

Quantum Algorithms for Clustering and Covert Factoring

by

Dhruv Gopalakrishnan

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Applied Science
in
Electrical and Computer Engineering (Quantum Information)

Waterloo, Ontario, Canada, 2024

© Dhruv Gopalakrishnan 2024

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

This thesis contains two projects:

1. A quantum algorithm for clustering called qLUE developed in collaboration with CERN inspired by CERN's CLUE Event Reconstruction algorithm. The material in the thesis pertaining to this project is based on *qLUE: A Quantum Clustering Algorithm for Multi-Dimensional Datasets* which has been submitted to the *Frontiers in Quantum Science and Technology* journal (preprint available [here](#). The text in this thesis about this project is reproduced from this preprint with minor modifications) of which Dhruv Gopalakrishnan is the first author. Felice Pantaleo, Wahid Redjeb and Antonio Di Pilato gave the classical CLUE algorithm and provided a starting point for qLUE. Dhruv Gopalakrishnan and Luca Dellantonio developed the quantum versions of the various qLUE subroutines that use the Grover algorithm as a blackbox. Dhruv Gopalakrishnan coded up and ran the classical numerical simulations for qLUE as well as provided the Grover oracle for the various qLUE subroutines.
2. A scheme for hiding Shor's algorithm in Hamiltonian Simulation and Ground State Estimation. The material in this thesis pertaining to this project is based on *Pragmatic Obfuscation of Factoring in Hamiltonian Simulation and Ground State Estimation*, under preparation. Michele Mosca conceived of this innovative research direction and laid the foundation for this research. Dhruv Gopalakrishnan did the literature survey and technical analysis for the resource estimates for various versions of the scheme.

Abstract

This thesis is composed of two projects – a quantum algorithm for clustering based on CERN’s event reconstruction algorithm, and a scheme to hide Shor’s algorithm in Hamiltonian Simulation and Ground State Estimation circuits.

Clustering algorithms are at the basis of several technological applications, and are fueling the development of rapidly evolving fields such as machine learning. In the recent past, however, it has become apparent that they face challenges stemming from datasets spanning several spatial dimensions. In fact, the best-performing clustering algorithms scale linearly in the number of points, but quadratically with respect to the local density of points. In this work, we introduce qLUE, a quantum clustering algorithm that scales linearly in both the number of points and their density. qLUE is inspired by CLUE, CERN’s algorithm developed to address the challenging time and memory budgets of Event Reconstruction (ER) in future High-Energy Physics experiments. As such, qLUE marries decades of development with the quadratic speedup provided by quantum computers. We numerically test qLUE in several scenarios, demonstrating its effectiveness and proving it to be a promising route to handle complex data analysis tasks – especially in high-dimensional datasets with high densities of points. The code we developed for these simulations is available [here](#).

The advent of large-scale quantum computers promises transformative advances across various fields including optimization, materials science, and cryptography. However, this also poses a threat to traditional cryptography, due to Shor’s algorithm, which efficiently factors large integers. The existence of this algorithm undermines widely-used cryptographic protocols based on integer factorization and discrete logarithms. Even with Post-Quantum Cryptography, attacks of the “save now, decrypt later” type can compromise the security of critical systems. Keeping this mind, we would like to develop quantum systems that are designed specifically for benign applications such as Hamiltonian Simulation or Ground State Estimation – which could be of importance to the pharmaceutical industry. However, it cannot be taken for granted that even such a system is secure from malicious users attempting to run Shor’s Algorithm. In this note we propose the idea of using known circuit-to-Hamiltonian mappings to hide Shor’s algorithm in Hamiltonian simulation and Ground state estimation circuits. We provide the resource estimates for these mappings, and also propose some methods to potentially reduce these overhead costs.

Acknowledgments

Firstly, I would like to thank my lead supervisor Prof. Michele Mosca and my co-supervisor Prof. Na Young Kim for their constant support throughout my graduate program.

I would like to especially thank Prof. Michele Mosca for providing me with the opportunity and environment to focus on exciting research through my graduate school journey.

I am grateful to all my co-authors - Felice Pantaleo, Antonio Di Pilato, Wahid Redjeb, and especially Michele Mosca and Luca Dellantonio for all their help and support.

I thank Amolak Ratan Kalra and Andrew Jena for their help with the Pragmatic Obfuscation Project, and Isaac Chuang for prompting us to delve deeper into the QSVT for this project.

I am grateful for my friends without whom life in Waterloo wouldn't have been half as eventful.

Finally, I would like to thank my parents, for constantly supporting me through difficult decisions in my life.

Dedication

This is dedicated to my family, friends and everyone out there working hard to be a force for good.

Table of Contents

Author’s Declaration	ii
Statement of Contributions	iii
Abstract	iv
Acknowledgments	v
Dedication	vi
List of Figures	ix
List of Tables	x
1 Introduction	1
1.1 Grover’s Algorithm	2
1.2 Shor’s Algorithm	2
2 qLUE	4
2.1 Introduction	4
2.2 Overview and Setting	7
2.3 Local Density Computation	10
2.4 Find Nearest Higher	13

2.5	Find Seeds, Outliers, and Assign Clusters	14
2.6	Results	16
2.7	Scoring metrics	16
2.7.1	Homogeneity and Completeness scores	16
2.7.2	Energy-aware Homogeneity and Completeness scores	17
2.7.3	Noise	18
2.7.4	Overlap	19
2.8	Non-centroidal Clusters	19
2.9	Conclusion and Outlook	20
3	Pragmatic Obfuscation of Factoring in Hamiltonian Simulation and Ground State Estimation	24
3.1	Overview	24
3.2	Review: Circuit-to-Hamiltonian Methods	25
3.2.1	Local Hamiltonian Simulation	25
3.2.2	Hamiltonian Ground State Estimation	27
3.3	Hamiltonian Obfuscation	31
3.3.1	Divide-and-Conquer	31
3.4	Conclusion	34
	References	35
A	APPENDICES	45
A.1	Grover Oracles for qLUE	45

List of Figures

1.1	The quantum subroutine of Shor’s algorithm	3
2.1	Introductory figure to illustrate classical and quantum scaling laws for state-of-the-art solutions to the Unstructured Search Problem	6
2.2	Pictorial representation of the main subroutines of qLUE	8
2.3	Algorithm flow for Local Density computation and for Assigning Clusters .	11
2.4	Algorithm for finding Nearest Highers in qLUE	22
2.5	Numerical results from qLUE simulated on a classical machine.	23
3.1	Divide and Conquer framework	31
3.2	Iterative phase estimation scheme. Here, $k = 2$	32
3.3	QSVT-Based Phase Estimation	32
A.1	Circuit used for Grover Oracle in qLUE’s Local Density Computation step	46
A.2	Grover Oracle for qLUE’s Cluster Assignment step	47

List of Tables

3.1	Resource Estimates for Circuit to Hamiltonian Simulation Constructs . . .	28
3.2	Resource Estimates for Circuit to Local Hamiltonian Ground State Estimation	30
3.3	Overall Resource Estimates for Divide-and-Conquer based Hamiltonian Simulation: note that qubit overheads in the second column are additive. For the actual qubit resource cost add $O(n^d)$ to the listed value	33

Chapter 1

Introduction

When discussing the revolutionary potential of quantum computers, two hallmark algorithms immediately come to mind – Shor’s algorithm [1] for factoring large prime products (and computing discrete logarithms) which provides an exponential speedup over its state-of-the-art classical counterpart and Grover’s algorithm [2] for finding an element in an unstructured database which provides a quadratic speedup over its classical counterpart.

This pair of algorithms serves as a window into the double-edged sword nature of quantum computers, especially large-scale fault-tolerant quantum computers that major firms such as IBM [3] and Google [4] are working towards steadily. Grover’s algorithm and its generalized Amplitude Amplification [5] subroutine are have several applications both directly in efficient search and indirectly in various fields like bio-informatics for DNA sequencing through an efficient Grover-based solution to the string matching problem [6]. Shor’s algorithm, on the other hand, poses a significant risk to current encrypted systems that are based on the classical intractability of computing discrete logarithms and factoring products of large prime numbers. One way to mitigate these threats is via the use of Post-Quantum Cryptography (PQC) for encryption and Quantum Key Distribution (QKD) for secure communication. However, there is still the threat of “save now, decrypt later” style attacks which makes currently encrypted data vulnerable.

In this thesis, we leverage Grover’s algorithm to provide an efficient solution to high-dimensional clustering inspired by CERN’s CLUE algorithm for event reconstruction. This work is explained in Chapter 2 where we describe our qLUE algorithm. This work highlights serves to highlight some of the positive aspects of building large-scale quantum computers. In Chapter 3, we look to open a research dialogue on mitigating the potential threats of large-scale quantum computers. We look at one level of mitigation where we use quantum

computers designed especially for benign applications such as Hamiltonian Ground State Estimation and Simulation and ask if such a setting suffices to protect against the threat of using Shor’s algorithm for cryptanalysis by malignant players.

The rest of this section is a brief overview of the Shor and Grover algorithms.

1.1 Grover’s Algorithm

Grover’s algorithm is a quantum algorithm to solve the Unstructured Search Problem. From a superposition of all states to be searched over, Grover’s algorithm involves successive applications of two operators U_ψ and U_P to ensure that the measurement result at the end of the algorithm gives the search output with high probability. We use this algorithm extensively in our work. The inset of Fig. 2.3 describes the flow of this algorithm for our Local Density computation step. For m points, this involves $O(\sqrt{m})$ successive applications of the operators

$$U_\psi = 2|\psi\rangle\langle\psi| - 1, \quad (1.1a)$$

$$U_P|x\rangle = \begin{cases} -|x\rangle & f(x) = 1 \\ |x\rangle & f(x) = 0 \end{cases}. \quad (1.1b)$$

Here, $f(x) = 1$ when the current point satisfies a desired condition.

1.2 Shor’s Algorithm

In this section we review some basic aspects of Shor’s algorithm and Circuit-to-Hamiltonian mapping methods.

Shor [1] provided a polynomial-time quantum algorithm for factoring and discrete logarithm computation that reduces these problems to the order-finding problem. The goal of the order finding problem is to find order r of a

$$a^r \equiv 1 \pmod{N}; a, r \in \mathbb{Z} \quad (1.2)$$

where N is the number to be factored in the case of factoring and the order r is the smallest integer that satisfies the above equation. For more details, please refer to [1, 7] for more details. The quantum order finding algorithm uses the well-known quantum




Figure 1.1: The quantum subroutine of Shor's circuit runs phase estimation with the eigenvector register initialized to a superposition of eigenstates that sum up to $|1\rangle$. U here is as defined in Eq. 1.3

phase estimation method as a subroutine, this is shown in Fig. 1.1. The circuit for phase estimation is given explicitly in-terms of unitaries in Fig. 1.1 where

$$U |k\rangle = \begin{cases} |ak \pmod N\rangle; & 0 \leq k < N \\ |k\rangle; & N \leq k < 2n \end{cases} \quad (1.3)$$

where n is the smallest integer such that $N < 2^n$.

The input to the algorithm consists of 2 sets of qubits or "registers". The first "estimate" register is initialized to $|0\rangle$ and the final eigenvalue estimate is obtained by measuring these registers at the end of the algorithm. The second "eigenvector register" stores the eigenvector of U throughout the algorithm. The first step of the algorithm begins with the preparation of a fully entangled state in the estimate register with the application of Hadamard gates. The next step is quantum modular exponentiation where controlled- U^{2^j} gates are applied with the control on the j^{th} qubit of the estimate register and the target on the eigenvector registers. Finally a $QFT^\dagger$ is applied at the end to unentangle the state and ensure that the j^{th} register can be measured to obtain the j^{th} significant bit of the eigenvalue phase (i.e, the j^{th} bit of binary $\phi = 0.1001\dots$ where $U|\psi\rangle = e^{i\phi}|\psi\rangle$. Note here that $\phi < 1$). Since Shor gave us the efficient quantum factoring algorithm, several improvements have emerged. [8–10]

The bottleneck of this algorithm is quantum modular exponentiation. But this step can be implemented efficiently in polynomial time using the methods in [11].

Chapter 2

qLUE

2.1 Introduction

Clustering is a data analysis technique that is crucial in several fields, owing to its ability to uncover hidden patterns and structures within large datasets. It is essential for simplifying complex data, improving data organization, and enhancing decision-making processes [12–15]. For instance, clustering has been applied in marketing [16, 17], where it helps segment customers for targeted advertising [18], and in biology, for classifying genes and identifying protein interactions [19–22]. In the realm of computer science and artificial intelligence, it is invaluable for image [23] and speech recognition [24, 25], as well as recommendation systems [26, 27] that personalize content for users. Finally, clustering techniques are pivotal for Event Reconstruction (ER), where data points that originated from the same “event” must be collected together. In High-Energy Physics, for instance, clustering algorithms reconstruct the trajectories of subatomic particles in collider experiments. It is expected that the endcap high granularity calorimeter (HGCAL) [28] being built for the CMS detector at the High Luminosity Large Hadron Collider will provide extremely large volumes of data that must be tackled by new generations of clustering algorithms such as CLUE. The discovery of the Higgs boson [29], awarded the Nobel prize in 2012, was made possible by such algorithms.

ER enables the interpretation of data obtained from particle collision events, including those occurring at the Large Hadron Collider (LHC) at CERN. Several clustering algorithms like DBScan, K-Means, and Hierarchical Clustering among others [30–32] can be employed for ER. Our work is based on CERN’s CLustering of Energy (CLUE) algorithm [33], which is adopted by the CMS collaboration [34–36]. It is designed for the future HG-

CAL detector due to the limitations of the currently employed algorithms. Despite these limitations, such algorithms are already at the basis of several discoveries, such as the doubly charged tetraquark [37], the observation of four-top quark production in proton-proton collisions [34] and the study of rare B meson decays to two muons [35].

