https://doi.org/10.26421%5C%2Fqic6.2-1>.
- [8] M. R. Alves. *VQE (GitHub repository for this project)*. <https://github.com/MafaldaRA/VQE>. 2021.
- [9] A. Anand, M. Degroote, and A. Aspuru-Guzik. “Natural evolutionary strategies for variational quantum computation”. In: *Machine Learning: Science and Technology* 2.4 (July 2021), p. 045012. issn: 2632-2153. doi: [10.1088/2632-2153/abf3ac](https://doi.org/10.1088/2632-2153/abf3ac). url: <http://dx.doi.org/10.1088/2632-2153/abf3ac>.
- [10] A. Arrasmith et al. *Effect of barren plateaus on gradient-free optimization*. 2020. arXiv: [2011.12245](https://arxiv.org/abs/2011.12245) [quant-ph].
- [11] R. J. Bartlett and M. Musiał. “Coupled-cluster theory in quantum chemistry”. In: *Rev. Mod. Phys.* 79 (1 Feb. 2007), pp. 291–352. doi: [10.1103/RevModPhys.79.291](https://doi.org/10.1103/RevModPhys.79.291). url: <https://link.aps.org/doi/10.1103/RevModPhys.79.291>.

- [12] P. Benioff. “The computer as a physical system: A microscopic quantum mechanical Hamiltonian model of computers as represented by Turing machines”. In: *Journal of Statistical Physics* 22.5 (May 1980), pp. 563–591. doi: [10.1007/bf01011339](https://doi.org/10.1007/bf01011339). url: <https://doi.org/10.1007%2Fbf01011339>.
- [13] E. van den Berg, Z. K. Mineev, and K. Temme. *Model-free readout-error mitigation for quantum expectation values*. 2021. arXiv: [2012.09738](https://arxiv.org/abs/2012.09738) [quant-ph].
- [14] M. Born and R. Oppenheimer. “Zur Quantentheorie der Molekeln”. In: *Annalen der Physik* 389.20 (1927), pp. 457–484. doi: [10.1002/andp.19273892002](https://doi.org/10.1002/andp.19273892002). url: <https://doi.org/10.1002%5C%2Fandp.19273892002>.
- [15] Z. Cai, X. Xu, and S. C. Benjamin. “Mitigating coherent noise using Pauli conjugation”. In: *npj Quantum Information* 6.1 (Feb. 2020). doi: [10.1038/s41534-019-0233-0](https://doi.org/10.1038/s41534-019-0233-0). url: <https://doi.org/10.1038%5C%2Fs41534-019-0233-0>.
- [16] M. Cerezo et al. “Cost function dependent barren plateaus in shallow parametrized quantum circuits”. In: *Nature Communications* 12.1 (Mar. 2021). issn: 2041-1723. doi: [10.1038/s41467-021-21728-w](https://doi.org/10.1038/s41467-021-21728-w). url: <http://dx.doi.org/10.1038/s41467-021-21728-w>.
- [17] M. Cerezo et al. *Variational Quantum Algorithms*. 2020. arXiv: [2012.09265](https://arxiv.org/abs/2012.09265) [quant-ph].
- [18] I. L. Chuang, N. Gershenfeld, and M. Kubinec. “Experimental Implementation of Fast Quantum Searching”. In: *Physical Review Letters* 80.15 (Apr. 1998), pp. 3408–3411. doi: [10.1103/physrevlett.80.3408](https://doi.org/10.1103/physrevlett.80.3408). url: <https://doi.org/10.1103%2Fphysrevlett.80.3408>.
- [19] O. Crawford et al. “Efficient quantum measurement of Pauli operators in the presence of finite sampling error”. In: *Quantum* 5 (Jan. 2021), p. 385. doi: [10.22331/q-2021-01-20-385](https://doi.org/10.22331/q-2021-01-20-385). url: <https://doi.org/10.22331%5C%2Fq-2021-01-20-385>.
- [20] T. D. Crawford et al. “PSI3: An open-source Ab Initio electronic structure package”. In: *Journal of Computational Chemistry* 28.9 (Apr. 2007), pp. 1610–1616. doi: [10.1002/jcc.20573](https://doi.org/10.1002/jcc.20573). url: <https://doi.org/10.1002%5C%2Fjcc.20573>.
- [21] D. Deutsch and R. Jozsa. “Rapid solution of problems by quantum computation”. In: *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences* 439.1907 (1992), pp. 553–558.
- [22] C. Developers. *Cirq*. Version v0.12.0. See full list of authors on Github: <https://github.com/quantumlib/Cirq/graphs/contributors>. Aug. 2021. doi: [10.5281/zenodo.5182845](https://doi.org/10.5281/zenodo.5182845). url: <https://doi.org/10.5281/zenodo.5182845>.
- [23] D. Dieks. “Communication by EPR devices”. In: *Physics Letters A* 92.6 (Nov. 1982), pp. 271–272. doi: [10.1016/0375-9601\(82\)90084-6](https://doi.org/10.1016/0375-9601(82)90084-6). url: <https://doi.org/10.1016%2F0375-9601%2882%2990084-6>.

