

# Quantum approximate optimization for hard problems in linear algebra

Ajinkya Borle<sup>1\*</sup>, Vincent E. Elfving<sup>2</sup> and Samuel J. Lomonaco<sup>1</sup>

<sup>1</sup> CSEE Department, University of Maryland, Baltimore County, Baltimore, USA

<sup>2</sup> Qu & Co BV, PO Box 75872, 1070 AW, Amsterdam, the Netherlands

\* [aborle1@umbc.edu](mailto:aborle1@umbc.edu)

## Abstract

The quantum approximate optimization algorithm (QAOA) by Farhi et al. is a quantum computational framework for solving quantum or classical optimization tasks. Here, we explore using QAOA for binary linear least squares (BLLS); a problem that can serve as a building block of several other hard problems in linear algebra, such as the non-negative binary matrix factorization (NBMF) and other variants of the non-negative matrix factorization (NMF) problem. Most of the previous efforts in quantum computing for solving these problems were done using the quantum annealing paradigm. For the scope of this work, our experiments were done on noiseless quantum simulators, a simulator including a device-realistic noise-model, and two IBM Q 5-qubit machines. We highlight the possibilities of using QAOA and QAOA-like variational algorithms for solving such problems, where trial solutions can be obtained directly as samples, rather than being amplitude-encoded in the quantum wavefunction. Our numerics show that even for a small number of steps, simulated annealing can outperform QAOA for BLLS at a QAOA depth of  $p \leq 3$  for the probability of sampling the ground state. Finally, we point out some of the challenges involved in current-day experimental implementations of this technique on cloud-based quantum computers.



Copyright A. Borle *et al.*

This work is licensed under the Creative Commons

[Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

Published by the SciPost Foundation.

Received 04-06-2021

Accepted 20-10-2021

Published 30-11-2021

doi:[10.21468/SciPostPhysCore.4.4.031](https://doi.org/10.21468/SciPostPhysCore.4.4.031)



Check for  
updates

## Contents

<b>1</b>	<b>Introduction</b>	<b>2</b>
<b>2</b>	<b>Background and related work</b>	<b>3</b>
2.1	Background	3
2.1.1	The binary linear least squares (BLLS) problem	3
2.1.2	Non-negative matrix factorization (NMF) and variants	4
2.1.3	Quadratic unconstrained binary optimization (QUBO)	5
2.1.4	The quantum approximate optimization algorithm (QAOA)	5
2.1.5	Implicit filtering optimization	7
2.2	Related work	7
<b>3</b>	<b>QAOA for BLLS</b>	<b>9</b>

3.1	Problem formulation	9
3.2	Mapping to quantum gates	9
3.2.1	For IBM Q specific gates	10
3.3	Experimental methods	11
3.3.1	Experiments with no noise	11
3.3.2	Experiments to compare no noise and shot-noise performance	11
3.3.3	Experiments on an IBM Q device	12
3.4	Results	12
3.4.1	Optimization trajectory for QAOA	13
3.4.2	QAOA results with no noise	14
3.4.3	No noise vs shot-noise optimization	14
3.4.4	Probability of sampling the ground state	16
3.4.5	Effect of noise on expectation value	17
3.4.6	Effect of shot number on optimization	20
3.4.7	IBM Q device performance	21
3.5	Discussion	21
<b>4</b>	<b>Conclusion</b>	<b>22</b>
<b>A</b>	<b>Detailed QUBO formulation for binary linear least squares</b>	<b>23</b>
<b>B</b>	<b>Example QAOA circuit for BLLS</b>	<b>23</b>
<b>C</b>	<b>Implementing two-qubit interactions on a QPU</b>	<b>24</b>
<b>D</b>	<b>Probability distribution of running a QAOA circuit for BLLS on a real QPU</b>	<b>25</b>
<b>E</b>	<b>Curve fit and simulated annealing details</b>	<b>26</b>
E.1	Curve fit	26
E.2	Simulated annealing experiments	27
	<b>References</b>	<b>27</b>

---

## 1 Introduction

The application of quantum computing to hard optimization problems is a candidate where quantum computing may eventually outperform classical computation [1–6]. At the time of writing this paper, noisy intermediate scale quantum (NISQ) computers [5] are being developed by several firms and research groups [7–13]. The two main approaches to quantum optimization are (i) the quantum annealing (QA) physical heuristic [6] and (ii) quantum approximate optimization algorithm (QAOA) [4] on the gate-model quantum computer [1].

In this paper, we explore and propose the use of QAOA for hard problems in linear algebra. In particular, we focus on the problem of binary linear least squares (BLLS) [14–16], a subroutine for solving certain types of non-negative matrix factorization (NMF) problems [17, 18], for example, the problem of non-negative binary matrix factorization (NBMF) [17, 19]. These are linear dimensionality reduction (LDR) tools useful in machine learning for learning a basis that would be used to model data [20], such as learning facial features in computer vision [21], topic modeling in documents [22] and properties of astronomical objects [23, 24]. Further-

more, BLS and its variant (binary compressive sensing [25]) have applications in signal processing [16, 25], image processing [15, 26] and can be a building block for other problems in linear algebra [27–32]. We hope that our work provides insights to fellow researchers to further explore the use of NISQ era methods [4, 33] for problems in linear algebra and numerical computation. The outline of our results are as follows:

- The technique seems to scale well in optimizing the expectation value as the problem size  $n$  increases (for a fixed QAOA circuit depth  $p$ ). This corresponds to getting ‘good solutions’ fast but not necessarily the most optimal one (refer Section 3.4.2).
- The probability of finding the optimal solution appears to drop exponentially as the problem size increases (for a fixed QAOA circuit depth  $p$ ). QAOA outperformed random sampling, but was beaten by simulated annealing. Further and extensive tests would be required to show if QAOA can deliver a quantum advantage for this criterion (refer Section 3.4.4).
- The implicit filtering (ImFil) classical optimizer backend [34] used for optimizing the QAOA parameters is robust to shot-noise (refer Sections 3.4.1, 3.4.3 and 3.4.6).
- Given a specific realistic quantum hardware noise model, a dense device coupling map would drastically reduce the errors in approximating the expectation value for this problem, when compared against a sparse one (refer Section 3.4.5).
- When compared against older quantum devices, the newer ones seem to be getting better at approximating the expectation value for a given set of QAOA parameters. However, the results they produce are still very noisy for QAOA to function effectively (refer Sections 3.4.1 and 3.4.7).

In Section 2, we cover the necessary background and related work for our paper. Section 3.1 is about formulating the BLS problem into a classical Ising problem and Section 3.2 goes over the details of the gates required for the QAOA ansatz. The experiments, results and discussion are detailed in Sections 3.3, 3.4 and 3.5 respectively. We finally conclude our paper in Section 4. We also have Appendices to complement and support the information in the main paper.