The efficiency of clustering algorithms, as illustrated by the CLUE algorithm [33], is crucial for handling extensive datasets. Initially designed for two-dimensional datasets, CLUE reduces the search complexity from $O(n^2)$ to $O(mn)$ through the use of local density and a tiling procedure, where n (m) represents the (average) number of points (per tile).

In the context of CLUE, where data sets are limited to two dimensions, m is small, making this approach to ER particularly effective. However, as the dimensionality of the dataset is incremented, the value of m generally increases exponentially. This is highlighted by Fig. 2.1(a), where for a d -dimensional lattice with a points per edge, m follows the relation $m = a^d$. This is a serious challenge to CLUE and classical clustering algorithms in general.

A first step towards extending CLUE to more dimensions is done by 3D-CLUE [33, 38]. In this work, data points from different layers of detectors are first projected onto a single $d = 2$ surface, where clustering is then performed. However, this projection from the original $d = 3$ dataset to a $d = 2$ surface comes at the cost of a slower algorithm since m becomes effectively larger. The solid lines in Fig. 2.1(b) show the increase in average points per tile in d -dimensional datasets made of the lattices in panel (a). While the improved performance of 3D-CLUE in ER tasks [33, 38] justifies the increased computational overhead, extending this enhancement to higher dimensions and larger datasets is challenging. Finding practical approaches to deal with datasets where d is large is therefore extremely important, not only for ER tasks but also in other fields such as gene analysis [39] and market segmentation in business [40].

Quantum computers provide a route to mitigate the complexity blowup arising from higher dimensional datasets. Ref. [41] addresses the task of jet clustering in High-Energy Physics, while Ref. [42] targets spectral clustering, which itself uses the efficient quantum analog of k -means clustering [43]. Other approaches include quantum k -medians clustering [44] and a quantum algorithm for density peak clustering [45].

In this work we develop qLUE, a CLUE-inspired quantum algorithm. Similarly to other quantum algorithms [46, 47], qLUE leverages the advantage provided by Grover Search [2]. A comparison of classical and quantum (Grover) runtimes is presented in Fig. 2.1(b), where the solid [dashed] lines refer to the classical $O(m)$ [quantum $O(\sqrt{m})$] scaling. As can be seen, the complexity advantage that Grover search provides can be substantial, particularly for large values of d or a .




Figure 2.1: Scaling of point density and complexities of classical and quantum algorithms for the unstructured search problem with dimension d . In (a), different d -dimensional lattices for $d = 1, 2, 3$ and $a = 3$ points per edge. In (b), best-known classical (solid lines) and quantum (dashed lines) algorithmic scalings for the Unstructured Search Problem [2] applied to square d -dimensional lattices with the values of a reported in the plot. Classically, the cost $O(m)$ reflects the need to iterate through all the m points to find the desired one. Grover achieves the same in $O(\sqrt{m})$ steps, providing a quadratic advantage. This advantage increases with the density of points in the considered dataset, which grows exponentially with respect to the dimension d according to $m = a^d$.

Overall, we find that qLUE performs well in a wide range of scenarios. With ER-inspired datasets as a specific example, we demonstrate that clusters are correctly reconstructed in typical experimental settings. Similar to other quantum approaches to clustering that rely on Grover Search [44, 48], qLUE also showcases a quadratic speedup compared to classical algorithms. The specific advantages of qLUE are its CLUE-inspired approach to cluster reconstruction (which has demonstrated extremely successful results [34, 35, 49]), and its consequent seamless integration with the classical framework currently employed by the CMS collaboration [33, 38, 50].

This chapter contains a description of the qLUE algorithm. Specifically, we provide a general overview of its subroutines - namely the Compute Local Density, Find Nearest Higher, and the Find Seeds, Outliers and Assign Clusters steps. We describe the results of our simulated version of qLUE on a classical computer in Sec. 2.6. In more detail, we explain the scoring metrics we use to quantify our results, and describe qLUE performance when the dataset is subject to noise and different clusters overlap. Conclusions and outlook are finally presented in Sec. 2.9.

qLUE is a quantum adaptation of CERN’s CLUE and 3D-CLUE algorithms [33, 38], that is specifically developed for ER, yet it is suitable to work with any (high dimensional) dataset. The main advantage of qLUE stems from employing Grover’s algorithm, which provides a quadratic speedup for the Unstructured Search Problem [2]. While qLUE is designed to work in arbitrary dimensions, for clarity we restrict ourselves to $d = 2$. This simplifies the following discussions and allows us to simulate qLUE with meaningful datasets on a classical computer. Generalizations to higher dimensions can be done following the steps outlined below. Furthermore, to provide a better connection with CLUE and 3D-CLUE, we employ a similar notation.

In Sec. 2.2, we offer an overview of the algorithm and its different subroutines. Sec. 2.3 is dedicated to the first subroutine of qLUE, namely, calculating the Local Density. We then explain how to determine the Nearest Higher ($\mathcal{N}_j$), Seeds, and Outliers in Sec. 2.4. Finally, Sec. 2.5 delves into the conclusive Cluster Assignment subroutine, where the points in the dataset are effectively hierarchically clustered.

2.2 Overview and Setting

As for CLUE and 3D-CLUE [33, 38], we consider a dataset with spatial coordinates and an energy for every point. Similar datasets can also be found in medical image analysis and segmentation [51, 52], in the analysis of asteroid reflectance spectra and hyperspectral astronomical imagery in astrophysics [53–55] and in gene analysis in bio-informatics [39, 56].

In $d = 2$ dimensions, the spatial coordinates X_j for point j are $X_j = [x_{j,1}, x_{j,2}]$, that are promptly generalized for larger values of d . Both CLUE and qLUE first perform tiling over the dataset to reduce the search and therefore enhance the efficiency of the algorithm. Tiling is the process of partitioning the dataset into a grid of rectangular tiles $\square_k$, where k is the tile index (see Fig. 2.2). Therefore, our input dataset comprises of point and tile indices j and k , respectively, the coordinates X_j , and a parameter E_j associated to each




Figure 2.2: Pictorial representation of the main subroutines of qLUE. In (a), the Local Density computation subroutine is represented. The consideration circle of radius d_c (light blue) centered at the base point j (black) contains all points (green) that satisfy $d_{i,j} \leq d_c$. This consideration circle intersects 2 tiles $\square_k$ (indexed by tile index k), highlighted in blue, that form the search space ($\mathcal{S}$). As per Eq. (2.2), the Local Density computation step determines the set of green points from all points in the search space (green and grey) and then computes the local density. In (b), we pictorially present the Find Nearest Higher ($\mathcal{N}_j$) subroutine. The consideration circle (green) around base point j (black) has radius d_m . This consideration circle, containing the green points as well as the Nearest Higher $\mathcal{N}_j$ (pink), intersects the 4 tiles highlighted in green, which form the search space $\mathcal{S}$. In (c), we describe the Find Seeds, Outliers and Assign Clusters subroutines. The seeds (red) and outliers (blue) are determined via Grover search on the dataset. In this specific example there are two clusters in the dataset whose non-seed points are in orange and purple, respectively. Followers (see main text) in these clusters are connected by dashed arrows. The Cluster Assignment subroutine is shown to be working on the orange cluster where the cluster $\mathcal{C}$ currently consists of the seed (red, dashed border) and the first of its followers (orange, dotted border). Followers are being found within the Dynamic Search Space (DSS, light red box with solid red border). The DSS is formed as the set of tiles $\square_k$ covered partially or fully by the minimum bounding box of the square windows that contains all the search spaces $\mathcal{S}$ of the points within $\mathcal{C}$.

point. Following CLUE's notation, we call E_j the energy, yet this should be considered as a label that can be employed to improve the clustering quality for any given dataset.

The tiling procedure of qLUE and CLUE enables searching only over Search Spaces $\mathcal{S}$ marked by the tiles in green in Fig. 2.2(a) as opposed to the full dataset. In case of CLUE, this allowed for an improvement in scaling from $O(n^2)$ to $O(mn)$. The scaling of qLUE is investigated in the following.

In this work, we employ a qRAM to store and access data, which is an essential building block for quantum computers. Following Ref. [57], we therefore assume that we can efficiently prepare the state

$$\sum_j |j\rangle \xrightarrow{\text{qRAM}} |j\rangle |D_j\rangle, \quad (2.1)$$

where D_j is the data associated with a given index j , e.g. the j^{th} point in the database. For convenience, here and throughout this paper we do not explicitly write the normalization factors of quantum states.

The qLUE algorithm consists of the following steps:

Local Density: The first step is to calculate the local density ρ_j of all points j [e.g, black point in Fig. 2.2(a)] that is defined by

$$\rho_j = E_j + \frac{1}{2} \sum_{d_{i,j} < d_c} E_i \quad (2.2)$$

and it is indicative of the energy in a neighborhood of point j . As can be seen from Eq. (2.2) and Fig. 2.2(a), ρ_j is a weighted sum over the energies E_i of all points i whose distance $d_{i,j} = \sqrt{\sum_{\alpha=1}^d (x_{i,\alpha} - x_{j,\alpha})^2}$ from the base point j is within a user-specified critical radius d_c that characterizes the consideration circle for the Local Density computation subroutine (light blue circle in the figure). As such, E_i is the energy of the i^{th} point which is $d_{i,j}$ away from point j . The choice of weight $1/2$ for E_j in the definition of ρ_j in Eq. (2.2) is empirically found to yield better performances for CLUE [33].

Find Nearest Higher: After calculating the local densities, we determine the nearest higher. The Nearest Higher $\mathcal{N}_j$ of a point j is the point nearest to j with a higher local density $\rho_{\mathcal{N}_j} > \rho_j$. As better explained in Sec. 2.5, the Nearest Higher are used to hierarchically cluster points together in the Cluster Assignment process at the end of qLUE. In Fig. 2.2(b), the Nearest Higher $\mathcal{N}_j$ of the base point j (black point) is the pink point.

Find Seeds, Outliers and Assign Clusters: As schematically represented in Fig. 2.2(c), seeds (red points) are the points whose distance $d_{j,\mathcal{N}_j}$ from their Nearest Higher $\mathcal{N}_j$ and whose local density ρ_j are lower bounded by user-defined thresholds. Outliers (blue points) are the points whose distance from Nearest Higher is similarly lower bounded but whose Local

Density has an upper threshold. As such a point j is

$$\text{a seed if } d_{\mathcal{N}_j,j} > d_c \text{ and } \rho_j > \tilde{\rho}, \quad (2.3a)$$

$$\text{an outlier if } d_{\mathcal{N}_j,j} > \delta d_c \text{ and } \rho_j < \tilde{\rho}. \quad (2.3b)$$

Here, δ is the Outlier Delta Factor that determines the upper bound on the allowed local density for outliers. Furthermore, $\tilde{\rho}$ is the critical density threshold – the lowest local density a point can have to be classified as a seed. Both δ and $\tilde{\rho}$ are user-specified and can be varied to enhance the quality of the output. Seeds are generally located in areas of high energy density, and will be employed as starting points to build clusters. Outliers are points that are likely to be noise in the dataset and are therefore discarded.

Once seeds and outliers are determined, the clusters are constructed. From the seeds, we iteratively combine “followers”. If point $\mathcal{N}_j$ is the Nearest Higher of point j , then point j is termed as $\mathcal{N}_j$ ’s follower. The follower of a point is most likely generated by the same process as the point itself (in the context of ER, by the same particle), and as such shall be included in the same cluster. In Fig. 2.2(c), the orange and purple points form two different clusters, and the followers of the points in the purple one are indicated by arrows.

2.3 Local Density Computation

In this section, we describe the subroutine (schematically represented in Fig. 2.3) that computes the Local Density ρ_j of the point j , as defined in Eq. (2.2). To perform the computation, all points i whose distance $d_{i,j}$ from point j is smaller than the threshold d_c need to be determined from the search space $\mathcal{S}$. This search space is the smallest set of tiles $\square_k$ required to cover the $d_{i,j} < d_c$ consideration circle. In Fig. 2.2(a), $\mathcal{S}$ is highlighted in light blue.

We shall refer to $\mathcal{S}$ as the local dataset that, as explained above, can be efficiently prepared with the qRAM [57]. To do so, we only require determining the tiles $\square_k$ that are in the search space, which can be done efficiently classically [33]. The initial state of this subroutine, after being prepared via the qRAM, is therefore

$$\sum_{k \in \mathcal{S}} \sum_{i \in \square_k} |i\rangle \xrightarrow{\text{qRAM}} \sum_{k \in \mathcal{S}} \sum_{i \in \square_k} |i\rangle |X_i, E_i\rangle, \quad (2.4)$$

where the index i is unique for each point in $\mathcal{S}$. $i \in \square_k$ indicate all indices within tile k [either of the light blue squares in Fig. 2.2(a)]. Ancillary qubits, omitted for clarity in Eq. (2.4), are employed within the Grover search (for more information, see App. 1.1).




Figure 2.3: Algorithm flow for Local Density computation and for Assigning Clusters. The quantum state is initialized in the green “Initialize” box. For Local Density Computation (Cluster Assignment), it comprises all points in $\mathcal{S}$ (in the DSS). The “Grover” (light blue) block performs U_ψ and U_P in succession $O(\sqrt{m})$ times, and returns all points satisfying the required condition. The inset considers the case of Local Density computation where the condition is $d_{i,j} < d_c$. For the cluster assignment step, we check if points in the DSS are followers of the points in the cluster $\mathcal{C}$ (see Sec. 2.5). The output of the Grover subroutine is then measured to yield an index that is checked for validity in the grey “Valid?” diamond. If the point satisfies the chosen condition, the Y branch is executed. Within the “Update” (light blue) step this point is then removed from either $\mathcal{S}$ or the DSS and stored to be returned in the “Return” orange box. Once all points are found, the “Valid?” condition triggers the N branch to terminate the algorithm. Depending on the chosen subroutine, the returned indices are employed to compute the Local Density from Eq. (2.2), or to construct $\mathcal{C}$.