-
- [24] E. Farhi, J. Goldstone, and S. Gutmann. *A Quantum Approximate Optimization Algorithm*. 2014. arXiv: [1411.4028](https://arxiv.org/abs/1411.4028) [quant-ph].
- [25] R. P. Feynman. “Simulating physics with computers”. In: *International journal of theoretical physics* 21.6 (1982), pp. 467–488.
- [26] B. T. Gard et al. “Efficient symmetry-preserving state preparation circuits for the variational quantum eigensolver algorithm”. In: *npj Quantum Information* 6.1 (Jan. 2020). issn: 2056-6387. doi: [10.1038/s41534-019-0240-1](https://doi.org/10.1038/s41534-019-0240-1). url: <http://dx.doi.org/10.1038/s41534-019-0240-1>.
- [27] M. R. Geller and M. Sun. *Efficient correction of multiqubit measurement errors*. 2020. arXiv: [2001.09980](https://arxiv.org/abs/2001.09980) [quant-ph].
- [28] P. Gokhale et al. *Minimizing State Preparations in Variational Quantum Eigensolver by Partitioning into Commuting Families*. 2019. arXiv: [1907.13623](https://arxiv.org/abs/1907.13623) [quant-ph].
- [29] P. Gokhale et al. “Optimization of Simultaneous Measurement for Variational Quantum Eigensolver Applications”. In: *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)*. IEEE, Oct. 2020. doi: [10.1109/qce49297.2020.00054](https://doi.org/10.1109/qce49297.2020.00054). url: <https://doi.org/10.1109/qce49297.2020.00054>.
- [30] *Google Colab, 2021*. <https://colab.research.google.com/>. 2021.
- [31] D. Gottesman. *An Introduction to Quantum Error Correction and Fault-Tolerant Quantum Computation*. 2009. arXiv: [0904.2557](https://arxiv.org/abs/0904.2557) [quant-ph].
- [32] E. Grant et al. “An initialization strategy for addressing barren plateaus in parametrized quantum circuits”. In: *Quantum* 3 (Dec. 2019), p. 214. issn: 2521-327X. doi: [10.22331/q-2019-12-09-214](https://doi.org/10.22331/q-2019-12-09-214). url: <http://dx.doi.org/10.22331/q-2019-12-09-214>.
- [33] D. Greenbaum and Z. Dutton. “Modeling coherent errors in quantum error correction”. In: *Quantum Science and Technology* 3.1 (Dec. 2017), p. 015007. doi: [10.1088/2058-9565/aa9a06](https://doi.org/10.1088/2058-9565/aa9a06). url: <https://doi.org/10.1088/2058-9565/aa9a06>.
- [34] D. J. Griffiths. *Introduction to quantum mechanics*. Vol. 2. Prentice Hall New Jersey, 1995.
- [35] H. R. Grimsley et al. “An adaptive variational algorithm for exact molecular simulations on a quantum computer”. In: *Nature Communications* 10.1 (July 2019). issn: 2041-1723. doi: [10.1038/s41467-019-10988-2](https://doi.org/10.1038/s41467-019-10988-2). url: <http://dx.doi.org/10.1038/s41467-019-10988-2>.
- [36] L. K. Grover. “A fast quantum mechanical algorithm for database search”. In: *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing - STOC '96*. ACM Press, 1996. doi: [10.1145/237814.237866](https://doi.org/10.1145/237814.237866). url: <https://doi.org/10.1145/237814.237866>.
- [37] S. Haroche and J. Raimond. “Quantum Computing: Dream or Nightmare?” In: *Physics Today* 49 (1996), pp. 51–52.