## 2 Background and related work

### 2.1 Background

#### 2.1.1 The binary linear least squares (BLS) problem

Given a matrix  $A \in \mathbb{R}^{m \times n}$ , an unknown column vector of variables  $x \in \{0, 1\}^n$  and a column vector  $b \in \mathbb{R}^m$  (Where  $m > n$ ). The linear BLS problem is to find the  $x$  that would minimize  $\|Ax - b\|$  the most. In other words, it can be described as:

$$\arg \min_x \|Ax - b\|. \quad (1)$$

The motivation behind choosing the BLS problem as a target problem is twofold: firstly, it is an NP-Hard problem [16] that makes it a relevant target for potential speedup. Secondly, it can act as a building block for other hard problems in linear algebra, such as the Non-negative Binary Matrix Factorization [17]. One reason why one may view BLS a building block for other problems is because multiple binary variables can be clubbed together for a fixed point

approximation of a real variable [18, 27–30, 35]. Amongst these, there are some problems that are NP-hard for which an approximate solution would be acceptable [18, 35]. In these cases, QAOA may be able to provide an improvement in finding approximate solutions against classical solvers and increase the probability of sampling the best solution.

### 2.1.2 Non-negative matrix factorization (NMF) and variants

NMF is a linear dimensionality reduction (LDR) tool that helps in data analysis. Given a matrix  $V \in \mathbb{R}_{\geq 0}^{m \times n}$  of  $n$  data points of dimension  $m$ , NMF extracts a basis  $W \in \mathbb{R}_{\geq 0}^{m \times r}$  containing  $r$  elements such that the linear space spanned by the basis  $W$  approximates the datapoints in  $V$  as best as they can [20]. In other words, for some matrix  $H \in \mathbb{R}_{\geq 0}^{r \times n}$ ,  $V \approx W \times H$ . The restrictions to non-negative number makes the factored matrices sparse and easily interpretable for certain domains where non-negative feature vectors do not apply [20]. Few such examples are:

**Computer vision:** NMF is useful for facial recognition as the problem suffers from high dimensionality of data. Here, NMF can extract a set of features  $W$  from the input/training data (matrix  $V$ ). Then for an image unseen by the model before (test), represented by column vector  $v^*$ , we'll use  $W$  and see if we can find some non-negative column vector  $h^*$  such that  $v^* \approx W \times h^*$ . NMF is found to be more robust than another LDR technique, Principal Component Analysis or PCA (which doesn't restrict the data to be non-negative), for image occlusions [21].

**Astronomy:** NMF has applications in astronomy primarily due to astrophysical signals being non-negative. As an example, NMF was used on spectroscopic observations [23] and the direct imaging observations [24] in order to study the common properties of astronomical objects and post-process astronomical observations.

**Text mining/Topic modeling:** Here, text documents can be converted into non-negative column vectors. One way to do this is by considering a dictionary of  $m$  words. Each matrix element  $V_{ij}$  in the dataset  $V$  would denote the number of occurrences of the  $i$ th word in the  $j$ th document. Applying NMF, a basis of topics  $W$  can be generated along with the matrix  $H$  that suggests the importance of each of the basis topics for each document (i.e. column) in  $V$ . This is useful to automate the extraction of abstract topics in a given set of documents [22].

The downside of NMF is that the factorization is NP-Hard in the general case [36].

**Non-negative binary matrix factorization (NBMF)** is a specialized version of the NMF problem; where entries in matrix  $H$  are restricted as  $H \in \{0, 1\}^{r \times n}$ . Here the matrix  $H$  is even sparser than in the general NMF [17]. The features learnt by the matrix  $W$  would be more holistic against those in NMF. As an example, In facial feature extraction, NMF is able to learn

---

**Algorithm 1** Alternating least squares for NBMF (pseudocode).

---

```

1: procedure MAIN( $V$ ) ▷  $V$  is the matrix to be factorized
2:   Randomly initialize the matrix  $H \in \{0, 1\}^{r \times n}$ 
3:   while not converged do
4:     for row  $i$  from 1 to  $m$  do
5:        $(W_i)^T \leftarrow \arg \min_{W_i^T} \|V_i^T - H^T(W_i)^T\|_2$  such that  $W \in \mathbb{R}_{\geq 0}^{m \times r}$ 
6:     end for
7:     for column  $j$  from 1 to  $n$  do
8:        $H_j \leftarrow \arg \min_{H_j} \|V_j - WH_j\|_2$  such that  $H \in \{0, 1\}^{r \times n}$ 
9:     end for
10:  end while
11:  return  $W, H$ 
12: end procedure

```

---

parts of faces from a given dataset, whereas the features extracted by NBMF are closer to a complete face [17].

BLLS can be used in order to solve the NBMF variant by using the alternating least squares method [37].

In Algorithm 1, line 5 is solved classically since efficient algorithms exist for it [38], and would probably be applied with some kind of regularization since NMF and its variants tend to generally be ill-posed problems [17, 39]. It is line 8 that is solved using BLLS (where QAOA would be applied).

In the past, quantum annealing was used as a subroutine within this algorithm to solve NBMF and other NMF related problems [17, 18]. Based on our work with this paper, QAOA can be an alternative to quantum annealing for NBMF, which can be explored in the future.

### 2.1.3 Quadratic unconstrained binary optimization (QUBO)

The QUBO objective function is as follows,

$$F(q) = \sum_a v_a q_a + \sum_{a < b} w_{ab} q_a q_b, \tag{2}$$

where  $q_a \in \{0, 1\}$ ,  $v_a$  and  $w_{ab}$  are real coefficients for the linear and quadratic parts of the function respectively. The QUBO objective function is NP-hard in nature [40]. One salient feature of this objective function is that many application domain problems map naturally to QUBO [17, 18, 27, 41, 42]. In the process of applying BLLS to gate model quantum devices, we use the QUBO formulation as an intermediate stage of expressing the problem.

### 2.1.4 The quantum approximate optimization algorithm (QAOA)

In 2014 Farhi et al. proposed an algorithm that uses both quantum and classical computation for solving optimization problems [4]. The potential advantage of using this algorithm is that it can be implemented by using low depth quantum circuits [43], making it suitable for NISQ devices. We here briefly summarize the QAOA formalism applied to binary optimization problems. For the required preliminaries of quantum computing, the authors recommend the textbook by Nielsen and Chuang [1].

One popular method of encoding an optimization problem to be solved using QAOA, is to first formulate the problem as an Ising Objective function.

$$F(\sigma) = \sum_a h_a \sigma_a + \sum_{a < b} J_{ab} \sigma_a \sigma_b, \tag{3}$$

$$\text{where } \sigma_a = 2q_a - 1. \tag{4}$$

Here  $\sigma_a \in \{-1, 1\}$ ,  $h$  and  $J$  are coefficients associated with individual and coupled binary variables respectively. The Ising model is a popular statistical mechanics model, associated primarily with ferromagnetism [44]. Because it has been shown to be NP-Complete in nature [45], the objective function associated with it can be used to represent hard problems [46]. It is important to note here that we still don't know if quantum computing can help solve NP-Complete problems efficiently [47]. Our hope for quantum algorithms, at the very least, is to be able to compete with classical heuristics when it comes to certain classes of hard problems.