At this stage, we must find the points i [green dots in Fig. 2.2(a)] that are within a radius of d_c from the base point j [black point in Fig. 2.2(a)]. As shown in Fig. 2.3, we perform Grover Search to prepare [58]

$$\sum_i |i\rangle |X_i, E_i\rangle \xrightarrow{\text{Grover}} \sum_{d_{i,j} < d_c} |i\rangle |X_i, E_i\rangle. \quad (2.5)$$

Here, the first register of the Grover output contains all points characterized by indices i such that $d_{i,j} < d_c$. As shown in the inset of the figure, the Grover Search consists of $O(\sqrt{m})$ repetitions (where m is the number of points in $\mathcal{S}$) of the U_ψ and U_P operators. U_P is the diffusion operator and U_ψ is the unitary associated with the oracle of Grover Search [2]. Further details regarding Grover Search and the unitaries we use for our algorithm can be found in App. 1.1.

When the algorithm is run, measurement either yields a point that satisfies this distance condition, or (if there are no valid indices left) an index that does not satisfy this condition. This is verified by the gray “Valid?” diamond in Fig. 2.3. The branched logic following this block ensures that the algorithm loops until all the required points are returned by the algorithm in the “Return” block.

Once we have obtained all indices i of points satisfying the distance condition ($d_{i,j} < d_c$), we perform the summation in Eq. (2.2). This is computed and stored in the original dataset for each point. The database is now updated using qRAM with local density values for all points where the j^{th} point in the database has the corresponding computed local density ρ_j .

The scaling of the subroutine that determines the local density of a single point is given by the number of points in the blue consideration circle in Fig. 2.2(a) such that $d_{i,j} < d_c$. If we say this number is p , $O(p)$ runs are required. This is therefore a $O(p\sqrt{m})$ algorithm as opposed to the $O(m)$ classical iterative algorithm for the Unstructured Search Problem. Since we do this for all points in the dataset, the full complexity of this step is $O(np\sqrt{m})$ as opposed to the classical $O(nm)$ version. Note that here, m is different for each point, but for the scaling expressions here and in other subroutines, one can use the number of points in the most dense tile of the dataset for m .

As a final remark, we highlight that it is in principle possible to design a unitary that computes the Local Density directly and stores the output in a quantum register. This unitary would remove the requirement of finding individually the indices i such that $d_{i,j} < d_c$, thus removing the overhead of p in $O(p\sqrt{m})$. However, designing this circuit is non-trivial and its depth may be large. This is therefore left for future investigations.

2.4 Find Nearest Higher

Here, we describe qLUE's subroutine for finding the Nearest Higher ($\mathcal{N}_j$) introduced in Sec. 2.2. As a reminder, $\mathcal{N}_j$ is the nearest point to the base point j whose local density $\rho_{\mathcal{N}_j}$ is more than the local density ρ_j of the base point, see Eq. (2.3a).

Similar to the initialization carried out for the Local Density Computation step, we use qRAM to initialize the quantum state

$$\sum_{k \in \mathcal{S}} \sum_{i \in \square_k} |i\rangle \xrightarrow{\text{qRAM}} \sum_{k \in \mathcal{S}} \sum_{i \in \square_k} |i\rangle |X_i\rangle |\rho_i\rangle. \quad (2.6)$$

Here, the indices i are within the tiles $\square_k$, as in Eq. (2.4), and $\mathcal{S}$ is the considered search space, schematically represented by the light green box in Fig. 2.2(b). This search space is determined from d_m as opposed to d_c , which is the user-defined threshold set to δd_c . Note that the energy E_i , employed for determining the densities ρ_i in Sec. 2.3, is hereon not required.

To find the Nearest Higher, we use a Grover-Enhanced Binary Search (GEBS) where each search step is enhanced by Grover's algorithm. The output of every Grover run,

$$\sum_{\substack{d_L < d_{i,j} < d_t, \\ \rho_i > \rho_b}} |i\rangle |X_i, \rho_i\rangle, \quad (2.7)$$

is a superposition over all points i whose distance $d_{i,j}$ from the base point j lies between the thresholds d_L and d_t . Furthermore, their local density ρ_i should be higher than that of the base ρ_j . At each step, d_L and d_t are updated based on whether a point satisfying the conditions in the gray diamond of Fig. 2.4(a) is found. Ancillary registers are used here as detailed in App. 1.1.

To better understand the algorithm, we provide a step-by-step walkthrough of the example in Fig. 2.4(b). The search space $\mathcal{S}$ is schematically represented by the inset in the right hand side, where each dot represents a point with a size that is proportional to its local density. The consideration circle (light green, dotted border) highlights all points within a radius $d_m = \delta d_c$. In this work, we set the outlier delta factor δ to 2. The consideration circle in the inset corresponds to $d_L = 0$ and $d_t = d_m$, shown in step (I). In the main panel, vertical lines refer to the steps (I-VI) of GEBS that are reported below, and schematically represent the distances of all points (colored dots) from the base point j (black one at the bottom).

GEBS starts with the higher threshold set as $d_t = d_m$ and the lower threshold $d_L = 0$ as shown in vertical line (I) of Fig. 2.4(b). Following the probabilistic nature of quantum mechanics, assume that the point with a red border indexed i is found after measuring the output of the Grover Search in Eq. (2.7). This triggers the updates in the Y branch in the diagram of Fig. 2.4(a), such that we assign $\mathcal{N}_j = i$ and update $d_t \mapsto (d_{i,j} + d_L)/2$. The point indexed i is then removed from the search space, as can be seen in (II). Now, since no point satisfies the conditions in the diamond of the flow diagram [see (II)] and d_t was just set to $(d_{i,j} + d_L)/2$, the B branch is carried out. This updates the thresholds d_t and d_L for the next iteration of the algorithm, see (III).

Now, assume that the new point with a red border is found [step (III)]. Updates in the Y branch of Fig. 2.4(a) are carried out again with a new index i and the search region is reduced to contain a single point. In the next step (IV), that point (yellow) is found and, for the third and last time, the nearest higher and the thresholds are triggered according to the Y branch. Next, since no point is found in (V), qLUE executes the updates in the B branch of the diagram. In the last iteration (VI), no points satisfy the desired conditions. The parameter d_t was just set to d_{t-1} , i.e, the subroutine just ran B which means that the A branch is now executed and $\mathcal{N}_j$ is returned.

The runtime complexity of the GEBS procedure, with m points in the search space $\mathcal{S}$, is $O(\alpha\sqrt{m})$ as opposed to $O(m)$ classically. The α term is due to the binary search procedure and depends on the size of the quantum register used to encode the distance. Specifically, for a chosen precision $2^{-\Delta}$ used for the positions of the points in the datasets, $\alpha = \Delta$. Again, since this is done for each point in the dataset, the full complexity of running this algorithm is $O(n\alpha\sqrt{m})$ as opposed to the $O(nm)$ classical solution.

2.5 Find Seeds, Outliers, and Assign Clusters

Once the Nearest Highers $\mathcal{N}_j$ are determined for all points j in the dataset, Seeds and Outliers are found via another Grover Search over all points in the dataset. As per the definition in Eq. (2.3a), Seeds [red points in Fig. 2.2(c)] are the points with highest local density within a neighborhood. Outliers [blue points in Fig. 2.2(c)] are mathematically described by Eq. (2.3b), are most likely noise, and therefore do not belong to any cluster.

Similar to the previous subroutines, the quantum registers for these procedures are initialized via qRAM. Seeds and outliers are then determined based on the corresponding conditions via Grover Search. Two quantum registers, the first marking whether a point is an outlier and the second to store the seed number – which is also the cluster number – are added to the quantum database.

The final subroutine of qLUE is the assignment of points to clusters. At this stage, outliers have been removed from the input dataset, as they have been already identified. The algorithm flow is the same as that of the Local Density step in Fig. 2.3. For a chosen seed s , we define $\mathcal{C}$ to be the set containing the indices of all points determined to be in the associated cluster at the end of this subroutine. To assign points to $\mathcal{C}$, we follow a procedure similar to that of the Local Density step in Fig. 2.3. In the “Initialize” step, $\mathcal{C}$ is initialized to $\{s\}$ and the quantum registers are initialized via qRAM to the state

$$\sum_{i \in \text{DSS}} |i\rangle \xrightarrow{\text{qRAM}} \sum_i |i\rangle |V_i\rangle, \quad (2.8a)$$

$$|V_i\rangle = |X_i, \rho_i, d_{\mathcal{N}_i, i}, X_{\mathcal{N}_i}\rangle. \quad (2.8b)$$

In the “Grover” block, we search over a superposition of points in the dataset which we call the Dynamic Search Space (DSS). The DSS differs from the search space $\mathcal{S}$ in the Local Density step as it is dynamic. This is because it depends on the points in $\mathcal{C}$, which are updated at each iteration. In Fig. 2.2(c), for instance, the red seed and the orange point both with black borders are the elements of the current $\mathcal{C}$. To find the DSS, a square window of edge $2d_m$ is first opened for every point in $\mathcal{C}$ (in the figure, the squares with the same border style as the corresponding points). A rectangular region (red box) is then obtained by finding the axis-aligned minimum bounding box for these windows. The set of tiles $\square_k$ covered partially or fully by this minimum bounding box is the DSS. For example, in Fig. 2.2(c), it comprises the 9 tiles highlighted in light red.

With a similar procedure as for the Local Density subroutine, the “Grover” block now systematically identifies all followers of all points within set $\mathcal{C}$. Here, in the “Update” step in Fig. 2.3, as the point found by the “Grover” block has passed the “Valid” condition, it is appended to $\mathcal{C}$. Once no more points are found, the “Return” block yields $\mathcal{C}$, following the same flow as the Local Density computation subroutine.

The complexity of the Cluster Assignment step is similar to the one of the Local Density Computation subroutine. The quantum advantage stems from the quadratic speedup provided by the Grover algorithm, which allows determining the follower faster if compared to CLUE. If there are f points in a cluster $\mathcal{C}$ and m points in the corresponding DSS, the classical complexity of the Cluster Assignment step is $O(m)$, while the quantum algorithm has a runtime of $O(f\sqrt{m})$. For the full Cluster Assignment step on the dataset, we know that since we are adding a single point to the cluster set for each non-outlier in the dataset. So there are at most n iterations of this search. However, the DSS increases arbitrarily in size for each follower found and appended to the cluster, increasing the number of points to search over. This variability makes it challenging to determine a precise worst-case

complexity for this step. Note that the sole difference between CLUE and qLUE in this step is the replacement of classical search in the DSS with Grover Search. Since this step is the bottleneck of CLUE as well as qLUE, qLUE provides an overall advantage.

2.6 Results

In this section, we test qLUE in multiple scenarios, each designed to investigate its performance for different settings. In Sec. 2.7, we introduce the scoring metrics used for our analysis. In Sec. 2.7.3, we describe the performance of the algorithm applied on a single cluster in a uniform noisy environment. In Sec. 2.7.4, we study the performance on overlapping clusters. Finally, in Sec. 2.8, we study the performance of qLUE on non-centroidal clusters with and without an energy profile.

2.7 Scoring metrics

2.7.1 Homogeneity and Completeness scores

We use the Homogeneity ($\mathcal{F}_H$) and Completeness ($\mathcal{F}_C$) scores [60] for quantifying the performance of qLUE. These metrics are defined in terms of the predicted cluster labels $\mathcal{C}_p$ obtained from qLUE, and the true cluster labels $\mathcal{C}_t$ of the generated dataset. $\mathcal{F}_H$ and $\mathcal{F}_C$ are defined based on the mutual information $I(\mathcal{C}_p : \mathcal{C}_t)$ and the Shannon entropy $H(\mathcal{C}_t)$ in addition to their joint Shannon entropy $H(\mathcal{C}_t, \mathcal{C}_p)$:

$$\mathcal{F}_H = \frac{I(\mathcal{C}_p : \mathcal{C}_t)}{H(\mathcal{C}_t)} \text{ and } \mathcal{F}_C = \frac{I(\mathcal{C}_p : \mathcal{C}_t)}{H(\mathcal{C}_p)}, \quad (2.9a)$$

$$H(\mathcal{C}_p) = - \sum_a P(\mathcal{C}_p \stackrel{?}{=} a) \log_2 P(\mathcal{C}_p \stackrel{?}{=} a), \quad (2.9b)$$

$$H(\mathcal{C}_t) = - \sum_b P(\mathcal{C}_t \stackrel{?}{=} b) \log_2 P(\mathcal{C}_t \stackrel{?}{=} b), \quad (2.9c)$$

$$H(\mathcal{C}_p, \mathcal{C}_t) = - \sum_a \sum_b P(a, b) \log_2 P(a, b), \quad (2.9d)$$

$$I(\mathcal{C}_p : \mathcal{C}_t) = H(\mathcal{C}_p) + H(\mathcal{C}_t) - H(\mathcal{C}_p, \mathcal{C}_t). \quad (2.9e)$$

Here, $P(\mathcal{C}_p \stackrel{?}{=} a)$ is the probability that qLUE classifies a point into cluster a . $P(\mathcal{C}_t \stackrel{?}{=} a)$ is the probability that qLUE assigns a point to its true cluster a . $P(a, b)$ is the probability that a point that is in cluster a in the true dataset is assigned to cluster b by qLUE.

qLUE applied to an input dataset yields homogeneity $\mathcal{F}_H = 1$ if all of the predicted clusters only contain data points that are members of a single true cluster. On the other hand, $\mathcal{F}_C = 1$ is obtained if all the data points that are members of a given true cluster are elements of the same reconstructed cluster. Therefore, these metrics are better suited to different scenarios. The impacts of noise and cluster overlap investigated in Secs. 2.7.3 and 2.7.4 are better captured by $\mathcal{F}_H$. Indeed, if qLUE incorrectly classifies noise points into predicted clusters, $\mathcal{F}_C$ is unaffected. On the other hand, $\mathcal{F}_C$ shall be employed when studying non-centroidal clusters in Sec. 2.8, since $\mathcal{F}_H = 1$ if one true cluster is divided by qLUE into many sub-clusters.