- [38] G. Harsha, T. Shiozaki, and G. E. Scuseria. "On the difference between variational and unitary coupled cluster theories". In: *The Journal of Chemical Physics* 148.4 (Jan. 2018), p. 044107. issn: 1089-7690. doi: [10.1063/1.5011033](https://doi.org/10.1063/1.5011033). url: <http://dx.doi.org/10.1063/1.5011033>.
- [39] Z. Holmes et al. *Connecting ansatz expressibility to gradient magnitudes and barren plateaus*. 2021. arXiv: [2101.02138](https://arxiv.org/abs/2101.02138) [quant-ph].
- [40] W. J. Huggins et al. "Efficient and noise resilient measurements for quantum chemistry on near-term quantum computers". In: *npj Quantum Information* 7.1 (Feb. 2021). issn: 2056-6387. doi: [10.1038/s41534-020-00341-7](https://doi.org/10.1038/s41534-020-00341-7). url: <http://dx.doi.org/10.1038/s41534-020-00341-7>.
- [41] IBM. *IBM Quantum Composer, 2021*. <https://quantum-computing.ibm.com/composer>. [Online; accessed 8-August-2021].
- [42] *IBM Quantum, 2021*. <https://quantum-computing.ibm.com/>. 2021.
- [43] *Is the Trotterized UCCSD Ansatz Chemically Well-Defined?* doi: [10.1021/acs.jctc.9b01083.s001](https://doi.org/10.1021/acs.jctc.9b01083.s001). url: <https://doi.org/10.1021/acs.jctc.9b01083.s001>.
- [44] J. K. Iverson and J. Preskill. "Coherence in logical quantum channels". In: *New Journal of Physics* 22.7 (Aug. 2020), p. 073066. issn: 1367-2630. doi: [10.1088/1367-2630/ab8e5c](https://doi.org/10.1088/1367-2630/ab8e5c). url: <http://dx.doi.org/10.1088/1367-2630/ab8e5c>.
- [45] A. F. Izmaylov, T.-C. Yen, and I. G. Ryabinkin. "Revising the measurement process in the variational quantum eigensolver: is it possible to reduce the number of separately measured operators?" In: *Chemical Science* 10.13 (2019), pp. 3746–3755. doi: [10.1039/c8sc05592k](https://doi.org/10.1039/c8sc05592k). url: <https://doi.org/10.1039/c8sc05592k>.
- [46] A. F. Izmaylov et al. "Unitary Partitioning Approach to the Measurement Problem in the Variational Quantum Eigensolver Method". In: *Journal of Chemical Theory and Computation* 16.1 (Nov. 2019), pp. 190–195. doi: [10.1021/acs.jctc.9b00791](https://doi.org/10.1021/acs.jctc.9b00791). url: <https://doi.org/10.1021/acs.jctc.9b00791>.
- [47] A. Jena, S. Genin, and M. Mosca. *Pauli Partitioning with Respect to Gate Sets*. 2019. arXiv: [1907.07859](https://arxiv.org/abs/1907.07859) [quant-ph].
- [48] P. Jordan and E. P. Wigner. "Über das Paulische Äquivalenzverbot". In: *The Collected Works of Eugene Paul Wigner*. Springer Berlin Heidelberg, 1993, pp. 109–129. doi: [10.1007/978-3-662-02781-3_9](https://doi.org/10.1007/978-3-662-02781-3_9). url: https://doi.org/10.1007/978-3-662-02781-3_9.
- [49] A. Kandala et al. "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets". In: *Nature* 549.7671 (Sept. 2017), pp. 242–246. doi: [10.1038/nature23879](https://doi.org/10.1038/nature23879). url: <https://doi.org/10.1038/nature23879>.
- [50] J.-S. Kim et al. *Hardware-efficient random circuits to classify noise in a multi-qubit system*. 2021. arXiv: [2104.10221](https://arxiv.org/abs/2104.10221) [quant-ph].

- [51] E. Knill, R. Laflamme, and W. H. Zurek. “Resilient quantum computation: error models and thresholds”. In: *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* 454.1969 (Jan. 1998), pp. 365–384. issn: 1471-2946. doi: [10.1098/rspa.1998.0166](https://doi.org/10.1098/rspa.1998.0166). url: <http://dx.doi.org/10.1098/rspa.1998.0166>.
- [52] R. Landauer. “The physical nature of information”. In: *Physics letters A* 217.4 (1996), pp. 188–193.
- [53] J. Lee et al. “Generalized Unitary Coupled Cluster Wave functions for Quantum Computation”. In: *Journal of Chemical Theory and Computation* 15.1 (Nov. 2018), pp. 311–324. issn: 1549-9626. doi: [10.1021/acs.jctc.8b01004](https://doi.org/10.1021/acs.jctc.8b01004). url: <http://dx.doi.org/10.1021/acs.jctc.8b01004>.
- [54] J. Lisenfeld et al. “Decoherence spectroscopy with individual two-level tunneling defects”. In: *Scientific Reports* 6.1 (Mar. 2016). issn: 2045-2322. doi: [10.1038/srep23786](https://doi.org/10.1038/srep23786). url: <http://dx.doi.org/10.1038/srep23786>.
- [55] J. Liu, Z. Li, and J. Yang. “An efficient adaptive variational quantum solver of the Schrödinger equation based on reduced density matrices”. In: *The Journal of Chemical Physics* 154.24 (June 2021), p. 244112. issn: 1089-7690. doi: [10.1063/5.0054822](https://doi.org/10.1063/5.0054822). url: <http://dx.doi.org/10.1063/5.0054822>.
- [56] J. M. Lourenço. *The NOVAthesis L^AT_EX Template User’s Manual*. NOVA University Lisbon. 2021. url: <https://github.com/joaomlourenco/novathesis/raw/master/template.pdf>.
- [57] J. K. L. MacDonald. “On the Modified Ritz Variation Method”. In: *Physical Review* 46.9 (Nov. 1934), pp. 828–828. doi: [10.1103/physrev.46.828](https://doi.org/10.1103/physrev.46.828). url: <https://doi.org/10.1103/physrev.46.828>.
- [58] E. Martín-López et al. “Experimental realization of Shor’s quantum factoring algorithm using qubit recycling”. In: *Nature Photonics* 6.11 (Oct. 2012), pp. 773–776. issn: 1749-4893. doi: [10.1038/nphoton.2012.259](https://doi.org/10.1038/nphoton.2012.259). url: <http://dx.doi.org/10.1038/nphoton.2012.259>.
- [59] J. R. McClean et al. “The theory of variational hybrid quantum-classical algorithms”. In: *New Journal of Physics* 18.2 (Feb. 2016), p. 023023. issn: 1367-2630. doi: [10.1088/1367-2630/18/2/023023](https://doi.org/10.1088/1367-2630/18/2/023023). url: <http://dx.doi.org/10.1088/1367-2630/18/2/023023>.
- [60] J. R. McClean et al. “Barren plateaus in quantum neural network training landscapes”. In: *Nature Communications* 9.1 (Nov. 2018). issn: 2041-1723. doi: [10.1038/s41467-018-07090-4](https://doi.org/10.1038/s41467-018-07090-4). url: <http://dx.doi.org/10.1038/s41467-018-07090-4>.
- [61] O. R. Meitei et al. “Gate-free state preparation for fast variational quantum eigensolver simulations”. In: *npj Quantum Information* 7.1 (Oct. 2021). doi: [10.1038/s41534-021-00493-0](https://doi.org/10.1038/s41534-021-00493-0). url: <https://doi.org/10.1038/s41534-021-00493-0>.