The problem then would be to maximize or minimize Eqn(3), depending on how it is set up. The (classical) Ising Hamiltonian, which naturally maps the Ising objective Eqn(3) to

---

**Algorithm 2** Quantum approximate optimization algorithm for minimization (pseudocode).

---

```

1: procedure MAIN( $\hat{B}, \hat{C}, p$ ) ▷ The main routine of the algorithm
2:    $\beta \leftarrow \{\emptyset\}, \gamma \leftarrow \{\emptyset\}, \text{expt\_val} \leftarrow \emptyset, \text{best\_res} \leftarrow \emptyset$ 
3:   Pick at random  $\beta \in [0, \pi]^p, \gamma \in [0, 2\pi]^p$ 
4:   while  $(\beta, \gamma)$  can be further optimized, or a limit is reached do
5:     Initialize  $\text{res\_set} \leftarrow \{\emptyset\}$ 
6:     for a fixed number of shots do
7:        $\text{res\_set} \leftarrow \text{res\_set} \cup \text{QAOA}(\hat{B}, \hat{C}, \beta, \gamma, p)$ 
8:     end for
9:     From  $\text{res\_set}$ , calculate the expectation value and store in  $\text{expt\_val}$ 
10:    Based on the  $\text{expt\_val}$ , pick new  $2p$  angles  $(\beta, \gamma)$  by classical optimization
11:  end while
12:  From the final  $\text{res\_set}$ , set the result with lowest energy,  $\text{best\_res} \leftarrow \min(\text{res\_set})$ 
13:  return  $\text{best\_res}$ 
14: end procedure
15: procedure QAOA( $\hat{B}, \hat{C}, \beta, \gamma, t$ )
16:  Initialize  $n$  qubits,  $|\psi\rangle \leftarrow |0\rangle^{\otimes n}$ 
17:  Apply Hadamard transform,  $|\psi\rangle = 1/\sqrt{2^n}(|0\rangle + |1\rangle)^{\otimes n}$ 
18:   $j \leftarrow 1$ 
19:  while  $j \leq t$  do
20:     $|\psi\rangle \leftarrow e^{-i\gamma_j \hat{C}} |\psi\rangle$ 
21:     $|\psi\rangle \leftarrow e^{-i\beta_j \hat{B}} |\psi\rangle$ 
22:     $j \leftarrow j + 1$ 
23:  end while
24:  Measure  $|\psi\rangle$  in standard basis and store in a classical register  $o$ 
25:  return  $o$ 
26: end procedure

```

---

qubits, can be expressed as:

$$\hat{C} = \sum_a h_a \hat{\sigma}_a^{(z)} + \sum_{a < b} J_{ab} \hat{\sigma}_a^{(z)} \hat{\sigma}_b^{(z)}, \quad (5)$$

$$\text{where } \hat{\sigma}_a^{(z)} = (\otimes_{i=1}^{a-1} \hat{I}) \otimes (\hat{\sigma}^{(z)}) \otimes (\otimes_{i=a+1}^n \hat{I}) \quad (6)$$

$$\text{and } \hat{\sigma}^{(z)} = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \quad (7)$$

Here, indices  $a, b, i$  label the qubits,  $n$  is the total number of qubits,  $\hat{\sigma}^{(z)}$  is the Pauli Z operator and  $I$  is the identity operator. The other type of Hamiltonian in the QAOA process is a summation of individual Pauli X operators for each qubit involved in the process, which intuitively represents a transverse field in the Ising model:

$$\hat{B} = \sum_a \hat{\sigma}_a^{(x)}, \quad (8)$$

$$\text{where } \hat{\sigma}_a^{(x)} = (\otimes_{i=1}^{a-1} \hat{I}) \otimes (\hat{\sigma}^{(x)}) \otimes (\otimes_{i=a+1}^n \hat{I}) \quad (9)$$

$$\text{and } \hat{\sigma}^{(x)} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}. \quad (10)$$

In QAOA, the qubits are first put in a uniform superposition over the computational basis states by applying a Hadamard gate, which maps  $|0\rangle \rightarrow (|0\rangle + |1\rangle)/\sqrt{2}$ , on every qubit. Then, the

Hamiltonian pair  $\hat{C}$  and  $\hat{B}$  is applied  $p$  number of times using a set of angles  $\gamma$  and  $\beta$ , where, for  $1 \leq l \leq p$ , each  $\gamma_l \in [0, 2\pi]$  and  $\beta_l \in [0, \pi]$ <sup>1</sup> [4]. The expectation value of the Hamiltonian  $\hat{C}$  with respect to the resultant state  $|\psi(\gamma, \beta)\rangle$  is calculated as

$$C(\gamma, \beta) = \langle \psi(\gamma, \beta) | \hat{C} | \psi(\gamma, \beta) \rangle. \quad (11)$$

A classical black-box optimizer then uses the expectation value as its input and suggests new  $\gamma$  and  $\beta$  sets (of length  $p$  each). The hope is that as the number of qubits (more specifically, variables)  $n$  involved in the optimization increases, if for circuit depth  $p \ll n$  we are able to efficiently sample the best solution, we would have an advantage in using QAOA over classical methods. Although Algorithm 2 is a summary of the QAOA method, we recommend readers the original paper [4] for further details.

### 2.1.5 Implicit filtering optimization

As mentioned before, QAOA requires us to give it the sets of angles  $\gamma$  and  $\beta$  in order to manipulate the state of the quantum system. The most common way to do this is to use classical black-box optimization techniques that do not need the derivative information of the problem [9, 48, 49]. Since the expectation value  $C(\gamma, \beta)$  of the objective function cost (or energy) would be approximate in nature, we need an optimization technique that can handle noisy data. The technique of our choice for this work is the implicit filtering algorithm [34].

In essence, implicit filtering or ImFil is a derivative-free, bounded black-box optimization technique that accommodates noise when it tries to suggest the best parameters to minimize the objective function. Various other techniques for noisy optimization exist, such as Bayesian Optimization [50], COMPASS [51], SPSA [52], etc. However, we found implicit filtering the best for our current efforts. For further details, we recommend the book by C.T Kelly on the topic [34].

## 2.2 Related work

One of the first applications of quantum computing for solving problems in the field of linear algebra is the HHL algorithm (named after its creators: Harrow, Hassidim and Lloyd) for solving a system of linear equations [53]. This was followed by works for solving linear least squares [54], preconditioned system of linear equations [55], recommendation systems [56] and many others [57–61]. Although the classical counterparts of the above mentioned algorithms run in polynomial time, the quantum algorithms mentioned above run in the polylog time complexity (provided the problem is well conditioned).

However, there are some caveats with such kind of algorithms [62]. Among the many caveats, we'd like to emphasize on the two that affect the practicality of their utility in the near future. Firstly, they require fault tolerant quantum computers whereas, at the time of writing this paper, we have just entered the NISQ era [5]. Secondly, for the algorithms focused on linear system of equations [53, 55, 57] and least squares [54, 58], the output data is encoded as a normalized vector of a quantum state  $|x\rangle$  (which means that the probability amplitudes of the basis states encode the data). This means that we need an efficient method to prepare the input data as a quantum state [62–65]; and the output will be a quantum state as well, which means it wouldn't be available for us in the classical world directly by performing measurement in the standard computational basis. This can be mitigated by either measuring the final state in a basis of our choice if our goal is to know some statistical information about  $x$  [53, 66] or learning certain values in  $x$  (though that will eliminate the exponential speedup [62]).

<sup>1</sup>This is true as long as all possible variable combinations have obj. function costs that are  $\geq 1$  in magnitude.