2.7.2 Energy-aware Homogeneity and Completeness scores

The standard metrics described in Sec. 2.7.1 are widely used for clustering tasks and also capture a lot of necessary information in our experiments. These metrics are ideal for situations where the given dataset has only the spatial coordinates in the dataset. However, we would like a metric that is able to use and capture some additional information contained in the energy variable of our dataset. This includes having a lower value of the metric if higher energy points such as seeds are misclassified as compared to lower energy points such as outliers. We here use modified, energy-aware versions [61] of the Homogeneity ($\mathcal{F}_H$) and Completeness ($\mathcal{F}_C$) scores [60]. These metrics are defined in terms of the predicted cluster labels $\mathcal{C}_p$ obtained from qLUE, and the true cluster labels $\mathcal{C}_t$ of the generated dataset. $\mathcal{F}_H$ and $\mathcal{F}_C$ are based on the (energy aware [61]) mutual information $I(\mathcal{C}_p : \mathcal{C}_t)$, the Shannon entropy $H(\mathcal{C}_t)$, and the joint Shannon entropy $H(\mathcal{C}_t, \mathcal{C}_p)$ [7]:

$$\mathcal{F}_H = \frac{I(\mathcal{C}_p : \mathcal{C}_t)}{H(\mathcal{C}_t)} \text{ and } \mathcal{F}_C = \frac{I(\mathcal{C}_p : \mathcal{C}_t)}{H(\mathcal{C}_p)}, \quad (2.10a)$$

$$H(\mathcal{C}_p) = - \sum_a \frac{E_a}{E} \log_2 \frac{E_a}{E}, \quad (2.10b)$$

$$H(\mathcal{C}_t) = - \sum_b \frac{E_b}{E} \log_2 \frac{E_b}{E}, \quad (2.10c)$$

$$H(\mathcal{C}_p, \mathcal{C}_t) = - \sum_a \sum_b \frac{E_{a,b}}{E} \log_2 \frac{E_{a,b}}{E}, \quad (2.10d)$$

$$I(\mathcal{C}_p : \mathcal{C}_t) = H(\mathcal{C}_p) + H(\mathcal{C}_t) - H(\mathcal{C}_p, \mathcal{C}_t). \quad (2.10e)$$

2.7.3 Noise

Here, we study the performance of qLUE for a single cluster in a noisy environment. We vary the number N_N of noise points sampled from a uniform distribution over a square region of fixed size. A cluster of N_C points with coordinates $X_j = [x_{j,1}, x_{j,2}]$ is sampled from the multivariate Gaussian distribution

$$pdf(X_j) = \frac{e^{-\frac{1}{2}(X_j - \mu_j)^T \Sigma^{-1} (X_j - \mu_j)}}{(2\pi)^{\frac{n}{2}} |\Sigma|^{\frac{1}{2}}}, \quad (2.11)$$

where $\mu = [\mu_{x_1}, \mu_{x_2}]^T$ is the mean of the distribution (set to $[0, 0]^T$ in our case) and Σ the covariance matrix. Here, we choose $\Sigma = \sigma \mathcal{I}$, with $\mathcal{I}$ being the identity matrix and σ a positive real number.

Examples of the generated clusters (in orange) and noise (in blue) are given in Fig. 2.5(a-b) for $N_N/N_C = 0.33$ at $\sigma = 32$ and $N_N/N_C = 1$ at $\sigma = 10$, respectively. The energy assigned to each point X_j in the cluster is given by $A \times pdf(X_j)$ [see Eq. (2.11)] with $A = 5 \times 10^2$. The energy of each noise point is randomly sampled between zero and one. This choice resembles the typical scenarios in ER tasks which CLUE [33] was designed for.

In Fig. 2.5(c), we show the variation of homogeneity score $\mathcal{F}_H$ with respect to the ratio N_N/N_C . We employ the values of σ reported in the legend, that are associated to different colors in the plot. As can be seen, the clustering performance is inversely proportional to both N_N/N_C and σ . When these parameters are small, the typical distance between cluster points is much smaller than that between noise points, and $\mathcal{F}_H$ approaches unity.

With a higher chance of labeling noisy points as within the cluster, however, $\mathcal{F}_H$ is lowered. As such, the degrading of $\mathcal{F}_H$ is proportional to the probability of a noise point being in the cluster region, which increases with both σ and N_N/N_C .

2.7.4 Overlap

Here, we consider the case of two circular clusters with N_1 and N_2 points respectively, each sampled from the multivariate Gaussian distribution in Eq. (2.11) and with $\Sigma = \sigma\mathcal{I}$. The energy profile is determined by $pdf(X_j)$ for coordinates X_j . The centers μ_1 and μ_2 (two instances of μ) are chosen to be $(r/2, 0)$ and $(-r/2, 0)$, respectively, such that the distance between the cluster centers is r .

In Fig. 2.5(d), we study the variation of homogeneity score $\mathcal{F}_H$ as a function of r/σ for several values of N_2/N_1 . The computed clusters for $r/\sigma = 2$ at $N_2/N_1 = 1$ and $r/\sigma = 2.67$ at $N_2/N_1 = 2$ are shown in panels (e) and (f), respectively, to showcase the typical scenarios considered here.

For all N_1/N_2 , $\mathcal{F}_H$ is zero for low r/σ (high overlap). There is then a region where $\mathcal{F}_H$ increases with r/σ and then saturates at unity for high r/σ (little to no overlap). When the two clusters are too nearby, i.e., $r/\sigma \ll 1$, they are in fact indistinguishable and qLUE labels all points together. Increasing the ratio r/σ makes the clusters move away from each other and thus qLUE can discern them. Importantly, large values of $\mathcal{F}_H$ are already attained when the clusters still have a significant overlap. In this scenario, employing the energy labels and the energy density considerably contributes to assigning the points to the right cluster. In fact, the nearest higher points are more likely to connect the points near or on the decision boundary with the more energetic core, thus separating the clusters better.

The performance of qLUE is also affected by the ratio N_1/N_2 . When one cluster contains more points than the other, it is more likely to “capture” points from the smaller. The resulting loss in homogeneity score $\mathcal{F}_H$ for low r/σ ratios is evident from Fig. 2.5(d), where it can be seen that clusters of similar sizes are better distinguished from each other.

2.8 Non-centroidal Clusters

Finally, we study the performance of qLUE on non-centroidal clusters. For this purpose, we use the Moons and Circles datasets in Fig. 2.5(g-j), generated using *scikit-learn* [59].

Two settings are considered - one where a uniform energy profile is applied over the points [panels (g-h)] and one where a linear gradient energy profile is employed [panels (i-j)].

In the latter case, for each cluster we assign the highest value of the energy to a single point and lower the energies of all other points proportionally to their x_2 coordinate. In the case of the moon dataset, $E = x_2$ for the upper moon (so the top point of the upper moon has the maximum energy in the cluster) and $E = 60 - x_2$ for the lower moon (so the bottom point has the highest energy in the cluster). For the circles, $E = |x_2 - 200|/10$ for the inner circle and $E = |x_2 + 100|/5$ for the outer one.

Since these datasets have no noise and are well separated, $\mathcal{F}_H$ is always one and we employ $\mathcal{F}_C$ to characterize the performance of qLUE. As in Fig. 2.5(g-h) the energy profile is uniform, and several points satisfy the seed condition. Therefore, qLUE groups each circle into several clusters, such that we obtain limited values for $\mathcal{F}_C$. On the contrary, cases with an energy profile assigned [Fig. 2.5(i-j)] results in less seeds that are better recognized by qLUE, and the completeness score $\mathcal{F}_C$ is considerably enhanced.

2.9 Conclusion and Outlook

We introduced qLUE, a novel quantum clustering algorithm designed to address the computational challenges associated with high-dimensional datasets. The significance of qLUE lies in its potential to efficiently cluster data leveraging quantum computing, mitigating the escalating computational complexity encountered by classical algorithms as dimensions increase. The algorithm’s ability to navigate high-dimensional spaces is particularly promising when the density of points is very large, such that local searches become too demanding for classical computers. Therefore, qLUE will be beneficial in multiple scenarios, ranging from quantum-enhanced machine learning [62, 63] to complex data analysis tasks [64].

According to our numerical results, qLUE works well and its performance is significantly enhanced when an energy profile is assigned. Specifically, we study qLUE in noisy environments, on overlapping clusters, and on non-centroidal datasets that are commonly used to benchmark clustering algorithms [65, 66]. In scenarios that are typically encountered in ER tasks, qLUE correctly reconstructs the true clusters to a high level of accuracy. On the other hand, an energy profile can significantly boost qLUE performance in the case of non-centroidal clusters. Our numerical results, backed up by the well-studied CLUE and by the quadratic speedup stemming from Grover search, make qLUE a promising candidate for addressing high-dimensional clustering problems [41, 42, 45].

As a first outlook, we identify the implementation of qLUE on NISQ hardware [67–73]. This requires a comprehensive consideration of real device constraints. Aspects such as circuit optimization [74], and the impact of noise will be critical and must be carefully addressed. Second, it is possible to improve the scaling of qLUE by devising a unitary that mitigates the need for repeating Grover’s algorithm and thereby eliminating the factors of p , α , and f in the scaling of the subroutines outlined in Secs. 2.3, 2.4 and 2.5 respectively. Finally, it is worth investigating variations of qLUE that improve the quality of clustering in different scenarios. For instance, one can devise more sophisticated criteria for the Nearest Higher or Local Density computation steps. One can also improve performance by performing exhaustive hyper-parameter searches or via hyper-parameter optimization algorithms [75].




Figure 2.4: (a) Diagrammatic representation of the algorithm. GEBS determines successive candidates for the “Nearest Higher” until the proper one is found. The quantum state in Eq. (2.6) is prepared in the “Initialize” step (green box). Grover Search (larger diamond) is then performed to find the points satisfying $d_L < d_{i,j} < d_t, \rho_i < \rho_j$. If this condition is satisfied (‘Y’ branch), d_t is updated and Grover run again. If not (‘N’ branch), control flows to the “?” diamond. The branch A is entered if the “?” condition is being checked for the first time or if branch B was just run. Branch B is entered if branch Y was just run. (b) The algorithm’s working is shown step-by-step (numbers at the bottom) for the search space $\mathcal{S}$ in the inset in the top right corner. The points are mapped to a line where the height represents the distance $d_{i,j}$ from the base point j (black dot at the bottom). The gray (orange) points are outside (inside) the green consideration circle with radius d_m [see also Fig. 2.2(b)]. At each step of GEBS, the thresholds d_L and d_t are updated according to the logic in panel (a). The dot with the red border indicates the current candidate for $\mathcal{N}_j$; when filled (empty) it is (not) found by Grover Search at that step. The yellow point is the Nearest Higher $\mathcal{N}_j$ found at the end of GEBS.




Figure 2.5: Numerical results from qLUE simulated on a classical machine. (a-c) qLUE’s performance in noisy environments. The dataset generated for these experiments and visualized in panels (a) and (b) consists of a cluster (noise) with $N_C = 750$ (N_N) points sampled from the Gaussian distribution in Eq. (2.11) (uniform distribution) over a square of size 500. The energy of noise points is sampled uniformly between zero and one, while each cluster point is assigned an energy that is the probability of being sampled multiplied by a factor $A = 500$. (a-b) Computed clusters at $N_N/N_C = 0.33, \sigma = 32$, and $N_N/N_C = 1, \sigma = 10$, respectively. In (c), $\mathcal{F}_H$ is plotted against N_N/N_C for the σ in the legend. (d-f) Performance for overlapping clusters. In (d), $\mathcal{F}_H$ vs r/σ is shown for $\sigma = 30$ and different ratios N_1/N_2 . Here, r is the distance between the centers of two clusters with $N_1 = 500$ and N_2 points, and we assign to each point an energy that is equal to its sampling probability in Eq. (2.11). (e-f) Computed clusters at $r/\sigma = 2.0, N_1/N_2 = 1$, and $r/\sigma = 2.67, N_1/N_2 = 2$, respectively. The shadowed regions in (c-d) represent the standard deviations of $\mathcal{F}_H$ over 30 iterations. (g-j) Performance over non-centroidal clusters of 500 points each generated from *scikit-learn* [59]. In (g-h) the points’ energy profile is uniform, while in (i-j) is varied linearly with respect to the distance such that each cluster has a single, most energetic point (see Sec. 2.8). For all experiments, d_c was set to 20 and $\tilde{\rho}$ was set to 25. (a-f) use the energy-aware metric in Eqs. (3.8) while in (g-j), since the energy profile is assigned by the user and is not part of the dataset itself, in the scoring process we set all points to have the same energy.

Chapter 3

Pragmatic Obfuscation of Factoring in Hamiltonian Simulation and Ground State Estimation

3.1 Overview

While quantum computing has the potential to revolutionize several industries, Shor’s algorithm poses a significant threat to present-day encryption systems that rely on the classical hardness of factoring and discrete logarithm computation. Keeping this in mind, one might consider developing quantum systems for specific benign applications like Hamiltonian Simulation (HS) and Hamiltonian Ground State Estimation (HGSE). Even with such a restriction, we study the possibility of hiding Shor’s algorithm in quantum circuits designed for HS and HGSE using known circuit-to-Hamiltonian mappings. We outline several methods to achieve this and list their resource estimates. Since these are problems that are likely to be of interest to the chemical and pharmaceutical industry, the methods discussed pose a potential threat.