- [62] M. Motta et al. *Low rank representations for quantum simulation of electronic structure*. 2018. arXiv: 1808.02625 [physics.comp-ph].
- [63] J. von Neumann. “Probabilistic Logics and the Synthesis of Reliable Organisms From Unreliable Components”. In: *Automata Studies*. (AM-34). Princeton University Press, Dec. 1956, pp. 43–98. doi: 10.1515/9781400882618-003. url: <https://doi.org/10.1515%5C%2F9781400882618-003>.
- [64] M. A. Nielsen and I. L. Chuang. “Bibliography”. In: *Quantum Computation and Quantum Information*. Cambridge University Press, 2010, pp. 649–664. doi: 10.1017/cbo9780511976667.023. url: <https://doi.org/10.1017%2Fcbo9780511976667.023>.
- [65] B. O’Gorman et al. *Generalized swap networks for near-term quantum computing*. 2019. arXiv: 1905.05118 [quant-ph].
- [66] J. L. Park. “The concept of transition in quantum mechanics”. In: *Foundations of Physics* 1.1 (1970), pp. 23–33. doi: 10.1007/bf00708652. url: <https://doi.org/10.1007%2Fbf00708652>.
- [67] A. Peruzzo et al. “A variational eigenvalue solver on a photonic quantum processor”. In: *Nature Communications* 5.1 (July 2014). issn: 2041-1723. doi: 10.1038/ncomms5213. url: <http://dx.doi.org/10.1038/ncomms5213>.
- [68] J. Preskill. *Quantum computing 40 years later*. 2021. arXiv: 2106.10522 [quant-ph].
- [69] J. Preskill. “Quantum Computing in the NISQ era and beyond”. In: *Quantum* 2 (Aug. 2018), p. 79. doi: 10.22331/q-2018-08-06-79. url: <https://doi.org/10.22331%2Fq-2018-08-06-79>.
- [70] S. Resch and U. R. Karpuzcu. *Quantum Computing: An Overview Across the System Stack*. 2019. arXiv: 1905.07240 [quant-ph].
- [71] J. Romero et al. “Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz”. In: *Quantum Science and Technology* 4.1 (Oct. 2018), p. 014008. doi: 10.1088/2058-9565/aad3e4. url: <https://doi.org/10.1088/2058-9565/aad3e4>.
- [72] N. C. Rubin, R. Babbush, and J. McClean. “Application of fermionic marginal constraints to hybrid quantum algorithms”. In: *New Journal of Physics* 20.5 (May 2018), p. 053020. doi: 10.1088/1367-2630/aab919. url: <https://doi.org/10.1088%2F1367-2630%2Faab919>.
- [73] M. Schuld et al. “Evaluating analytic gradients on quantum hardware”. In: *Physical Review A* 99.3 (Mar. 2019). issn: 2469-9934. doi: 10.1103/physreva.99.032331. url: <http://dx.doi.org/10.1103/PhysRevA.99.032331>.
- [74] S. Sheldon et al. “Characterizing errors on qubit operations via iterative randomized benchmarking”. In: *Physical Review A* 93.1 (Jan. 2016). issn: 2469-9934. doi: 10.1103/physreva.93.012301. url: <http://dx.doi.org/10.1103/PhysRevA.93.012301>.