With respect to quantum annealing, O'Malley and Vesselinov's paper in 2016 [27] was one of the first that proposed to solve linear least squares. Other works in this domain were for solving specific NMF problems [17–19], polynomial system of equations [28], underdetermined binary linear systems [67] and polynomial least squares [30]. It's hard to speculate about speedups analytically with (i) D-wave's noisy implementation of quantum annealing [68] and (ii) the problem of exponential gap-closing between the problem Hamiltonian's ground state and its excited states [69]. In the work by O'Malley and Vesselinov [17], they used a time to target benchmark in which classical solvers (Tabu search [70] and Gurobi [71] in their comparison) have to match or find better solutions than the ones returned by quantum annealing (not necessarily the optimal solution) in the same amount of time. The D-wave quantum annealer was able to beat those classical solvers for the benchmark, but the authors also mention that a combination of the two classical techniques would probably perform better than the D-wave by compensating for each other's shortcomings. The other important result in subsequent papers [18, 19] was to show that combining reverse and forward annealing improved results over just using forward annealing for most cases. Golden and O'Malley saw an improvement of 12% over forward annealing [19], but that came at the price of having at least 7 reverse annealing runs per QUBO (which was reported to have the quantum runtime of 29 forward anneals). It is important to note that certain quantum inspired algorithms may perform just as well or better than quantum annealers for such highly dense problems in terms of variable interactions [72]. The above mentioned quantum annealing techniques use the classical Ising model's diagonal Hamiltonian for problem formulation. This means that measuring the post annealing quantum state in computational basis gives us a bitstring which directly encodes the vector  $x$  (which we hope to be the best solution to Eqn(1)), unlike a lot of gate-model algorithms like the ones mentioned above [53–55, 57, 58] that encode the normalized solution vector in the amplitudes of  $|x\rangle$ . The exception to this rule is the work by An and Lin, where they explore using adiabatic quantum computing, not based on the Ising model, to produce amplitude encoded results for linear system of equations [73].

NISQ-compatible algorithms for efficiently solving linear algebra problems are highly desirable as of the time of writing this paper. The work by Chen et. al [74] proposes a hybrid algorithm that uses quantum random walks for solving a particular type of linear system, producing a classical result in  $O(n \log n)$ . However, the closest related works to ours are the recent papers that employ variational algorithms [73, 75–77]. The major difference however, is that, in those papers: i) The output is encoded as the vector of probability amplitudes of the quantum state  $|x\rangle$  and ii) The problems explored thus far are convex in nature and solved in polynomial time classically.

We in this paper implement QAOA on similar problems which were implemented on D-wave's quantum annealer previously, and therefore briefly mention a comparison here. The standard QAOA circuit strategy can be seen as similar to a bang-bang quantum annealing schedule, where cost and driver Hamiltonians are alternated. The quantum alternating operator ansatz extends this with more general operators [78] than available to Ising Hamiltonian annealers. Furthermore, QAOA is a gate-based quantum computational algorithm, a type of framework which promises universal programmability in terms of mapping an arbitrary problem graph to a qubit layout, even if the latter is not all-to-all connected. Conversely, in quantum annealing architectures mapping the logical problem qubits to a graph of physical hardware qubits can be a significant challenge in the general case [79, 80]. Our work is certainly not the first in applying QAOA to various relevant computational problems, and we refer the reader to a small list of examples [81–83]; in this work, we make an attempt to highlight some of the salient features and challenges of QAOA in the context of problems applicable to linear algebra and numerical analysis.

### 3 QAOA for BLLS

Before we go deeper, we here set the context of how QAOA will be used in this work. Rather than treating QAOA as an approximation algorithm with theoretical guarantees for the quality of solution obtained, it is used as a heuristic supported by empirical results.

#### 3.1 Problem formulation

O'Malley and Vesselinov first gave a QUBO formulation for the BLLS problem [27]. The details of how that is done is in Appendix A. Referring back to Eqn(1), if  $A \in \mathbb{R}^{m \times n}$ ,  $b \in \mathbb{R}^m$  and  $x \in \{0, 1\}^n$ , we can refine Eqn(2) to be

$$F(x) = \sum_j v_j x_j + \sum_{j < k} w_{jk} x_j x_k, \tag{12}$$

$$\text{where } v_j = \sum_i A_{ij}(A_{ij} - 2b_i) \tag{13}$$

$$\text{and } w_{jk} = 2 \sum_i A_{ij} A_{ik}. \tag{14}$$

Which means that the number of qubits depends only upon the size of the column vector  $x$ . All the rows in Matrix  $A$  and vector  $b$  are preprocessed classically in order to produce the coefficients of the QUBO problem.

By the equivalence stated in Eqn(4), we can then convert the problem into an Ising objective function (plus an offset value, irrelevant for optimization)

$$F(\sigma) = \sum_j h_j \sigma_j + \sum_{j < k} J_{jk} \sigma_j \sigma_k + \text{offset}, \tag{15}$$

$$\text{where } \sigma_j = 2x_j - 1. \tag{16}$$

#### 3.2 Mapping to quantum gates

Using the  $h, J$  coefficients from Eqn(15) along with the mapping to a classical Ising Hamiltonian given in Eqn(6) we get:

$$\hat{C} = \sum_j h_j \hat{\sigma}_j^{(z)} + \sum_{j < k} J_{jk} \hat{\sigma}_j^{(z)} \hat{\sigma}_k^{(z)}. \tag{17}$$

Because the individual components of Eqn(17) commute, we can express the Hamiltonian simulation of  $\hat{C}$  with an angle  $\gamma_l$  as follows

$$e^{-i\gamma_l \hat{C}} = \prod_j e^{(-ih_j \gamma_l) \hat{\sigma}_j^{(z)}} \prod_{j < k} e^{(-iJ_{jk} \gamma_l) \hat{\sigma}_j^{(z)} \hat{\sigma}_k^{(z)}}. \tag{18}$$

Similarly, the exponential of hamiltonian  $B$  can be broken down as

$$e^{-i\beta_l \hat{B}} = \prod_j e^{(-i\beta_l) \hat{\sigma}_j^{(x)}}. \tag{19}$$

In order to realize Eqn(18) and Eqn(19), we use the following gates

$$R_x(\omega) = e^{-i\frac{\omega}{2}\hat{\sigma}^{(x)}} = \begin{pmatrix} \cos \omega/2 & -i \sin \omega/2 \\ -i \sin \omega/2 & \cos \omega/2 \end{pmatrix}, \quad (20)$$

$$R_z(\omega) = e^{-i\frac{\omega}{2}\hat{\sigma}^{(z)}} = \begin{pmatrix} e^{-i\omega/2} & 0 \\ 0 & e^{i\omega/2} \end{pmatrix}, \quad (21)$$

$$\text{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}. \quad (22)$$

While Eqn(20) is the only gate needed to realize Eqn(19), Eqn (21) alone can merely help with the single qubit components of Eqn(18). For the components that require two qubit interaction, the following gate combination (expressed diagrammatically) is used as a template

$$e^{(-iJ_{1,2}\gamma_1)\hat{\sigma}_1^{(z)}\hat{\sigma}_2^{(z)}} = \begin{array}{c} \bullet \\ \text{---} \\ \oplus \end{array} \boxed{R_z(2\gamma_1 J_{1,2})} \begin{array}{c} \bullet \\ \text{---} \\ \oplus \end{array}. \quad (23)$$

While Eqn(23) shows the ZZ interactions for adjacent qubits, this strategy can be generalized to any pair of qubits in the system. Appendix B provides an example of a QAOA circuit for BLLS.