From the theoretical perspective, quantum obfuscation is well studied [76–79]. Several theoretic results regarding formal obfuscatability of classical and quantum programs exist [76, 77, 80]. There has also been some related work on frameworks that aim to hide quantum circuits by garbling [77] where the key idea is to encode a quantum circuit such that no information about the quantum circuit can be recovered from the input and the encoding. The scheme has perfect correctness and information-theoretic security but has the drawback of a double-exponential worst-case circuit depth scaling for the encoding.

Such a framework enables the user to obfuscate the algorithm’s intent but does not provide a benign explanation. Another natural setting to study ideas closely related to obfuscation is Universal Blind Quantum Computation [78] which provides a framework to allow for quantum computation to be carried out such that the quantum hardware provider is blind to the inputs, outputs and computation and thus these are known only by the user. More recently in [79] Bartusek, et. al proposed a scheme to obfuscate pseudo-deterministic quantum circuits. This scheme is potentially powerful enough to formally obfuscate Shor’s algorithm based on the assumption of quantum hardness of Learning With Errors (LWE).

These general results help understand the art of the possible for obfuscating an implementation of Shor’s algorithm. However, adversaries do not need to deploy provably obfuscated circuits, especially if no constructive detection method is known, or practical, or actually being deployed in practice.

There is a natural trade-off between detectability of the obfuscated factoring circuit and the computational overhead to implement the obfuscated circuit.

In order to get a better handle on the pragmatic question of whether a specific quantum computation may in fact be hiding an implementation of Shor’s factoring algorithm, we examine some concrete examples without requiring formal guarantees of obfuscation and examine the resource requirements.

Sec. 1.2 contains a review of Shor’s algorithm. We provide a review of circuit-to-Hamiltonian mappings in Sec. 3.2 where we explain the procedure of mapping a quantum circuit to a Local Hamiltonian Simulation in Sec. 3.2.1, and the mapping to the Local Hamiltonian Ground State Estimation problem under the Adiabatic Computing framework in Sec. 3.2.2. In Sec. 3.3.1, we present our Divide-and-Conquer scheme which allows for better resource costs as compared to the methods outlined in Sections 3.2.1 and 3.2.2.

3.2 Review: Circuit-to-Hamiltonian Methods

In this section, we review the relevant Circuit-to-Hamiltonian methods which will later be used to obfuscate Shor’s Algorithm in Sec. 1.2 into Local Hamiltonian Simulation (LHS), 3.2.1 and Local Hamiltonian Ground State Estimation (LHGSE) 3.2.2.

3.2.1 Local Hamiltonian Simulation

In this sub-section we review the literature on mapping quantum circuits to Hamiltonian simulations. Concretely, this means that given a circuit U $|\psi\rangle$ with input $|\psi\rangle$, we want to

determine the Hamiltonian H for which simulation for time t would give us the output $U|\psi\rangle$. For the sake of simplicity in this section, we use $|\psi\rangle = |0\rangle^{\otimes n}$. The state $U|0\rangle^{\otimes n} = U_D U_{D-1} \dots U_1 |0\rangle^{\otimes n} = |0\rangle_d$ is the output of our depth- D quantum circuit on input $|0\rangle_d$.

Let us now look at Feynman's construction [81, 82] that performs this mapping. The main idea behind this construction is to have two separate registers - a clock that keeps track of the time-step t (depth) and the other stores the corresponding circuit output $U_t \dots U_1 |0\rangle_d$ at that time step. It thus makes sense to define a bipartite space,

$$\mathcal{H} = \mathcal{H}_c \otimes \mathcal{H}_d, \quad (3.1)$$

where $\mathcal{H}_c$ is the clock space that keeps track of the time steps and $\mathcal{H}_d$ is the data space that tracks the state of the circuit's action on the input at a given time.

We shall use the form of Feynman's Hamiltonian construction discussed in [82] and reproduced here in Eq. (3.2).

$$H_F = \sum_{t=1}^D |t\rangle \langle t-1|_c \otimes U_t + |t-1\rangle \langle t|_c \otimes U_t^\dagger. \quad (3.2)$$

One can see that evolving this Hamiltonian for time t by performing $e^{iH_F t}$ would give us $1 + iH_F t - (H_F t)^2/2 \dots$ when expanded using Taylor series. Applying this to $|\mathbf{0}\rangle_c \otimes |\mathbf{0}\rangle_d$ would give us a state that is in the span of $|\psi_t\rangle$

$$|\psi_t\rangle = |t\rangle_c \otimes U_t \dots U_1 |\mathbf{0}\rangle_d, t = 0, \dots, D. \quad (3.3)$$

This can be seen by noting that $H_F^i |\mathbf{0}\rangle$ in the i^{th} term of the Taylor expansion will give us terms in $\text{span}(|\psi_t\rangle)$. The state $|\mathbf{0}\rangle_c \otimes |\mathbf{0}\rangle_d$ can only be taken to $|j\rangle \otimes U_j U_{j-1} \dots U_1 |\mathbf{0}\rangle_d$. Applications of the $|t-1\rangle \langle t|_c \otimes U_t^\dagger$ type terms will only take $|j\rangle \otimes U_j U_{j-1} \dots U_1 |\mathbf{0}\rangle_d$ to $|j-1\rangle \otimes U_{j-1} U_{j-2} \dots U_1 |\mathbf{0}\rangle_d$ which lies in the span.

Since the result is in the span outlined above, if we measure register $|d\rangle$, we will obtain the circuit output $U_D \dots U_1 |\mathbf{0}\rangle_d$.

There are several methods such as Trotterization, Quantum Walks [83, 84] and the Quantum Singular Value Transform [85] to perform the simulation $e^{iH_F t}$ via a quantum circuit. An efficient method to perform this simulation outlined in [82] takes advantage of the fact that $H_F |\psi_j\rangle = \psi_{j-1} + \psi_{j+1}$ which is similar to a continuous time quantum walk on a line as discussed in [84]. Thus, the mixing time is given in the time averaged form where we pick a time randomly from 0 to $\Theta(N^2 \log N)$ so that we get close enough to the limiting distribution. This gives a $\Theta(N^{-1})$ probability of measuring the circuit in the $|D\rangle \otimes U_D U_{D-1} \dots U_1 |\mathbf{0}\rangle_d$ state. See [82, 84] for a complete discussion.

There are several ways to realize $|t\rangle_{clock}$. Related work is usually concerned with local clocks since Local Hamiltonians are of special interest, especially for the Local Hamiltonian problem [82, 86, 87] but for our purposes, we also note that a non-local (logarithmic locality) binary-encoded clock can also be used. The locality of H_F is at most 2 more than the locality of H_{clock} due to U_i being at most 2-local for all i .

Table 3.1 lists the complexities of the salient clock constructions that optimize either for locality, runtime, qubit overhead or the ability to arbitrarily increase the success probability of measuring $|D\rangle$ in the clock register and therefore obtaining the result of the quantum circuit computation $U_D \dots U_1 |\mathbf{0}\rangle_d$. We know that the depth of the Shor circuit is poly-logarithmic in the number of qubits, with the precise exponent and constants depending on various details such as which multiplication algorithm is used. So let's say the depth of the n qubit Shor circuit is n^d . The $O(\log n)$ -local construction uses a binary clock and thus has an $O(\log n)$ qubit overhead. The running time of this construction is $O(n^d)$ based on the results in [82] and the complexity of the continuous-time quantum walk on a line [84]. One might want to minimize the locality of this Hamiltonian so that it appears less suspicious since naturally occurring Hamiltonians (e.g. molecular Hamiltonians) are k -local for small constant k . That being said, the use of the Jordan-Wigner [88] or Bravyi-Kitaev [89] transformations could explain nonlocality, even of naturally occurring Hamiltonians. A unary pulse clock [82] allows for a 4-local construction (2-local clock and a locality of 2 in U_i on the data register) but this Hamiltonian has the same spectral properties as the quantum walk on a line that governed Feynman's clock so the mixing time is the same, i.e., $\Theta(n^{2d})$. The 5-local construction uses the unary domain-wall clock which has a spectral gap of $\Omega(D^{-2})$ and therefore an $O(n^d)$ runtime. The reduction in runtime could in some situations justify the locality trade-off and so we believe this construction is a suitable candidate for our purposes.

The quantum walk complexity [84] result is used in computing the running times in all of the constructions listed in Table. 3.1. In case one is not satisfied with the $\Theta(D^{-1})$ success probability offered by the quantum walk method, one can use the multi-cog construction in [82] at the cost of locality. The locality of the clock Hamiltonian in this case can be brought down to 8, making H_f 10-local at best.

3.2.2 Hamiltonian Ground State Estimation

Finding the ground state of a Hamiltonian is a task that is ubiquitous across quantum chemistry and therefore obfuscating Shor's algorithm into LHGSE would offer significant opportunities for covert malicious use of quantum hardware. While finding the exact

Table 3.1: Resource Estimates for Circuit to Hamiltonian Simulation Constructs: note that qubit overheads in the second column are additive. For the actual qubit resource cost add $O(n^d)$ to the listed value

Locality	Qubit overhead wrt $O(n)$ Shor	Arbitrarily increase success prob?	T (simulation time randomly chosen from between 0 and T)
4 [82]	$O(n^d)$	No	$\Theta(n^{3d})$
5 [82]	$O(n^d)$	No	$\Theta(n^{2d})$
10 [82]	$O(n^d)$	Yes	$\text{poly}(n^d)$
$O(\log(n))$	$O(\log(n))$	No	$\Theta(n^{2d})$

ground state is QMA-hard [87], algorithms such as Adiabatic Quantum Computing [86] can provide approximate solutions with quantifiable error-bounds. These error-bounds affect the run-time of the Adiabatic algorithm. In fact, Adiabatic Quantum Computing is proven to be equivalent to Universal Quantum Computing [86]. It is also possible to solve the Ground State Estimation problem variationally using a Variational Quantum Eigensolver (VQE) [90] which can also be run on a fault-tolerant computer [91].

Adiabatic Quantum Computing

The formulation of Adiabatic Quantum Computing [86] maps an n qubit, depth- D quantum circuit to the ground state of a Hamiltonian. The ground state can be prepared to error ϵ (in l_2 -norm) by an Adiabatic Computation where we simulate a certain time-dependent Hamiltonian for time $\Omega(T)$ [86, 92–94]. Concretely, if we want to determine the ground state of target Hamiltonian H_f , we should start with a Hamiltonian H_i whose unique ground state $|g\rangle_i$ is known. Then the evolution is performed on time-dependent Hamiltonian $H(s) = (1 - s)H_i + sH_f$. T is given by

$$T \geq \Omega\left(\frac{\|H_i - H_f\|^{1+\delta}}{\epsilon^\delta \min_{s \in [0,1]} (\Delta^{(2+\delta)} H(s))}\right) \|H_{max}\|, \quad (3.4)$$

where $\|\cdot\|$ is the spectral norm, $\Delta(H(s))$ is the spectral gap of $H(s)$, and $H_{max} = \max_s \|H(s)\|$

Now that we have a recipe for finding the ground state of a Hamiltonian, let us look at the original construct to map a circuit to the ground state of a Hamiltonian [87]. Here,

the strategy is still to keep track of time steps and the circuit output at each time step as in the earlier Feynman construction. So, we start our construction in the same Hilbert space as defined in Eq. (3.1). The ground state of our constructed Hamiltonian that we determine would be the so-called “history” state

$$|g\rangle = \frac{1}{\sqrt{D+1}} \sum_{t=0}^D |\psi\rangle_t, \quad (3.5)$$

where $|\psi_t\rangle$ is defined in Eq. (3.3). However, we use a unary clock to encode the time-step, i.e., $|t\rangle = |1^t 0^{D-t}\rangle$

Compared to Feynman’s strategy of making sure all powers of H_F were mappings within the span, Kitaev’s strategy differs in that we add terms to penalize unwanted states such that the unique ground state of the Hamiltonian H_f is the history state in Eq. 3.5. We use the description of Kitaev’s 5-Local Hamiltonian in the context of Adiabatic Computing in [86] reproduced here:

$$H_i = H_{ci} + H_{in} + H_c, \quad (3.6a)$$

$$H_f = \frac{1}{2} \sum_{t=0}^D H_t + H_{in} + H_c. \quad (3.6b)$$

We shall define and motivate each of these terms along the lines of the definition in [86]. Notice that our unary encoding disallows all states with a $|01\rangle$ in the bitstring while all other 2-qubit bitstrings may be present. So, $H_c = \sum_{t=0}^D |01\rangle \langle 01|_{t,t+1}^c$ penalizes all $|01\rangle$ bitstrings. The c superscript indicates that H_c is defined on $\mathcal{H}_c$. Since the input to the circuit is $|\mathbf{0}\rangle^c \otimes |\mathbf{0}\rangle^d$, we have $H_{in} = |0\rangle \langle 0|_1^c \otimes \sum_{i=1}^n |1\rangle \langle 1|_i^d$ that penalizes all states other than $|\mathbf{0}\rangle^d$ at time step 0. Finally for H_i , we have $H_{ci} = |1\rangle \langle 1|_1$ that ensures the ground state of H_i has the clock in $|\mathbf{0}\rangle^c$ as required. For H_f , we have $\frac{1}{2} \sum_{t=0}^D H_t$ checking that the transition from $|t\rangle$ to $|t+1\rangle$ corresponds to the application of U_t in the data space $\mathcal{H}_d$ thus making the history state the unique ground state with eigenvalue 0. Concretely,

$$\begin{aligned} H_t = & |100\rangle \langle 100|_{t-1,t,t+1}^c \otimes I - |110\rangle \langle 100|_{t-1,t,t+1} \otimes U_t - \\ & |110\rangle \langle 100|_{t-1,t,t+1} \otimes U_t^\dagger + |110\rangle \langle 110|_{t-1,t,t+1}^c \otimes I. \end{aligned} \quad (3.7)$$

Note that this definition of H_t will not work at the boundaries $t = 1$ and $t = D$. For these cases, the Hamiltonian uses two clock qubits instead of three, i.e, $H_1 = |00\rangle \langle 00|_{1,2} \otimes$

$I + |10\rangle\langle 10|_{1,2} \otimes -|10\rangle\langle 00|_{1,2} \otimes U_1 - |00\rangle\langle 10| \otimes U_1^\dagger$. Similarly for H_D , U_1 is replaced by U_D , $|00\rangle$ is replaced by $|10\rangle$ and $|10\rangle$ is replaced by $|11\rangle$.