- [75] V. O. Shkolnikov et al. *Avoiding symmetry roadblocks and minimizing the measurement overhead of adaptive variational quantum eigensolvers*. 2021. arXiv: [2109.05340](https://arxiv.org/abs/2109.05340) [quant-ph].
- [76] P. W. Shor. “Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer”. In: *SIAM Journal on Computing* 26.5 (Oct. 1997), pp. 1484–1509. issn: 1095-7111. doi: [10.1137/S0097539795293172](https://doi.org/10.1137/S0097539795293172). url: <http://dx.doi.org/10.1137/S0097539795293172>.
- [77] P. W. Shor. “Scheme for reducing decoherence in quantum computer memory”. In: *Phys. Rev. A* 52 (4 Oct. 1995), R2493–R2496. doi: [10.1103/PhysRevA.52.R2493](https://doi.org/10.1103/PhysRevA.52.R2493). url: <https://link.aps.org/doi/10.1103/PhysRevA.52.R2493>.
- [78] D. R. Simon. “On the Power of Quantum Computation”. In: *SIAM Journal on Computing* 26.5 (Oct. 1997), pp. 1474–1483. doi: [10.1137/S0097539796298637](https://doi.org/10.1137/S0097539796298637). url: <https://doi.org/10.1137/S0097539796298637>.
- [79] A. Skolik et al. *Layerwise learning for quantum neural networks*. 2020. arXiv: [2006.14904](https://arxiv.org/abs/2006.14904) [quant-ph].
- [80] I. O. Sokolov et al. “Quantum orbital-optimized unitary coupled cluster methods in the strongly correlated regime: Can quantum algorithms outperform their classical equivalents?” In: *The Journal of Chemical Physics* 152.12 (Mar. 2020), p. 124107. issn: 1089-7690. doi: [10.1063/1.5141835](https://doi.org/10.1063/1.5141835). url: <http://dx.doi.org/10.1063/1.5141835>.
- [81] A. Steane. “Multiple particle interference and quantum error correction”. In: *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* 452.1954 (Nov. 1996), pp. 2551–2577. issn: 1471-2946. doi: [10.1098/rspa.1996.0136](https://doi.org/10.1098/rspa.1996.0136). url: <http://dx.doi.org/10.1098/rspa.1996.0136>.
- [82] M. Steudtner and S. Wehner. “Fermion-to-qubit mappings with varying resource requirements for quantum simulation”. In: *New Journal of Physics* 20.6 (June 2018), p. 063010. doi: [10.1088/1367-2630/aac54f](https://doi.org/10.1088/1367-2630/aac54f). url: <https://doi.org/10.1088/1367-2630/aac54f>.
- [83] K. J. Sung et al. “Using models to improve optimizers for variational quantum algorithms”. In: *Quantum Science and Technology* 5.4 (Oct. 2020), p. 044008. issn: 2058-9565. doi: [10.1088/2058-9565/abb6d9](https://doi.org/10.1088/2058-9565/abb6d9). url: <http://dx.doi.org/10.1088/2058-9565/abb6d9>.
- [84] A. Szabo and N. S. Ostlund. *Modern quantum chemistry: introduction to advanced electronic structure theory*. Courier Corporation, 2012.
- [85] H. L. Tang et al. “Qubit-ADAPT-VQE: An Adaptive Algorithm for Constructing Hardware-Efficient Ansätze on a Quantum Processor”. In: *PRX Quantum* 2.2 (Apr. 2021). issn: 2691-3399. doi: [10.1103/prxquantum.2.020310](https://doi.org/10.1103/prxquantum.2.020310). url: <http://dx.doi.org/10.1103/PRXQuantum.2.020310>.

- [86] W. G. Unruh. “Maintaining coherence in quantum computers”. In: *Physical Review A* 51.2 (Feb. 1995), pp. 992–997. doi: [10.1103/physreva.51.992](https://doi.org/10.1103/physreva.51.992). url: <https://doi.org/10.1103/physreva.51.992>.
- [87] A. V. Uvarov and J. D. Biamonte. “On barren plateaus and cost function locality in variational quantum algorithms”. In: *Journal of Physics A: Mathematical and Theoretical* 54.24 (May 2021), p. 245301. issn: 1751-8121. doi: [10.1088/1751-8121/abfac7](https://doi.org/10.1088/1751-8121/abfac7). url: <http://dx.doi.org/10.1088/1751-8121/abfac7>.
- [88] G. Verdon et al. *Learning to learn with quantum neural networks via classical neural networks*. 2019. arXiv: [1907.05415](https://arxiv.org/abs/1907.05415) [quant-ph].
- [89] V. Verteletskyi, T.-C. Yen, and A. F. Izmaylov. “Measurement optimization in the variational quantum eigensolver using a minimum clique cover”. In: *The Journal of Chemical Physics* 152.12 (Mar. 2020), p. 124114. issn: 1089-7690. doi: [10.1063/1.5141458](https://doi.org/10.1063/1.5141458). url: <http://dx.doi.org/10.1063/1.5141458>.
- [90] T. Volkoff and P. J. Coles. “Large gradients via correlation in random parameterized quantum circuits”. In: *Quantum Science and Technology* 6.2 (Jan. 2021), p. 025008. issn: 2058-9565. doi: [10.1088/2058-9565/abd891](https://doi.org/10.1088/2058-9565/abd891). url: <http://dx.doi.org/10.1088/2058-9565/abd891>.
- [91] L.-W. Wang and A. Zunger. “Solving Schrödinger’s equation around a desired energy: Application to silicon quantum dots”. In: *The Journal of Chemical Physics* 100.3 (Feb. 1994), pp. 2394–2397. doi: [10.1063/1.466486](https://doi.org/10.1063/1.466486). url: <https://doi.org/10.1063/1.466486>.
- [92] S. Wang et al. *Noise-Induced Barren Plateaus in Variational Quantum Algorithms*. 2021. arXiv: [2007.14384](https://arxiv.org/abs/2007.14384) [quant-ph].
- [93] D. Wecker, M. B. Hastings, and M. Troyer. “Progress towards practical quantum variational algorithms”. In: *Physical Review A* 92.4 (Oct. 2015). issn: 1094-1622. doi: [10.1103/physreva.92.042303](https://doi.org/10.1103/physreva.92.042303). url: <http://dx.doi.org/10.1103/PhysRevA.92.042303>.
- [94] J. D. Whitfield, J. Biamonte, and A. Aspuru-Guzik. “Simulation of electronic structure Hamiltonians using quantum computers”. In: *Molecular Physics* 109.5 (Mar. 2011), pp. 735–750. issn: 1362-3028. doi: [10.1080/00268976.2011.552441](https://doi.org/10.1080/00268976.2011.552441). url: <http://dx.doi.org/10.1080/00268976.2011.552441>.
- [95] M. Wilson et al. *Optimizing quantum heuristics with meta-learning*. 2019. arXiv: [1908.03185](https://arxiv.org/abs/1908.03185) [quant-ph].
- [96] W. K. Wootters and W. H. Zurek. “A single quantum cannot be cloned”. In: *Nature* 299.5886 (Oct. 1982), pp. 802–803. doi: [10.1038/299802a0](https://doi.org/10.1038/299802a0). url: <https://doi.org/10.1038/299802a0>.