### 3.2.1 For IBM Q specific gates

Our experiments were done on IBM Q devices (ibmq\_london and ibmq\_atheens) available to us through the IBM Q Network. The former has the following basis gates  $\{U_1, U_2, U_3, \text{CNOT}, I\}$  and the latter has  $\{CX, ID, RZ, SX, X\}$ . Although most of these gates are defined and established in the literature well [1], we need to elaborate more on  $U_1, U_2$  and  $U_3$  as they are IBM Q specific gates.

$$U_1(\lambda) = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\lambda} \end{pmatrix}, \quad (24)$$

$$U_2(\lambda, \phi) = \begin{pmatrix} 1/\sqrt{2} & -e^{i\lambda}/\sqrt{2} \\ e^{i\phi}/\sqrt{2} & e^{i(\lambda+\phi)}/\sqrt{2} \end{pmatrix}, \quad (25)$$

$$U_3(\theta, \phi, \lambda) = \begin{pmatrix} \cos \theta/2 & -e^{i\lambda} \sin \theta/2 \\ e^{i\phi} \sin \theta/2 & e^{i(\lambda+\phi)} \cos \theta/2 \end{pmatrix}. \quad (26)$$

We can implement Eqn(20) and Eqn(21) [84] as

$$R_x(\omega) = U_3(\omega, -\frac{\pi}{2}, \frac{\pi}{2}), \quad (27)$$

$$R_z(\omega) = U_3(\pi, 0, \pi)U_1(-\frac{\omega}{2})U_3(\pi, 0, \pi)U_1(\frac{\omega}{2}). \quad (28)$$

Another practical consideration to be taken is the qubit connectivity of a real quantum computer. As the number of qubits increase, it is safe to assume that full connectivity between physical qubits is not feasible to engineer. This means that for distant-qubits to interact with each other, we would need logical qubit replacement using SWAP operations. Appendix C elaborates on this with a demonstration with IBM Q gates.

### 3.3 Experimental methods

The dataset used in our experiments was randomly generated (seeded for reproducibility) consisting of  $A \in \mathbb{R}^{40 \times n}$ ,  $b \in \mathbb{R}^{40}$  and  $x \in \{0, 1\}^n$  where  $n \in \{3, 4, 5, 9, 10\}$  is the size of the problem. Due to the exponential nature of simulating quantum computation on classical hardware and the limited resources in our hands, we decided to limit our problem sizes to 10 qubits. Instead of allocating computational resources to even larger qubit numbers, we focus on: i) increasing the number of problem instances we average over, ii) comparing exact wavefunction versus bitstring-sampling experiments, iii) comparing QAOA performance (waveform) with simulated annealing and random sampling. All values for  $A$  in the dataset are generated by uniformly sampling floating point approximations of real values in the interval  $[-1.0, 1.0]$ , and then rounding the values to 3 decimal places. For each value of  $n$ , we generate 100 test cases with 40 cases in which  $Ax^* = b$ , where the best solution  $x^*$  is sampled randomly and  $b \leftarrow Ax^*$ . The other 60 cases have  $Ax^* \neq b$ , where  $b$  is generated similarly to  $A$  and the best solution  $x^*$  is found by going through all  $2^n$  possible values for  $x$ . This is done to cover both scenarios of the least squares problem. The matrix  $A$  is a sparse matrix having density of 0.2, this was done because sparse matrices have a lot of applications in numerical computation and machine learning [85–87].

We use the QISKIT [88] SDK to write our own implementation of the QAOA algorithm. As mentioned before, ImFil [34] is our black-box optimizer of choice. The only parameter of ImFil we choose to control here is the budget, which governs the maximum iteration limit. The rest of the ImFil parameters for our experiments use their default values. Similarly, unless explicitly stated, all qiskit parameters values taken are default as well. In the later sections, the term quantum virtual machine (QVM) is used to denote classical simulators broadly (unless specified to be of a particular kind). Similarly, the term quantum processing unit (QPU) is used to denote a quantum device. All classical simulations were conducted on standard x86-64 based laptops. Following is a list of the main experiments we conducted, with a few slightly-modified ones described later on in the results section as applicable.

#### 3.3.1 Experiments with no noise

Our first set of experiments on the dataset were done on a simulator with the statevector QVM, giving us the exact waveform. This means that we are able to compute the exact expectation value  $C(\gamma, \beta)$  for the set of angles  $\gamma$  and  $\beta$ . These experiments help assess the performance of QAOA in a perfectly noiseless environment for a large dataset.

The above set of experiments were done for  $p = 1, 2$  and  $3$  with random starting points: 20 for  $p = 1$ , 40 for 2 and 60 for 3 (seeded for reproducibility). Our preliminary study suggested a budget of 200 iterations for  $p = 1, 2$  and 400 iterations for  $p = 3$  respectively. This ensured that at least 70% of our tests converged within the budget while being computationally feasible. At the end of the process, for each problem, the best result from all the starting points is chosen and recorded.

#### 3.3.2 Experiments to compare no noise and shot-noise performance

For our next set of experiments, we use measurement based results on the simulator. Each circuit is run a number of times, specified by the ‘shots’ parameter. This means that the expectation value we get for a given  $\gamma$  and  $\beta$  is approximate in nature. Thus, while quantum circuit simulation itself is noiseless and deterministic in producing the same wavefunction before taking each shot, a finite number of shots is sampled from the resulting wavefunction output probability distribution, introducing a stochastic component. In a real quantum device, one is always limited to this finite tomography, as one has no direct access to the qubit register’s quantum wavefunction.

Since in a real quantum device, we do not have access to the qubit register's waveform, simulations with shot-noise are important to conduct. Each experiment was done 10 times per shot value. The shot values chosen for these experiments are in the set  $\{2^i | n-2 \leq i \leq n+2, i \in \mathbb{Z}\}$ . We chose this range in order to observe the performance in the limit of perfectly reproducing the wavefunction.

Also, the problem instances chosen for this set of experiments are a random subset of the original dataset. For each problem size  $n \in \{3, 4, 5, 9, 10\}$ , we randomly choose 5 problems of the 100 problems (while maintaining the 2 : 3 ratio of the problems by their type). This is done because doing the shot-noise experiments on the original dataset would be computationally infeasible for the limited computational resources at our disposal, since each shot-noise experiment is at least 50 times slower than its statevector counterpart.

The parameters of these experiments have also been modified accordingly. They were done for  $p = 1, 2$  and  $3$ , for a budget of 200 iterations with random starting points: 5 for  $p = 1$ , 10 for  $2$  and 15 for  $3$  (seeded for reproducibility), with the best result being chosen and recorded. For fair comparison, this subset of problem instances was also run with the statevector QVM for the same parameters.

### 3.3.3 Experiments on an IBM Q device

Based on the results of the first two sets of experiments, we design our experiments for the 5 qubit IBM Q devices 'ibmq\_london' and 'ibmq\_athens'. In a real device like this one, the qubits face decoherence issues, coherent gate errors, control errors, incoherent gate errors, leakage, cross-talk, readout noise and more. The first set of IBM Q experiments was to run QAOA for problems with  $n = 5$ , for parameters  $p = 1$ , budget of 200 iterations and a shot value of 1024. The reason for choosing these parameters for QAOA is to take into account the gate depth limitation and noisy computation, thus choosing the minimal number of qubits while still covering a non-trivial problem graph structure, which can still be easily verified with classical methods at this size. The next set of experiments was to take the  $\gamma$  and  $\beta$  from the results of the statevector experiments done in Section 3.3.2 (where  $n = 5$ ) and to try and recreate the distribution and expectation values using the quantum computer.