For the same reasons as outlined in the previous section, we might be concerned with the locality of the Hamiltonian used. The best known method [95] gives a mapping from the n^d -depth Shor circuit to a 2-local Hamiltonian whose ground state can be found by the Adiabatic algorithm in $\text{poly}(n^d, \epsilon^{-1})$ time.

Table 3.2: Resource Estimates for Circuit to Local Hamiltonian Ground State Estimation: note that qubit overheads in the second column are additive. For the actual qubit resource cost add $O(n^d)$ to the listed value

Locality	Qubit overhead wrt $O(n)$ Shor	Arbitrarily increase success prob?	Running time
2 [96]	$O(n^d)$	No	$\text{poly}(n^d, \epsilon^{-1})$
3 [86]	$O(n^d)$	No	$O(n^{14d} \epsilon^{-16})$
4 [82]	$O(n^d)$	Yes	$\text{poly}(n^{2d})$
$O(\log(n))$	$O(\log(n))$	No	$\Theta(\epsilon^{-1} n^{6d})$

Variational Quantum Eigensolvers (VQEs)

VQEs [97, 98] have gained prominence as classical-quantum hybrid algorithms that can potentially be useful for Hamiltonian Ground State estimation in the NISQ era where quantum computers do not have a large number of qubits.

Given a Hamiltonian H , we would like to find the ground energy E_0 of this Hamiltonian. VQEs use the fact that the expected energy $\langle \psi | H | \psi \rangle \geq E_0$ for all states $|\psi\rangle$. For some ground state(s), $\langle \psi | H | \psi \rangle = E_0$. VQEs attempt to find this state and thereby the ground state energy by using parameterized guesses for the ground state. Concretely, the following algorithm is used:

First, an initial guess $|\psi(\theta)\rangle$ is prepared on a quantum computer where θ is the classical parameterization of the state. Second, the measurement $\langle \psi(\theta) | H | \psi(\theta) \rangle$ is performed to get the expected energy of the parameterized state. Now, a classical optimization algorithm decides the parameters for the quantum state to be prepared for the next iteration. The algorithm proceeds until convergence, i.e, until the difference between expected energies obtained at consecutive steps is less than an appropriate threshold.




Figure 3.1: The divide and conquer framework involves dividing the Quantum Phase Estimation subroutine in Shor's algorithm into segments and performing them on separate quantum hardware (QC_i)

While the original VQE is a classical-quantum hybrid algorithm, works such as [99, 100] indicate promising potential for fully quantum eigensolvers.

If we are to use VQEs for pragmatic obfuscation of Shor's algorithm into Hamiltonian simulation, we would use the same Hamiltonians as described in Sec. 3.2.2 with the localities and qubit overheads as listed in Table 3.2.2. As mentioned earlier, VQEs run until convergence. However, the number of steps a VQE takes to achieve this convergence is not predictable in general. So it is not possible to provide runtime upper bounds for VQEs.

3.3 Hamiltonian Obfuscation

In this section we will modify the above constructions to reduce overhead via a computing scheme that involves several quantum hardware that can communicate among each other via quantum teleportation. In the case of the earlier circuit-to-Hamiltonian mappings, although factoring is not immediately obvious, the inquisitive quantum hardware provider might be able to reverse engineer the Hamiltonian from the circuits.

3.3.1 Divide-and-Conquer

In order to provide a reasonable obfuscation of the Hamiltonian at lower overhead, one can think of breaking up the Shor circuit into segments as diagrammatically shown in Fig. 3.1 where each segment would directly map to a Local Hamiltonian evolution. The output from one segment is the input of the next. Since each of these segments runs on different quantum hardware, the required quantum states are communicated among them




Figure 3.2: Iterative phase estimation scheme. Here, $k = 2$




Figure 3.3: QSVT-Based Phase Estimation

via quantum teleportation. In general, one might be suspicious about how such a setup could be used for benign applications but one could envision this being used in a fault-tolerant VQE setup where each segment is an ansatz that represents an improvement over the previous ansatz segment.

Fig. 3.2 shows the iterative phase estimation procedure that allows us to use the measurements of the previous qubits to control the rotation gate on the next. This allows us to reduce the number of required qubits. A malicious party could claim that each chunk is attempting to obtain energy estimates of a molecule via measurements in the “estimate” register and is using said energy estimates to tweak subsequent ansatzes.

Each block would also contain a constant number of $c-U^{2^i}$ gates so that the precision of energy estimations of the Hamiltonian make sense from a quantum chemistry standpoint.

In the iterative phase estimation scheme, one might find the classically controlled rotations suspicious. Using the QSVT [101] for phase estimation gets rid of this problem by using the classical outputs to control the unitary acting on the eigenvector register instead of an appropriate rotation before measuring the phase qubit [101].

In the QSVT construction [101],

$$A_j(\theta) = \frac{1}{2}(I + e^{-2\pi i\theta}U^{2^j}), \quad (3.8a)$$

$$\Theta(x - c) = \begin{cases} -1 & x < c \\ 0 & x = c \\ 1 & x > c \end{cases}. \quad (3.8b)$$

The resource costs for the Divide-and-Conquer scheme for hiding the n^d -depth Shor circuit are outlined in Table 3.3. Since the obfuscation overhead scales with the size of the circuit being obfuscated, obfuscating components of the circuit leads to smaller overheads. (See the space-time resource scaling in Table 3.1 and the success probability scaling described in Sec. 3.2.1). The qubit overhead reduces by a factor of n since we divide the circuit into n segments. For each segment, the runtime now scales with n^{d-1} as opposed to n^d for the full circuit. For example, the 4-local construct would have a per-segment runtime of $\Theta(n^{3d-3})$ implying an overall runtime of $\Theta(n^{3d-2})$ as there are n segments. The last column of Table 3.3 is populated using such analyses for all the listed constructs. The success probability for each segment is $\Theta(n^{1-d})$ as opposed to $\Theta(n^{-d})$ for the full circuit. If we mitigate this error simply using multiple runs then the advantage is not significant as we would use $\Theta(n^{d-1})$ runs for each segment and since we have n segments we would use $\Theta(n^d)$ runs which is the same as the number of runs needed if we did not divide the circuit. However, using other error-correction schemes could potentially result in an advantage on this front as well.

Table 3.3: Overall Resource Estimates for Divide-and-Conquer based Hamiltonian Simulation: note that qubit overheads in the second column are additive. For the actual qubit resource cost add $O(n^d)$ to the listed value

Locality	Qubit overhead wrt $O(n)$ Shor	Arbitrarily increase success prob?	T (simulation time randomly chosen from between 0 and T)
4	$O(n^{d-1})$	No	$\Theta(n^{3d-2})$
5	$O(n^{d-1})$	No	$\Theta(n^{2d-1})$
10	$O(n^{d-1})$	Yes	$npoly(n^{d-1})$
$O(\log(n))$	$O(\log(n))$	No	$\Theta(n^{2d-1})$

3.4 Conclusion

In this paper, we proposed the idea of pragmatic obfuscation of Shor’s algorithm into Hamiltonian Simulation and Ground State Estimation using known circuit-to-Hamiltonian mappings. We listed the resource estimates for such a scheme and concluded that the resource overhead for such a scheme is quite high with the state-of-the-art mappings, especially if Hamiltonian locality is a factor.

The Divide-and-Conquer scheme described above, while possible for Ground State estimation, would not use Local Hamiltonians for each segment as the required Hamiltonian $H(s)$ (defined in Sec. 3.2.2) is dependent on the input state of the circuit (as can be seen in Eq. (A.1) which would determine the locality of the Hamiltonian. One could think of using distributed, fault-tolerant setups for such an attack where instead of aiming for locality, one would aim for maximum qubit utilization on the hardware but we leave this for future work.

One could also explore other quantum algorithms for Ground State Estimation where the initial state used has some overlap with the final state and the spectral gap between the ground and first excited state is lower bounded [102–104].

Overall, the approaches uncovered in this foray into the art of the possible for obfuscating Shor’s algorithm require resources that vastly exceed the resources needed for a direct unobfuscated implementation of Shor’s algorithm. Further investigations may bring the overhead closer to the realm of the possible.

References

- [1] Peter W. Shor. “Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer”. *SIAM Journal on Computing* 26.5 (1997), pp. 1484–1509. DOI: [10.1137/s0097539795293172](https://doi.org/10.1137/s0097539795293172).
- [2] Lov K. Grover. “A Fast Quantum Mechanical Algorithm for Database Search”. *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*. STOC '96. Philadelphia, Pennsylvania, USA: Association for Computing Machinery, 1996, 212–219. DOI: [10.1145/237814.237866](https://doi.org/10.1145/237814.237866).
- [3] Jay Gambetta. *Expanding the IBM Quantum roadmap to anticipate the future of quantum-centric supercomputing*. 2022.
- [4] Google. *Our focused and responsible approach to quantum computing*.
- [5] Gilles Brassard, Peter Høyer, Michele Mosca, and Alain Tapp. *Quantum amplitude amplification and estimation*. 2002. DOI: [10.1090/conm/305/05215](https://doi.org/10.1090/conm/305/05215).
- [6] Pradeep Niroula and Yunseong Nam. “A quantum algorithm for string matching”. *npj Quantum Information* 7.1 (2021), p. 37. DOI: [10.1038/s41534-021-00369-3](https://doi.org/10.1038/s41534-021-00369-3).
- [7] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010.
- [8] Martin Ekerå and Johan Håstad. *Quantum algorithms for computing short discrete logarithms and factoring RSA integers*. 2017.
- [9] Oded Regev. *An Efficient Quantum Factoring Algorithm*. 2024.
- [10] Clémence Cheignard, Pierre-Alain Fouque, and André Schrottenloher. *Reducing the Number of Qubits in Quantum Factoring*. Cryptology ePrint Archive, Paper 2024/222. <https://eprint.iacr.org/2024/222>. 2024.
- [11] Igor L. Markov and Mehdi Saeedi. “Constant-Optimized Quantum Circuits for Modular Multiplication and Exponentiation”. *arXiv e-prints*, arXiv:1202.6614 (Feb. 2012), arXiv:1202.6614. DOI: [10.48550/arXiv.1202.6614](https://doi.org/10.48550/arXiv.1202.6614).

- [12] Zuguang Gu and Daniel Hübschmann. “SimplifyEnrichment: A Bioconductor Package for Clustering and Visualizing Functional Enrichment Results”. *Genomics, Proteomics and Bioinformatics* 21.1 (June 2022), pp. 190–202. DOI: [10.1016/j.gpb.2022.04.008](https://doi.org/10.1016/j.gpb.2022.04.008).
- [13] Jelili Oyelade et al. “Data Clustering: Algorithms and Its Applications”. *2019 19th International Conference on Computational Science and Its Applications (ICCSA)*. 2019, pp. 71–81. DOI: [10.1109/ICCSA.2019.000-1](https://doi.org/10.1109/ICCSA.2019.000-1).
- [14] Tong Wu, Xinwang Liu, Jindong Qin, and Francisco Herrera. “Balance Dynamic Clustering Analysis and Consensus Reaching Process With Consensus Evolution Networks in Large-Scale Group Decision Making”. *IEEE Transactions on Fuzzy Systems* 29.2 (2021), pp. 357–371. DOI: [10.1109/TFUZZ.2019.2953602](https://doi.org/10.1109/TFUZZ.2019.2953602).
- [15] Giulia Caruso, Stefano Antonio Gattone, Francesca Fortuna, and Tonio Di Battista. “Cluster Analysis as a Decision-Making Tool: A Methodological Review”. *Decision Economics: In the Tradition of Herbert A. Simon’s Heritage*. Ed. by Edgardo Bucciarelli, Shu-Heng Chen, and Juan M. Corchado. Cham: Springer International Publishing, 2018, pp. 48–55.
- [16] Girish Punj and David W. Stewart. “Cluster Analysis in Marketing Research: Review and Suggestions for Application”. *Journal of Marketing Research* 20.2 (1983), pp. 134–148. DOI: [10.1177/002224378302000204](https://doi.org/10.1177/002224378302000204).
- [17] Jih-Jeng Huang, Gwo-Hshiung Tzeng, and Chorng-Shyong Ong. “Marketing segmentation using support vector clustering”. *Expert Systems with Applications* 32.2 (2007), pp. 313–317. DOI: <https://doi.org/10.1016/j.eswa.2005.11.028>.
- [18] Xiaohui Wu et al. “Probabilistic Latent Semantic User Segmentation for Behavioral Targeted Advertising”. *Proceedings of the Third International Workshop on Data Mining and Audience Intelligence for Advertising*. ADKDD ’09. Paris, France: Association for Computing Machinery, 2009, 10–17. DOI: [10.1145/1592748.1592751](https://doi.org/10.1145/1592748.1592751).
- [19] Pratik Dutta, Sriparna Saha, Sanket Pai, and Aviral Kumar. “A Protein Interaction Information-based Generative Model for Enhancing Gene Clustering”. *Scientific Reports* 10.1 (2020), p. 665. DOI: [10.1038/s41598-020-57437-5](https://doi.org/10.1038/s41598-020-57437-5).
- [20] Wai-Ho Au, K.C.C. Chan, A.K.C. Wong, and Yang Wang. “Attribute clustering for grouping, selection, and classification of gene expression data”. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* 2.2 (2005), pp. 83–101. DOI: [10.1109/TCBB.2005.17](https://doi.org/10.1109/TCBB.2005.17).