-
- [97] T.-C. Yen, V. Verteletskyi, and A. F. Izmaylov. “Measuring All Compatible Operators in One Series of Single-Qubit Measurements Using Unitary Transformations”. In: *Journal of Chemical Theory and Computation* 16.4 (Mar. 2020), pp. 2400–2409. doi: [10.1021/acs.jctc.0c00008](https://doi.org/10.1021/acs.jctc.0c00008). url: <https://doi.org/10.1021%2Facs.jctc.0c00008>.
- [98] Y. S. Yordanov, D. R. M. Arvidsson-Shukur, and C. H. W. Barnes. “Efficient quantum circuits for quantum computational chemistry”. In: *Physical Review A* 102.6 (Dec. 2020). issn: 2469-9934. doi: [10.1103/PhysRevA.102.062612](https://doi.org/10.1103/PhysRevA.102.062612). url: <http://dx.doi.org/10.1103/PhysRevA.102.062612>.
- [99] Y. S. Yordanov et al. “Qubit-excitation-based adaptive variational quantum eigensolver”. In: *Communications Physics* 4.1 (Oct. 2021). issn: 2399-3650. doi: [10.1038/s42005-021-00730-0](https://doi.org/10.1038/s42005-021-00730-0). url: <http://dx.doi.org/10.1038/s42005-021-00730-0>.
- [100] P. Zhao et al. “High-Contrast ZZ Interaction Using Superconducting Qubits with Opposite-Sign Anharmonicity”. In: *Phys. Rev. Lett.* 125 (20 Nov. 2020), p. 200503. doi: [10.1103/PhysRevLett.125.200503](https://doi.org/10.1103/PhysRevLett.125.200503). url: <https://link.aps.org/doi/10.1103/PhysRevLett.125.200503>.

Implementation of Unitary Pauli String Exponentials

We wish to obtain a circuit to implement e^{-itH_k} for H_k a Pauli string ($\bigotimes_i P_i$, $P_i \in \{I_i, Z_i, X_i, Y_i\}$). This can be done from the circuit used for $H_k = \bigotimes_i Z_i$. The latter, as it was explained in subsection 2.1.3, only needs to apply a phase shift conditional on the parity of the computational basis states. In that section, a simple circuit that fit the purpose was presented, consisting of two ladders of CNOT gates (for computing and subsequently uncomputing the parity) and a single-qubit rotation.

The circuit for $H_k = \bigotimes_i P_i$ can be obtained by a simple modification. For any qubit i acted on by P_i in H_k , we find a basis rotation U_i that obeys

$$U_i^\dagger Z_i U_i = P_i$$

The unitary U_i will be different from the identity for any qubit that is acted on by a Pauli operator P_i different from Z_i . Examples of U_i include the Hadamard gate (for $P_i = X_i$) and a rotation around the X axis by $\pi/2$ (for $P_i = Y_i$).

This allows us to rewrite the expression as

$$e^{-itH_k} = e^{-it \bigotimes_i P_i} = e^{-it \bigotimes_i U_i^\dagger Z_i U_i}$$

Decomposing the exponential series, we obtain

$$\begin{aligned} &= \sum_{n=0}^{\infty} \frac{(-it \bigotimes_i U_i^\dagger Z_i U_i)^n}{n!} \\ &= I_1 \otimes I_2 \otimes I_3 \otimes I_4 - it \bigotimes_i U_i^\dagger Z_i U_i + \frac{(it)^2}{2} \left(\bigotimes_i U_i^\dagger Z_i U_i \right)^2 + \dots \end{aligned}$$

Using the properties of the tensor product ($(a \otimes b)(c \otimes d) = ac \otimes bd$) and the fact that the U_i are unitary ($U_i^\dagger U_i = U_i U_i^\dagger = I$), we can isolate an exponential series with argument free of the U_i :

$$\bigotimes_i U_i^\dagger \left(\sum_{n=0}^{\infty} \frac{(-it \bigotimes_i Z_i)^n}{n!} \right) \bigotimes_i U_i = \bigotimes_i U_i^\dagger \left(e^{-it \bigotimes_i Z_i} \right) \bigotimes_i U_i$$

So concluding, we have the identity

$$e^{-it \bigotimes_i P_i} = \bigotimes_i U_i^\dagger \left(e^{-it \bigotimes_i Z_i} \right) \bigotimes_i U_i$$

And we can apply any operator of the form $e^{-it} \otimes_i P_i$ by applying the circuit for $e^{-it} \otimes_i Z_i$, preceded and succeeded by single qubit rotations.