## 3.4 Results

Two metrics we use here to quantify the performance of QAOA in the simulations are (i) the probability of sampling the best possible solution (or the ground state of Hamiltonian  $\hat{C}$ ) and (ii) the relative error  $\varepsilon_r$  of the expectation value  $C(\gamma, \beta)$  (from Eqn(11)) with respect to the ground state energy  $E_{gs}$ . We define  $\varepsilon_r$  as

$$\varepsilon_r = \left| \frac{C(\gamma, \beta) - E_{gs}}{E_{gs}} \right|. \quad (29)$$

We note that  $\varepsilon_r$  is zero when the expectation value is exactly equal to the groundstate energy, and can grow beyond 1 depending on the spectral width as compared to the groundstate cost.




Figure 1: One instance of convergence behaviour for the QAOA optimization routines run on an exact statevector simulator, shot-noise simulator with  $2^5 = 32$  shots, and two 5-qubit IBM Q processors: `ibmq_london` (QPU 1) and `ibmq_athens` (QPU 2) with  $2^{10} = 1024$  shots each. The problem in question has 5 binary variables mapped to 5 qubits and is done for a QAOA depth of  $p = 1$ . All  $C(\gamma, \beta)$  values dealing with shot-noise (QVM and QPU) are approximate in nature. For comparison, we include the exact ground state energy. The data in the figure is a result of the experiments described in Sections 3.3.2 and 3.3.3.

### 3.4.1 Optimization trajectory for QAOA

In Figure 1, we see an example of how QAOA with ImFil performs on a BLLS problem. As the iterations progress, the fluctuations in the energy expectation value also reduces. This happens either till the black box optimizer converges to a solution (depending on default internal parameters in our case) or the iterations have reached the maximum threshold (governed by the budget). Here the experiments done with the statevector QVM, which has access to the exact energy expectation value, sets the baseline for the other modes of experiments. One interesting thing you will see is that the statevector QVM can not always converge to the ground state energy  $E_{gs}$ . This is because the results are computed for QAOA depth  $p = 1$ . For low circuit depth  $p$ , no guarantees are known to converge to the exact groundstate. As an upper bound, Farhi's theorem [4] states that  $C(\gamma, \beta)$  can reach  $E_{gs}$  as  $p \rightarrow \infty$ , in the general case, assuming the optimal circuit angles are known. While our simulations containing shot-noise due to measurement do relatively well against statevector results, the experiments on a real quantum device are mixed. At the time of writing, the IBM Q devices we tested on did not approximate the theoretically-optimal QAOA result distribution very well, but it still finds the best solution every time. This means they failed to converge using expectation values  $C(\gamma, \beta)$  they produced, but randomly happened to sample the ground state during the process. We have discussed this further in Section 3.4.7 and have included the results for reference.




Figure 2: Comparison of converged final relative error  $\epsilon_r$  (median), for  $p \in \{1, 2, 3\}$ , as a function of problem size  $n \in \{3, 4, 5, 9, 10\}$ . Simulations performed with access to the exact wave form (statevector QVM) done on 100 problem instances per  $n$  with an optimization budget of 200 iterations for  $p \in \{1, 2\}$  and 400 iterations for  $p = 3$ . The dashed curves in the plot are fitted to the experimental data. The information presented in this figure is the output of experiments described in Section 3.3.1.

### 3.4.2 QAOA results with no noise

Before we study the results of QAOA for BLLS with shot-noise, it is important to evaluate the theoretical performance of the same without any noise at all. Figure 2 shows the relative error growth with respect to the problem size  $n$  for the experiments described in Section 3.3.1. We use median as measure of central tendency and median absolute deviation (MAD) for our error bars. Simulations larger than  $p = 3$  take a lot more time for the complete dataset and were computationally infeasible for this project.

Using the information from the medians calculated, we attempt to fit curves for different values of  $p$ . Based on the curve fit, we find that the growth in  $\epsilon_r$  to be polynomial in nature (as described by  $\epsilon_r = a \times n^b$ , where  $a$  and  $b$  are coefficients, see Appendix E for details). This polynomial growth may be partly attributed to some spin configurations with energies close to the ground state. You can see that going from  $p = 1$  to  $p = 2$  decreases the relative error moderately. The difference in performance between  $p = 2$  and  $p = 3$  is more modest, particularly for the larger problem sizes  $n$ . There may be room for further improvement if we allow a larger simulation time budget, for example by tightening the classical optimizer’s convergence parameters and increasing the number of initial starting points for the optimizer. Further rigorous experimentation would be required to draw definite conclusions about the scaling based on such numerics.

### 3.4.3 No noise vs shot-noise optimization

Figure 3 shows us how QAOA with ImFil performs for the parameters described in Section 3.3.2 for statevector and measurement based results for  $2^{n+2}$  shots. We have 5 different problem instances per  $n \in \{3, 4, 5, 9, 10\}$ . The reason for choosing  $2^{n+2}$  shots for this comparison was to try and see if the optimizer could replicate the statevector results given plentiful shots. In subsequent figures, we’ll be showing the performance of the optimizer with fewer number of shots.




Figure 3: Comparison of converged final relative error (median), for  $p \in \{1, 2, 3\}$ , as a function of problem size  $n \in \{3, 4, 5, 9, 10\}$ , done on 5 problem instances per  $n$  with a budget of 200 iterations. (TOP) shows the results from the optimization having access to the exact waveform (statevector simulator), while (BOTTOM) shows the results using a shot-noise simulator with  $2^{n+2}$  shots. For drawing the bottom plot, we use the best angles found using shot-based optimization and used the statevector QVM for one more run at those angles in order to compute the exact expectation value and corresponding relative error. These results are from the experiments described in Section 3.3.2.

We can see the similarities between the top and bottom plots in Figure 3. The main difference however, seems to be the result and error bar overlap between the results of  $p = 1$  and 2. While the two lines are close to each other in the statevector results up to problem size of 9, the measurement-based results for the two parameters are extremely close to each other (when considering median and MAD). This could be attributed to the noise due to approximate results, or due to the small number of experiments we average over, as detailed in Section 3.3.2.

Another effect of the smaller dataset here is that the relative error's growth doesn't seem fully monotonic to the problem size, unlike in Section 3.4.2. However, it still shows a general upward trajectory. Finding optimal parameters for  $p = 4$  and upwards become computationally infeasible due to the time required, even for the smaller dataset we worked with in Figure 3. This can be attributed to the search space growing rapidly as  $p$  is increased [4].

### 3.4.4 Probability of sampling the ground state

We compare the probability of sampling the ground state of the BLS problem (henceforth also referred to as the success probability) for QAOA with two classical methods: random sampling and simulated annealing. For the simulated annealing experiments, we chose 10 steps, an exponential decay schedule, and ran it on problem sizes  $n$  from 3 to 20 (with 100 problems per  $n$ ). The choice of 10 steps was made in order to make a reasonable comparison against our QAOA results, given the size of the problem and the QAOA depths considered. For this problem size, simulated annealing is very cheap computationally on a classical computer, as compared to simulating the quantum circuit with a classical computer. This allows us to consider statistics of a larger number of random problems per  $n$ . See Appendix E for further details.