- [21] Jianxin Wang, Min Li, Youping Deng, and Yi Pan. “Recent advances in clustering methods for protein interaction networks”. *BMC Genomics* 11.3 (2010), S10. DOI: [10.1186/1471-2164-11-S3-S10](https://doi.org/10.1186/1471-2164-11-S3-S10).
- [22] Sitaram Asur, Duygu Ucar, and Srinivasan Parthasarathy. “An ensemble framework for clustering protein–protein interaction networks”. *Bioinformatics* 23.13 (July 2007), pp. i29–i40. DOI: [10.1093/bioinformatics/btm212](https://doi.org/10.1093/bioinformatics/btm212).
- [23] G.B. Coleman and H.C. Andrews. “Image segmentation by clustering”. *Proceedings of the IEEE* 67.5 (1979), pp. 773–785. DOI: [10.1109/PROC.1979.11327](https://doi.org/10.1109/PROC.1979.11327).
- [24] Kishore Kumar R, Lokendra Birla, and Sreenivasa Rao K. “A robust unsupervised pattern discovery and clustering of speech signals”. *Pattern Recognition Letters* 116 (2018), pp. 254–261. DOI: <https://doi.org/10.1016/j.patrec.2018.10.035>.
- [25] Jianlong Chang, Lingfeng Wang, Gaofeng Meng, Shiming Xiang, and Chunhong Pan. “Deep Adaptive Image Clustering”. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. 2017.
- [26] Andriy Shepitsen, Jonathan Gemmell, Bamshad Mobasher, and Robin Burke. “Personalized Recommendation in Social Tagging Systems Using Hierarchical Clustering”. *Proceedings of the 2008 ACM Conference on Recommender Systems*. RecSys ’08. Lausanne, Switzerland: Association for Computing Machinery, 2008, 259–266. DOI: [10.1145/1454008.1454048](https://doi.org/10.1145/1454008.1454048).
- [27] Vincent Schickel-Zuber and Boi Faltings. “Using Hierarchical Clustering for Learning Theontologies Used in Recommendation Systems”. *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD ’07. San Jose, California, USA: Association for Computing Machinery, 2007, 599–608. DOI: [10.1145/1281192.1281257](https://doi.org/10.1145/1281192.1281257).
- [28] “The Phase-2 Upgrade of the CMS Endcap Calorimeter” (2017). DOI: [10.17181/CERN.IV8M.1JY2](https://doi.org/10.17181/CERN.IV8M.1JY2).
- [29] G. Aad et al. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. *Physics Letters B* 716.1 (2012), pp. 1–29. DOI: [10.1016/j.physletb.2012.08.020](https://doi.org/10.1016/j.physletb.2012.08.020).
- [30] F D Amaro et al. “Directional iDBSCAN to detect cosmic-ray tracks for the CYGNO experiment”. *Measurement Science and Technology* 34.12 (2023), p. 125024. DOI: [10.1088/1361-6501/acf402](https://doi.org/10.1088/1361-6501/acf402).

- [31] S A Rodenko, A G Mayorov, V V Malakhov, I K Troitskaya, and on behalf of PAMELA collaboration. “Track reconstruction of antiprotons and antideuteron in the coordinate-sensitive calorimeter of PAMELA spectrometer using the Hough transform”. *Journal of Physics: Conference Series* 1189.1 (2019), p. 012009. DOI: [10.1088/1742-6596/1189/1/012009](https://doi.org/10.1088/1742-6596/1189/1/012009).
- [32] Christoph Dalitz et al. “Automatic trajectory recognition in Active Target Time Projection Chambers data by means of hierarchical clustering”. *Computer Physics Communications* 235 (2019), pp. 159–168. DOI: <https://doi.org/10.1016/j.cpc.2018.09.010>.
- [33] Marco Rovere, Ziheng Chen, Antonio Di Pilato, Felice Pantaleo, and Chris Seez. “CLUE: A Fast Parallel Clustering Algorithm for High Granularity Calorimeters in High-Energy Physics”. *Frontiers in Big Data* 3 (2020). DOI: [10.3389/fdata.2020.591315](https://doi.org/10.3389/fdata.2020.591315).
- [34] A. Hayrapetyan et al. “Observation of four top quark production in proton-proton collisions at $\sqrt{s}=13\text{TeV}$ ”. *Physics Letters B* 847 (2023), p. 138290. DOI: <https://doi.org/10.1016/j.physletb.2023.138290>.
- [35] A. Tumasyan et al. “Measurement of the $B_s^0 \rightarrow \mu^+ \mu^-$ decay properties and search for the $B^0 \rightarrow \mu^+ \mu^-$ decay in proton-proton collisions at $\sqrt{s}=13\text{TeV}$ ”. *Physics Letters B* 842 (2023), p. 137955. DOI: <https://doi.org/10.1016/j.physletb.2023.137955>.
- [36] Aram Hayrapetyan et al. *Search for new physics with emerging jets in proton-proton collisions at $\sqrt{s} = 13 \text{ TeV}$* . Tech. rep. Submitted to the Journal of High Energy Physics. All figures and tables can be found at <http://cms-results.web.cern.ch/cms-results/public-results/publications/EXO-22-015> (CMS Public Pages). Geneva: CERN, 2024.
- [37] R. Aaij et al. “First Observation of a Doubly Charged Tetraquark and Its Neutral Partner”. *Phys. Rev. Lett.* 131 (4 2023), p. 041902. DOI: [10.1103/PhysRevLett.131.041902](https://doi.org/10.1103/PhysRevLett.131.041902).
- [38] Erica Brondolin. *CLUE a clustering algorithm for current and future experiments*. Tech. rep. Geneva: CERN, 2022.
- [39] Md Rezaul Karim et al. “Deep learning-based clustering approaches for bioinformatics”. *Briefings in Bioinformatics* 22.1 (Feb. 2020), pp. 393–415. DOI: [10.1093/bib/bbz170](https://doi.org/10.1093/bib/bbz170).
- [40] Jian Zhou, Linli Zhai, and Athanasios A. Pantelous. “Market segmentation using high-dimensional sparse consumers data”. *Expert Systems with Applications* 145 (2020), p. 113136. DOI: <https://doi.org/10.1016/j.eswa.2019.113136>.

- [41] Annie Y. Wei, Preksha Naik, Aram W. Harrow, and Jesse Thaler. “Quantum algorithms for jet clustering”. *Phys. Rev. D* 101 (9 2020), p. 094015. DOI: [10.1103/PhysRevD.101.094015](https://doi.org/10.1103/PhysRevD.101.094015).
- [42] Iordanis Kerenidis and Jonas Landman. “Quantum spectral clustering”. *Phys. Rev. A* 103 (4 2021), p. 042415. DOI: [10.1103/PhysRevA.103.042415](https://doi.org/10.1103/PhysRevA.103.042415).
- [43] Iordanis Kerenidis, Jonas Landman, Alessandro Luongo, and Anupam Prakash. “q-means: A quantum algorithm for unsupervised machine learning”. *Advances in Neural Information Processing Systems*. Ed. by H. Wallach et al. Vol. 32. Curran Associates, Inc., 2019.
- [44] Esma Aïmeur, Gilles Brassard, and Sébastien Gambs. “Quantum Clustering Algorithms”. *Proceedings of the 24th International Conference on Machine Learning*. ICML ’07. Corvalis, Oregon, USA: Association for Computing Machinery, 2007, 1–8. DOI: [10.1145/1273496.1273497](https://doi.org/10.1145/1273496.1273497).
- [45] Duarte Magano, Lorenzo Buffoni, and Yasser Omar. “Quantum density peak clustering”. *Quantum Machine Intelligence* 5.1 (2023), p. 9. DOI: [10.1007/s42484-022-00090-0](https://doi.org/10.1007/s42484-022-00090-0).
- [46] Davide Nicotra et al. *A quantum algorithm for track reconstruction in the LHCb vertex detector*. 2023.
- [47] Cenk Tüysüz et al. “Particle Track Reconstruction with Quantum Algorithms”. *EPJ Web of Conferences* 245 (2020). Ed. by C. Doglioni et al., p. 09013. DOI: [10.1051/epjconf/202024509013](https://doi.org/10.1051/epjconf/202024509013).
- [48] Diogo Pires, Pedrame Bargassa, João Seixas, and Yasser Omar. *A Digital Quantum Algorithm for Jet Clustering in High-Energy Physics*. 2021.
- [49] CMS Collaboration. *Review of top quark mass measurements in CMS*. 2024.
- [50] CMS Collaboration. *Development of the CMS detector for the CERN LHC Run 3*. 2023.
- [51] Bahjat F Qaqish, Jonathon J O’Brien, Jonathan C Hibbard, and Katie J Clowers. “Accelerating high-dimensional clustering with lossless data reduction”. *Bioinformatics* 33.18 (May 2017), pp. 2867–2872. DOI: [10.1093/bioinformatics/btx328](https://doi.org/10.1093/bioinformatics/btx328).
- [52] H.P. Ng, S.H. Ong, K.W.C. Foong, P.S. Goh, and W.L. Nowinski. “Medical Image Segmentation Using K-Means Clustering and Improved Watershed Algorithm”. *2006 IEEE Southwest Symposium on Image Analysis and Interpretation*. 2006, pp. 61–65. DOI: [10.1109/SSIAI.2006.1633722](https://doi.org/10.1109/SSIAI.2006.1633722).

- [53] L. Galluccio, O. Michel, P. Bendjoya, and E. Slezak. “Unsupervised Clustering on Astrophysics Data: Asteroids Reflectance Spectra Surveys and Hyperspectral Images”. *Classification and Discovery in Large Astronomical Surveys*. Ed. by Coryn A. L. Bailer-Jones. Vol. 1082. American Institute of Physics Conference Series. Dec. 2008, pp. 165–171. DOI: [10.1063/1.3059034](https://doi.org/10.1063/1.3059034).
- [54] Michael J. Gaffey. “Space weathering and the interpretation of asteroid reflectance spectra”. *Icarus* 209.2 (2010), pp. 564–574. DOI: <https://doi.org/10.1016/j.icarus.2010.05.006>.
- [55] Angela F. Gao et al. “Generalized Unsupervised Clustering of Hyperspectral Images of Geological Targets in the Near Infrared”. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*. 2021, pp. 4294–4303.
- [56] Jelili Oyelade et al. “Clustering Algorithms: Their Application to Gene Expression Data”. *Bioinformatics and Biology Insights* 10 (2016). PMID: 27932867, BBI.S38316. DOI: [10.4137/BBI.S38316](https://doi.org/10.4137/BBI.S38316).
- [57] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. “Quantum Random Access Memory”. *Physical Review Letters* 100.16 (2008). DOI: [10.1103/physrevlett.100.160501](https://doi.org/10.1103/physrevlett.100.160501).
- [58] Gilles Brassard, Peter Høyer, Michele Mosca, and Alain Tapp. *Quantum amplitude amplification and estimation*. 2002. DOI: [10.1090/conm/305/05215](https://doi.org/10.1090/conm/305/05215).
- [59] Fabian Pedregosa et al. *Scikit-learn: Machine Learning in Python*. 2018.
- [60] Andrew Rosenberg and Julia Hirschberg. “V-measure: A conditional entropy-based external cluster evaluation measure”. *Proceedings of the 2007 joint conference on empirical methods in natural language processing and computational natural language learning (EMNLP-CoNLL)*. 2007, pp. 410–420.
- [61] Jekaterina Jaroslavceva. “A New Trackster Linking Algorithm Based on Graph Neural Networks for the CMS Experiment at the Large Hadron Collider at CERN”. Presented 14 Jul 2023. Prague, Tech. U., 2023.
- [62] Amine Zeguendry, Zahi Jarir, and Mohamed Quafafou. “Quantum Machine Learning: A Review and Case Studies”. en. *Entropy (Basel)* 25.2 (Feb. 2023).
- [63] Tobias Haug, Chris N Self, and M S Kim. “Quantum machine learning of large datasets using randomized measurements”. *Machine Learning: Science and Technology* 4.1 (2023), p. 015005. DOI: [10.1088/2632-2153/acb0b4](https://doi.org/10.1088/2632-2153/acb0b4).