Hamiltonian of H2

This appendix contains the Hamiltonian of the H_2 molecule at an interatomic distance of 0.37\AA , obtained via OpenFermion [2]. The utilized basis set was *sto-3g*. Because it is a minimal basis set and no approximations were used, the molecule is represented by four spin-orbitals/qubits.

The choice to present the Hamiltonian of the simplest molecule was made due to the quartic growth of the number of terms with the number of orbitals: the Hamiltonian of a molecule such as *LiH* would already have almost two thousand operators.

The molecular Hamiltonian in the second quantization formalism is the following:

$$\begin{aligned}
&0.7151043390810812+ \\
&a_0^\dagger a_0 \cdot (-1.2533097866459773)+ \\
&a_1^\dagger a_1 \cdot (-1.2533097866459773)+ \\
&a_2^\dagger a_2 \cdot (-0.47506884877217576)+ \\
&a_3^\dagger a_3 \cdot (-0.47506884877217576)+ \\
&a_0^\dagger a_0^\dagger a_0 a_0 \cdot 0.3373779634072241+ \\
&a_0^\dagger a_0^\dagger a_2 a_2 \cdot 0.09060523100759854+ \\
&a_0^\dagger a_1^\dagger a_1 a_0 \cdot 0.3373779634072241+ \\
&a_0^\dagger a_1^\dagger a_3 a_2 \cdot 0.09060523100759854+ \\
&a_0^\dagger a_2^\dagger a_0 a_2 \cdot 0.09060523100759854+ \\
&a_0^\dagger a_2^\dagger a_2 a_0 \cdot 0.3318557006754069+ \\
&a_0^\dagger a_3^\dagger a_1 a_2 \cdot 0.09060523100759854+ \\
&a_0^\dagger a_3^\dagger a_3 a_0 \cdot 0.3318557006754069+ \\
&a_1^\dagger a_0^\dagger a_0 a_1 \cdot 0.3373779634072241+ \\
&a_1^\dagger a_0^\dagger a_2 a_3 \cdot 0.09060523100759854+ \\
&a_1^\dagger a_1^\dagger a_1 a_1 \cdot 0.3373779634072241+ \\
&a_1^\dagger a_1^\dagger a_3 a_3 \cdot 0.09060523100759854+ \\
&a_1^\dagger a_2^\dagger a_0 a_3 \cdot 0.09060523100759854+ \\
&a_1^\dagger a_2^\dagger a_2 a_1 \cdot 0.3318557006754069+ \\
&a_1^\dagger a_3^\dagger a_1 a_3 \cdot 0.09060523100759854+ \\
&a_1^\dagger a_3^\dagger a_3 a_1 \cdot 0.3318557006754069+
\end{aligned}$$

$$\begin{aligned}
& a_2^\dagger a_0^\dagger a_0 a_2 \cdot 0.3318557006754069+ \\
& a_2^\dagger a_0^\dagger a_2 a_0 \cdot 0.09060523100759854+ \\
& a_2^\dagger a_1^\dagger a_1 a_2 \cdot 0.3318557006754069+ \\
& a_2^\dagger a_1^\dagger a_3 a_0 \cdot 0.09060523100759854+ \\
& a_2^\dagger a_2^\dagger a_0 a_0 \cdot 0.09060523100759854+ \\
& a_2^\dagger a_2^\dagger a_2 a_2 \cdot 0.348825752245232+ \\
& a_2^\dagger a_3^\dagger a_1 a_0 \cdot 0.09060523100759854+ \\
& a_2^\dagger a_3^\dagger a_3 a_2 \cdot 0.348825752245232+ \\
& a_3^\dagger a_0^\dagger a_0 a_3 \cdot 0.3318557006754069+ \\
& a_3^\dagger a_0^\dagger a_2 a_1 \cdot 0.09060523100759854+ \\
& a_3^\dagger a_1^\dagger a_1 a_3 \cdot 0.3318557006754069+ \\
& a_3^\dagger a_1^\dagger a_3 a_1 \cdot 0.09060523100759854+ \\
& a_3^\dagger a_2^\dagger a_0 a_1 \cdot 0.09060523100759854+ \\
& a_3^\dagger a_2^\dagger a_2 a_3 \cdot 0.348825752245232+ \\
& a_3^\dagger a_3^\dagger a_1 a_1 \cdot 0.09060523100759854+ \\
& a_3^\dagger a_3^\dagger a_3 a_3 \cdot 0.348825752245232
\end{aligned}$$