Figure 4: Comparison of success probabilities of sampling the ground state, for QAOA at depths  $p = 1, 2, 3$ , simulated annealing and random sampling. (TOP) shows the success probability per-query and (BOTTOM) shows the cumulative success probability after 10 queries. The box plots indicate the nature of the experimental results (with medians as triangles) while curves in the figure are fitted to them. The QAOA results are from the statevector (no noise) experiments as described in Section 3.3.1 and the approach for comparing it with simulated annealing and random sampling is described in Section 3.4.4.

The success probabilities for QAOA are calculated from the outputs of the statevector experiments described in Section 3.3.1. More specifically, these would be the success probabilities from sampling the optimized waveform  $|\psi(\gamma, \beta)\rangle$ .

While the success probability of uniform random sampling is easy to calculate per query as  $1/2^n$ , things are not as straightforward with simulated annealing. The main reason being that simulated annealing requires  $k$  number of steps to arrive at a result. Assigning  $k = 1$  would be equivalent to random sampling, but by this logic, assigning  $k > 1$  would yield us a cumulative success probability for random sampling.

Thus, our proposed comparison method assumes an effective probability  $\mathcal{P}_{\text{eff}}$  which is the success probability per query for simulated annealing. This  $\mathcal{P}_{\text{eff}}$  is not directly observed, but calculated by extrapolation from simulated annealing run on  $k > 1$  steps. After  $k$  steps, the cumulative success probability would be:

$$\text{cumulative success prob.} = 1 - (1 - \mathcal{P}_{\text{eff}})^k. \quad (30)$$

Doing a curve fit modeled on Eqn(30) for the results of simulated annealing with  $k = 10$  steps yields us the value of  $\mathcal{P}_{\text{eff}}$  which would be the effective success probability, for simulated annealing, per query. Conversely, curve fits were done on the QAOA per query success probabilities and were extrapolated to cumulative probabilities for 10 queries. This can be conceptualized as the success probability of measuring the ground state at least once from an optimized QAOA quantum state when prepared and measured 10 times. The resulting information is plotted in Figure 4 along with the box plots from the experiments. We refer the interested reader to Appendix E for the details.

As Figure 4 suggests, the success probability for all methods decrease exponentially as the size of the problem increases. The comparison of QAOA with uniform random sampling for BLLS corroborates with previous work done on applying QAOA to a different problem [9]. Even with  $p = 1$ , QAOA performs much better than uniform random sampling. Simulated annealing on BLLS however, yields a higher success probability than QAOA at depth  $p = 3$  when  $n > 8$  (for 10 steps). This empirical result is supported by the theoretical work on QAOA for MAXCUT [89] which indicates that there exists classes of graphs for which quantum advantage is not possible for  $p < 6$ . Another work focusing on numerical simulations on MAXCUT problems suggests any QAOA advantage would need quantum computers to have in the order of hundreds of qubits at the very least [90]. That being said, while our result is interesting, it is not conclusive on its own. As future work, we recommend testing on larger problems, for higher  $p$  and different simulated annealing steps and schedules.

It is one thing to calculate probability, it is another to sample the best solution (or ground state) from a quantum state after QAOA. Figure 5 displays the number of experimental instances where we sample the ground state, at least once, for a particular set of parameters, across various problem sizes and instances (for shot-noise experiments in Section 3.3.2). We contrast this with the analytical results of getting the ground state by uniform random sampling. Here, we see that for optimization done with upto  $2^n$  shots, QAOA has a clear advantage over random sampling. This can be explained by the mechanism of QAOA, which selectively amplifies those bitstring sampling probabilities which have the lowest energy, while suppressing those with higher energy. In this way, the success probabilities may be greatly enhanced over the naive random sampling from the uniform probability distribution.

### 3.4.5 Effect of noise on expectation value

While Figure 1 indicates to us that the current generation of devices are unable to do the entire variational optimization procedure for this problem very effectively, studying the effects of noise on approximating  $C(\gamma, \beta)$  for a particular set of angles gives us information that may




Figure 5: In this bar graph, we collect the number of experiment instances in which we observe the exact ground state bitstring, at least once (on the Y axis) for  $n \in \{3, 4, 5, 9, 10\}$ . Each  $n$  has 5 problem instances, which is repeated 10 times for a given shot value. We compare the QAOA shot-noise simulator experiments (color-coded with the number of binary variables,  $n$ ), with the results one would expect randomly sampling from a uniform distribution  $shots$  times (black bars, labeled with rand, x-axis positioning corresponding to its colour-labeled partner). The QAOA data in this figure comes from sampling the circuit with optimized angle sets  $\gamma^*$  and  $\beta^*$ , after acquiring them by optimizing for a QAOA depth of  $p = 3$  and a budget of 200 iterations, as described in Section 3.3.2.

be useful for future research and development. Here, we investigate the impact of noise by approximating a problem’s expectation value on a noiseless waveform simulator (statevector QVM), a quantum device (QPU), and a range of increasingly noisy simulations in between the two.

In order to do this, we first need a template for the noise model that we can use to make QVMs (for noisy classical simulations). What IBM Q allows through Qiskit [88] is to produce a simplified approximate noise model based on the properties of a real QPU (ibmq\_atheas in our case). We take this extracted noise model as our template and scale each of the quantum and readout errors down in powers of two (upto  $2^{-7}$  for our experiment). Then, from our problem set of 5 variables, we randomly choose 30 problems where the ground state bitstrings do not have all zeros or ones in them. The reason being that all-zeros or all-ones bitstrings might be comparatively easier to sample from a QAOA circuit. For our selected problems, we take the  $(\gamma, \beta)$  pairs responsible for producing the best  $C(\gamma, \beta)$  in their noiseless waveform simulations (circuit depth of  $p = 1$ ), and use them with the various scaled noise models and the device to approximate the expectation value. In order to reduce the effect of shot-noise, we calculated the expectation value with 8192 shots. Since actual devices are limited by qubit connectivity, this also needs to factor in our noisy simulations as it increases the gate operations and the overall circuit depth (see Section 3.2.1 and Appendix C). This information about how physical qubits are connected is referred to as a coupling map of a device. For our experiment, all noisy simulations were done either assuming all-to-all connectivity, or with the device limited-coupling map restriction.

Figure 6 plots the relative errors of  $C(\gamma, \beta)$  (with respect to the problem ground state) for the noiseless waveform, quantum device and noisy simulation expectation values that interpolate between the two. Though our approximation of the device noise model often does not




Figure 6: (TOP) Comparison of relative error  $\epsilon_r$  with respect to the scaled noise model simulations for a typical 5 qubit problem instance (log scale). Noiseless waveform simulation as well as QPU results are also plotted for reference. (BOTTOM) Relative error factor  $\epsilon_r/\epsilon_r^0$  with respect to the scaled noise model simulations (log scale). Here,  $\epsilon_r^0$  refers to the relative error observed in the case of the noiseless waveform simulator and  $\epsilon_r/\epsilon_r^0$  measures  $\epsilon_r$  as a factor of  $\epsilon_r^0$ . Values plotted on the y axis are medians (error bars are in Median Absolution Deviations) calculated from the values of the 30 selected 5 qubit problems. The QPU result in the bottom plot has a MAD of 1.04. Both plots are made with results from a QAOA depth of  $p = 1$ , for device-limited coupling map restrictions (QPU conn.) and all-to-all connectivity (ideal conn.).

mimic the exact effect of the device noise on expectation values, for the scope of this work, it serves its purpose well since the true device result and the unscaled noise model results are not wildly divergent. From the figure, we see that for this type of a problem (and parameters), the relative error in  $C(\gamma, \beta)$  that we observe while using the device is  $\approx 2.5$  times the relative error of waveform simulations, statistically speaking. An in-depth study on the effect of various noise types would be required to make a definite statement on the effect of noise

on QAOA (as it relates to this problem).