- [64] Ilya Sinayskiy Maria Schuld and Francesco Petruccione. “An introduction to quantum machine learning”. *Contemporary Physics* 56.2 (2015), pp. 172–185. DOI: [10.1080/00107514.2014.964942](https://doi.org/10.1080/00107514.2014.964942).
- [65] Prayag Tiwari, Shahram Dehdashti, Abdul Karim Obeid, Massimo Melucci, and Peter Bruza. *Kernel Method based on Non-Linear Coherent State*. 2020.
- [66] Kazuhisa Fujita. “Approximate spectral clustering using both reference vectors and topology of the network generated by growing neural gas”. en. *PeerJ Comput Sci* 7 (Aug. 2021), e679.
- [67] Alessio Celi et al. “Emerging Two-Dimensional Gauge Theories in Rydberg Configurable Arrays”. *Phys. Rev. X* 10 (2 2020), p. 021057. DOI: [10.1103/PhysRevX.10.021057](https://doi.org/10.1103/PhysRevX.10.021057).
- [68] Hannes Bernien et al. “Probing many-body dynamics on a 51-atom quantum simulator”. *Nature* 551.7682 (2017), pp. 579–584. DOI: [10.1038/nature24622](https://doi.org/10.1038/nature24622).
- [69] Henning Labuhn et al. “Tunable two-dimensional arrays of single Rydberg atoms for realizing quantum Ising models”. *Nature* 534.7609 (2016), pp. 667–670. DOI: [10.1038/nature18274](https://doi.org/10.1038/nature18274).
- [70] Frank Arute et al. “Quantum supremacy using a programmable superconducting processor”. *Nature* 574.7779 (2019), pp. 505–510. DOI: [10.1038/s41586-019-1666-5](https://doi.org/10.1038/s41586-019-1666-5).
- [71] B. P. Lanyon et al. “Universal Digital Quantum Simulation with Trapped Ions”. *Science* 334.6052 (2011), pp. 57–61. DOI: [10.1126/science.1208001](https://doi.org/10.1126/science.1208001).
- [72] S. Debnath et al. “Demonstration of a small programmable quantum computer with atomic qubits”. *Nature* 536.7614 (2016), pp. 63–66. DOI: [10.1038/nature18648](https://doi.org/10.1038/nature18648).
- [73] A. D. Córcoles et al. “Demonstration of a quantum error detection code using a square lattice of four superconducting qubits”. *Nature Communications* 6.1 (2015), p. 6979. DOI: [10.1038/ncomms7979](https://doi.org/10.1038/ncomms7979).
- [74] Beatrice Nash, Vlad Gheorghiu, and Michele Mosca. “Quantum circuit optimizations for NISQ architectures”. *Quantum Science and Technology* 5.2 (2020), p. 025010. DOI: [10.1088/2058-9565/ab79b1](https://doi.org/10.1088/2058-9565/ab79b1).
- [75] Jia Wu et al. “Hyperparameter Optimization for Machine Learning Models Based on Bayesian Optimization”. *Journal of Electronic Science and Technology* 17.1 (2019), pp. 26–40. DOI: <https://doi.org/10.11989/JEST.1674-862X.80904120>.
- [76] Gorjan Alagic and Bill Fefferman. *On Quantum Obfuscation*. 2016.

- [77] Zvika Brakerski and Henry Yuen. *Quantum Garbled Circuits*. 2020.
- [78] IEEE, 2009. DOI: [10.1109/focs.2009.36](https://doi.org/10.1109/focs.2009.36).
- [79] James Bartusek, Fuyuki Kitagawa, Ryo Nishimaki, and Takashi Yamakawa. *Obfuscation of Pseudo-Deterministic Quantum Circuits*. 2023.
- [80] Sanjam Garg et al. “Candidate Indistinguishability Obfuscation and Functional Encryption for all Circuits”. *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*. 2013, pp. 40–49. DOI: [10.1109/FOCS.2013.13](https://doi.org/10.1109/FOCS.2013.13).
- [81] Richard Feynman. “Quantum mechanical computers”. *Optics news* 11.2 (1985), pp. 11–20.
- [82] Libor Caha, Zeph Landau, and Daniel Nagaj. “Clocks in Feynman’s computer and Kitaev’s local Hamiltonian: Bias, gaps, idling, and pulse tuning”. *Phys. Rev. A* 97 (6 2018), p. 062306. DOI: [10.1103/PhysRevA.97.062306](https://doi.org/10.1103/PhysRevA.97.062306).
- [83] Andrew M. Childs. “On the Relationship Between Continuous- and Discrete-Time Quantum Walk”. *Communications in Mathematical Physics* 294.2 (Oct. 2009), 581–603. DOI: [10.1007/s00220-009-0930-1](https://doi.org/10.1007/s00220-009-0930-1).
- [84] Daniel Reitzner, Daniel Nagaj, and Vladimír Bužek. “Quantum Walks”. *Acta Physica Slovaca. Reviews and Tutorials* 61.6 (Dec. 2011). DOI: [10.2478/v10155-011-0006-6](https://doi.org/10.2478/v10155-011-0006-6).
- [85] John M. Martyn, Zane M. Rossi, Andrew K. Tan, and Isaac L. Chuang. “Grand Unification of Quantum Algorithms”. *PRX Quantum* 2.4 (Dec. 2021). DOI: [10.1103/prxquantum.2.040203](https://doi.org/10.1103/prxquantum.2.040203).
- [86] Dorit Aharonov et al. *Adiabatic Quantum Computation is Equivalent to Standard Quantum Computation*. 2005.
- [87] Alexei Y. Kitaev, A. H. Shen, and Mikhail N. Vyalyi. “Classical and Quantum Computation”. *Graduate Studies in Mathematics*. 2002.
- [88] P. Jordan and E. Wigner. “Über das Paulische Äquivalenzverbot”. *Zeitschrift für Physik* 47.9 (1928), pp. 631–651. DOI: [10.1007/BF01331938](https://doi.org/10.1007/BF01331938).
- [89] Sergey B. Bravyi and Alexei Yu. Kitaev. “Fermionic Quantum Computation”. *Annals of Physics* 298.1 (2002), pp. 210–226. DOI: <https://doi.org/10.1006/aphy.2002.6254>.
- [90] M. Cerezo, Kunal Sharma, Andrew Arrasmith, and Patrick J. Coles. “Variational quantum state eigensolver”. *npj Quantum Information* 8.1 (2022), p. 113. DOI: [10.1038/s41534-022-00611-6](https://doi.org/10.1038/s41534-022-00611-6).

- [91] Hasan Sayginel, Francois Jamet, Abhishek Agarwal, Dan E. Browne, and Ivan Rungger. *A fault-tolerant variational quantum algorithm with limited T-depth*. 2023.
- [92] Ben W. Reichardt. “The quantum adiabatic optimization algorithm and local minima”. *Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing*. STOC ’04. Chicago, IL, USA: Association for Computing Machinery, 2004, 502–510. DOI: [10.1145/1007352.1007428](https://doi.org/10.1145/1007352.1007428).
- [93] G Nenciu. “On the adiabatic theorem of quantum mechanics”. *Journal of Physics A: Mathematical and General* 13.2 (1980), p. L15. DOI: [10.1088/0305-4470/13/2/002](https://doi.org/10.1088/0305-4470/13/2/002).
- [94] Albert Messiah. *Quantum Mechanics John Wiley & Sons*. 1968.
- [95] Julia Kempe, Alexei Kitaev, and Oded Regev. *The Complexity of the Local Hamiltonian Problem*. 2005.
- [96] Julia Kempe and Oded Regev. *3-Local Hamiltonian is QMA-complete*. 2003.
- [97] Jules Tilly et al. “The Variational Quantum Eigensolver: A review of methods and best practices”. *Physics Reports* 986 (2022). The Variational Quantum Eigensolver: a review of methods and best practices, pp. 1–128. DOI: <https://doi.org/10.1016/j.physrep.2022.08.003>.
- [98] Alberto Peruzzo et al. “A variational eigenvalue solver on a photonic quantum processor”. *Nature Communications* 5.1 (2014), p. 4213. DOI: [10.1038/ncomms5213](https://doi.org/10.1038/ncomms5213).
- [99] Bozhi Wang et al. *A powered full quantum eigensolver for energy band structures*. 2023.
- [100] Shijie Wei, Hang Li, and GuiLu Long. *A Full Quantum Eigensolver for Quantum Chemistry Simulations*. 2020.
- [101] John M. Martyn, Zane M. Rossi, Andrew K. Tan, and Isaac L. Chuang. “Grand Unification of Quantum Algorithms”. *PRX Quantum* 2 (4 2021), p. 040203. DOI: [10.1103/PRXQuantum.2.040203](https://doi.org/10.1103/PRXQuantum.2.040203).
- [102] Lin Lin and Yu Tong. “Near-optimal ground state preparation”. *Quantum* 4 (2020), p. 372. DOI: [10.22331/q-2020-12-14-372](https://doi.org/10.22331/q-2020-12-14-372).
- [103] Alan Aspuru-Guzik, Anthony D. Dutoi, Peter J. Love, and Martin Head-Gordon. “Simulated Quantum Computation of Molecular Energies”. *Science* 309.5741 (2005), pp. 1704–1707. DOI: [10.1126/science.1113479](https://doi.org/10.1126/science.1113479).
- [104] Guoming Wang, Daniel Stilck França, Ruizhe Zhang, Shuchen Zhu, and Peter D. Johnson. “Quantum algorithm for ground state energy estimation using circuit depth with exponentially improved dependence on precision” (2023).

- [105] Raphael Seidel, Nikolay Tcholchev, Sebastian Bock, Colin Kai-Uwe Becker, and Manfred Hauswirth. *Efficient Floating Point Arithmetic for Quantum Computers*. 2021.

Appendix A

APPENDICES

A.1 Grover Oracles for qLUE

In this appendix, we describe the oracles U_P (recall the working of Grover Search described in Sec. 1.1) needed for the various subroutines of qLUE.

In the context of Local Density Computation, $f(x) = 1$ for a point if it lies in the critical radius d_c . If this condition is not satisfied, $f(x) = 0$.

To implement the operators in Eqs. (1.1), we require A(dd) and M(ultiply) circuits. We use the ones introduced in Ref. [105], which perform the following operations

$$A |X_i\rangle |0\rangle = |X_i\rangle |x_{1_i} + x_{2_i}\rangle, \quad (\text{A.1a})$$

$$M |X_i\rangle |0\rangle = |X_i\rangle |x_{1_i} \cdot x_{2_i}\rangle. \quad (\text{A.1b})$$

For the local density step, the quantum circuit U_P to implement the search function f is given in Fig. A.1. The overall idea is to compute the Euclidean distance between every input point and the base point and check if this distance is higher than d_c . The $|sign\rangle$ qubit stores the output of this computation. The F gates are X_{sign} gates which are X on the sign qubit and act as the identity on every other qubit, such that $F|x_{1_i}\rangle = |-x_{1_i}\rangle$. These are used such that the first and second levels of Add gates A compute $x_{i_1} - x_{i_2}$ into $|a_i\rangle$ and $|b_i\rangle$ respectively. The M multiply gates then set the $|m_i\rangle$ states to $(x_{i_1} - x_{i_2})^2$ taking a_i and b_i as inputs. An A gate next acts on m_1 and m_2 to set $|ans\rangle$ to $(x_{1_1} - x_{1_2})^2 + (x_{2_1} - x_{2_2})^2$. The A_{sign} gate is a subcircuit of the addition circuit A that finally computes the sign of $(x_{1_1} - x_{1_2})^2 + (x_{2_1} - x_{2_2})^2 - d_c^2$ and stores it in $|sign\rangle$. Thus, $|sign\rangle = |f(x)\rangle$ [with $f(x)$ as in Eq. (1.1b)] for the Local Density computation step.




Figure A.1: Circuit corresponding to the marking operator U_P in Eq. (1.1b). First, it computes the Euclidean distance $d_{1,2}$ between points 1 and 2, as defined in Sec. 2.2. Then, it compares $d_{1,2}$ to the threshold d_c and marks the index based on the sign qubit $|sign\rangle$. A and B are defined in Eq. (A.1) and in Ref. [105]. Bars before gates are employed to indicate that the corresponding qubits are unaffected. In the figure, the first subscript on theregister contains the spatial coordinate and the second subscript contains the index. So, for example, $|x_{1_i}\rangle$ contains the x_1 coordinate of the i^{th} point. The flip gates F flip the sign bit of the input. The ancillas a_1 and b_1 contain the output of computation $x_{1_i} - x_{1_j}$, while a_2 and b_2 contain the output of computation $x_{2_i} - x_{2_j}$, after the first two levels of A gates. The M gates then set m_1 and m_2 to $(x_{1_i} - x_{1_j})^2$ and $(x_{2_i} - x_{2_j})^2$ respectively. The pre-final A gate sets the $|ans\rangle$ state to $(x_{1_i} - x_{1_j})^2 + (x_{2_i} - x_{2_j})^2$. The final A_{sign} gate computes $(x_{1_i} - x_{1_j})^2 + (x_{2_i} - x_{2_j})^2 - d_c^2$. However, for this final A_{sign} gate, we need only compute the sign qubit of the computation and so only the sign qubit computation subcircuit of A features in A_{sign} .

For the Nearest Higher Procedure, a similar circuit can be used with additional registers $|\rho_i\rangle$ for the Local Density and $|\tilde{\rho}\rangle$ for the critical density threshold $\tilde{\rho}$. d_c is replaced by d_L and the signs of $|d_{i,j}^2\rangle - |ans\rangle$ and $\rho_i - \tilde{\rho}$ are additionally computed in order to find only the points that lie between $d_{i,j}$ and d_L and satisfy the $\tilde{\rho}$ threshold. These signs can be

combined with the Toffoli gate to compute the required function $f(x)$

For Seeds and Outliers, the condition is similar to that in the Nearest Higher (GEBS) procedure in that the distance criterion is only a lower bound as opposed to a window. The GEBS blackbox without the upper distance bound computation can be used.

For the Cluster Assignment step, the operator U_P needs to check if the Nearest Higher index of the points over which Grover Search is carried out lies within $\mathcal{C}$. The U_P operator for this can be generated using X and C^nZ gates. C^nZ gates map $|1\rangle^{\otimes n}$ to $-|1\rangle^{\otimes n}$ and act as the identity on all the other elements of the computational basis. Let us consider a simple example where the dynamic search space DSS has 8 points indexed from 000 to 111 in binary. So, the Nearest Higher index will always be in this range. If $\mathcal{C}$ consists of indices 0(000) and 5(101), we would use U_P as defined in Fig. A.2. This flips the phase of 000 and 101 in the input superposition as required (following Eq. (A.1b)). As shown in Fig. A.2, operators for each index can be sequentially applied to check for multiple indices in $\mathcal{C}$.




Figure A.2: Circuit for U_P to check if the Nearest Higher indices are in $\mathcal{C}$, which here has 2 elements: 5(101) and 0(000). In the above circuit, the first part of the circuit flips the phase of the state if the register is in the $|101\rangle$ state that corresponds to index 5. The second part flips the phase of the state if the register is in the $|000\rangle$ state that corresponds to index 0. Note that the X gates performed at the end of the two subcircuits are uncomputation steps that ensure that the input to the subsequent subcircuits are as required and unaffected by our initial X gates.