Once we apply the Jordan-Wigner transformation [2.57](#) to obtain a Hamiltonian acting on qubits, we obtain:

$$\begin{aligned}
& -0.09706626816762856+ \\
& -0.04530261550379927 \cdot X_0 X_1 Y_2 Y_3+ \\
& 0.04530261550379927 \cdot X_0 Y_1 Y_2 X_3+ \\
& 0.04530261550379927 \cdot Y_0 X_1 X_2 Y_3+ \\
& -0.04530261550379927 \cdot Y_0 Y_1 X_2 X_3+ \\
& 0.17141282644776898 \cdot Z_0+ \\
& 0.1686889817036120 \cdot Z_0 Z_1+ \\
& 0.12062523483390417 \cdot Z_0 Z_2+ \\
& 0.16592785033770344 \cdot Z_0 Z_3+ \\
& 0.17141282644776903 \cdot Z_1+ \\
& 0.16592785033770344 \cdot Z_1 Z_2+ \\
& 0.12062523483390417 \cdot Z_1 Z_3+ \\
& -0.22343153690813572 \cdot Z_2+
\end{aligned}$$

$$0.174412876122616 \cdot Z_2 Z_3 + \\ -0.22343153690813572 \cdot Z_3$$

Adapt-VQE Gradient Formula

This appendix will show how the gradient formula from Adapt-VQE can be derived.

First of all, it is important to remember that the pool operators \hat{A}_i^\dagger are anti-hermitian:

$$\hat{A}_i^\dagger = -\hat{A}_i$$

This condition forces the unitarity of their parameterized and exponentiated versions that will actually appear in the ansatz. Since the variational parameters θ_i are real, we can see that

$$\left(e^{\theta_i \hat{A}_i}\right)^\dagger = e^{-\theta_i \hat{A}_i}$$

We want to calculate the derivative of the energy with respect to each possible variational parameter θ_i , as an indicator of how much the corresponding operator will impact the energy. For the i th operator, this derivative can be calculated as

$$\begin{aligned} & \frac{\partial}{\partial \theta_i} \left\langle \psi^{(n)} \left| e^{-\theta_i \hat{A}_i} \hat{H} e^{\theta_i \hat{A}_i} \right| \psi^{(n)} \right\rangle \\ &= \left\langle \psi^{(n)} \left| e^{-\theta_i \hat{A}_i} (-\hat{A}_i) \hat{H} e^{\theta_i \hat{A}_i} \right| \psi^{(n)} \right\rangle + \left\langle \psi^{(n)} \left| e^{-\theta_i \hat{A}_i} \hat{H} \hat{A}_i e^{\theta_i \hat{A}_i} \right| \psi^{(n)} \right\rangle \\ &= \left\langle \psi^{(n)} \left| e^{-\theta_i \hat{A}_i} [\hat{H}, \hat{A}_i] e^{\theta_i \hat{A}_i} \right| \psi^{(n)} \right\rangle \end{aligned}$$

All operators will be initialized with null coefficients: when a new operator $e^{\theta_i \hat{A}_i}$ is added to the ansatz, the corresponding variational parameter θ_i is set to zero. Thus, the relevant derivative is that calculated at this point, and we obtain 4.1 as desired.

$$\left. \frac{\partial E^{(n)}}{\partial \theta_i} \right|_{\theta_i=0} = \left\langle \psi^{(n)} \left| [\hat{H}, \hat{A}_i] \right| \psi^{(n)} \right\rangle \quad (\text{C.1})$$

The expectation value of the commutator can be measured efficiently on the quantum computer.

Given that \hat{H} is hermitian and \hat{A}_i is anti-hermitian, we can also write the commutator in 4.1 as

$$\begin{aligned} \left\langle \psi^{(n)} \left| [\hat{H}, \hat{A}_i] \right| \psi^{(n)} \right\rangle &= \left\langle \psi^{(n)} \left| \hat{H} \hat{A}_i \right| \psi^{(n)} \right\rangle - \left\langle \psi^{(n)} \left| \hat{A}_i \hat{H} \right| \psi^{(n)} \right\rangle \\ &= \left\langle \psi^{(n)} \left| \hat{H} \hat{A}_i \right| \psi^{(n)} \right\rangle - \left(\left\langle \psi^{(n)} \left| \hat{H} (-\hat{A}_i) \right| \psi^{(n)} \right\rangle \right)^\dagger \end{aligned}$$

$$= \langle \psi^{(n)} | \hat{H} \hat{A}_i | \psi^{(n)} \rangle + \left(\langle \psi^{(n)} | \hat{H} \hat{A}_i | \psi^{(n)} \rangle \right)^\dagger = 2\text{Re} \langle \psi^{(n)} | \hat{H} \hat{A}_i | \psi^{(n)} \rangle$$

In the sampling noise free, matrix algebra based simulations, formula C.2 was used instead of 4.1.

$$\left. \frac{\partial E^{(n)}}{\partial \theta_i} \right|_{\theta_i=0} = 2\text{Re} \langle \psi^{(n)} | \hat{H} \hat{A}_i | \psi^{(n)} \rangle \quad (\text{C.2})$$