Another observation we can make is the how an ideal coupling map, where every qubit is connected to every other qubit, gives us significantly better results than the device’s qubit map, for the very same noise model. This is mainly due to the additional SWAP gates needed for each qubit to interact with all the other qubits for the latter.

### 3.4.6 Effect of shot number on optimization

For QAOA to become practical, the shot number chosen for the computation has to be far less than the number of eigenstates for our cost Hamiltonian ( $2^n$  for our case). For this work, we chose not to randomly guess a shot number value but rather get an understanding of the optimization performance for a set of shot numbers in  $\{2^i | n - 2 \leq i \leq n + 2, i \in \mathbb{Z}\}$ . We hope this helps all future research work in finding better estimates for the least amount of shots required for QAOA, especially for these type of applications. In Figure 7 we see the optimization result for a problem instance where  $n = 5$ . The shot optimization is compared with noiseless statevector optimization. Here it is important to point that we optimized using the stochastic blackbox (using a given shot number) and then calculate the exact expectation value using the wavefunction (i.e, with the statevector QVM) in order to assess the true value of relative error. For the most part, our experiments show that as the problem size increases, we see the optimizer do well even with  $2^{n-2}$  shots. This seems to indicate that the number of shots required to get a good optimization may not be exponential in terms of the problem size, or at least with a smaller exponent than applying random sampling from a uniform distribution. Further research is needed.




Figure 7: Relative error of the optimized QAOA output distribution (y-axis), as a function of number of shots (x-axis), for a problem instance of the  $n = 5$  case. Circles indicate median, and we plot the error-bars (using MAD) when approximating the expectation value with a finite number of shots. QVM:shot-noise refers to the values after the shot-noise optimization and QVM:exact denotes the relative error found after noiseless statevector optimization and is  $\approx 0.2$ . The QAOA depth for all values in the plot is  $p = 1$ . Data represented in the figure is an output of an experiment done in Section 3.3.2.

### 3.4.7 IBM Q device performance

We briefly mentioned our results on the quantum processors in Section 3.4.1. The good news is that the IBM Q devices were always able to find the best solution for our optimization experiments. But that came at the price of taking 1024 shots for each QAOA iteration, which is relatively expensive for a problem size of 5 qubits. When we lowered the shot number, the optimal bitstring was not always sampled and the convergence deteriorated further. The immediate cause of why the optimization process on the device was not close to the simulation results, is the inability of the device to approximate the distribution of the measured bitstrings (for a given circuit).

We shall use the quantum volume (QV) [91] measure for comparing the devices we ran experiments on. In short, it is a scalar number which quantifies the largest square quantum circuit (depth = width) that the QPU can successfully implement. When we consider a pair of  $(\gamma, \beta)$  angles obtained after noiseless optimization; `ibmq_athens`, a device with a QV of 32, comes closer to the actual probability distribution compared to the distribution produced by `ibmq_london` with a QV of 16. For reference, we provide an example of this in Appendix D for the readers. However, even the QPU with QV of 32 wasn't able to efficiently do its part in the variational procedure. And while this result is not surprising for QAOA on the current generation of cloud accessible quantum computing devices [92–95], it is important to emphasize on it particularly for readers from other domains who may be interested in QAOA for their own applications.

It is important to take into account the entire context here. Firstly, the problem graph of our use case is fully connected. Due to the sparse connectivity of the on-chip qubits in the device, logical qubits have to be swapped around a number of times for them to be able to entangle with each other (for the ZZ interactions of our problem). This makes the median gate depth for the final (transpiled) circuits to be about 42 for `ibmq_london` (t-shaped coupling map) and 38 for `ibmq_athens` (linear coupling map). Since two qubit gate fidelity is still low (at the time of writing), the error propagates across the circuit. Secondly, due to the large circuit depth on the real device, we need to take decoherence into account. Thirdly, readout-errors were not considered here and they have significant impact on the noise in the qubit measurement results. It should also be emphasized that in this work, we primarily focused on how to model the BLLS problem using QAOA. Thus, the experiments on the real devices were done “as is”, in order to demonstrate the near-term implementability, without any error mitigation [96]. This could be looked at for future work.

## 3.5 Discussion

We can see the various possibilities and potential advantages QAOA may provide in solving BLLS and similar problems. However, there are challenges that need to be addressed. These are both theoretical and practical in nature.

One theoretical challenge is the proper pre-processing of the problem Hamiltonian by scaling and shifting the coefficients of the objective function, such that we optimally make use of the parameter space  $\beta_l \in [0, \pi], \gamma_l \in [0, 2\pi]$  (most of the problems in the dataset did not suffer from this issue, as we found the default scaling to work well already). However, scaling the problem way beyond necessity also creates issues as the energy landscape is periodic in nature [4]. Thus, one possible way is to use scaling as a heuristic within the QAOA process, and treat it as a hyperparameter to optimize over.

Another challenge, which is both theoretical and practical in nature, is the full connectivity in our problem and in most hard optimization problems in general [97]. Computationally non-trivial problems typically require a high degree of graph connectivity (for instance, planar graphs are easy to solve classically [98]). Simultaneously, a high connectivity poses a challenge

in quantum chip implementation because not all gatesets implement non-nearest neighbour interactions natively. Those need then be implemented effectively by means of a swap network approach [95]. For future work, one option is to modify the problem formulation by not considering the ZZ interactions of a pair of qubits, if its coefficient's magnitude falls below a user defined threshold. This can potentially make it easier for QAOA to run, but its effectiveness in finding the ground state would come under scrutiny. Nonetheless, it can make for interesting future work. Also, error mitigation techniques and readout error correction will also help in improving the results [95]. It would take a combination of the above mentioned approaches to improve performance on a real device.

Challenges aside, we observe the effects of noise on producing the expectation value for a set of  $(\gamma, \beta)$  angles. Although we see that a device with QV of 32 is able to do a closer approximation of the expectation value (for QAOA depth of  $p = 1$ ) than a device with a QV of 16, it is still not good enough for the variational optimization procedure. We also see the potential of a better coupling map for significantly improving the results for the very same noise model. Thus, the solution for improving the results for problems in this application domain need not be considered in unidimensional terms (improving the noise only or coupling only).

One of the next steps would be to explore if QAOA can be valuable in applications that require BLLS as a subroutine, such as NBMF [17]. Another step can be to try the BLLS problem on other types of quantum computers [9–13] to see how different hardware implementations fare. Finally, from a practical standpoint, it may also be beneficial to apply a modified QAOA-like ansatz based on just nearest-neighbor connected CNOT gates, along with a larger number of parameters [99, 100], or apply a SWAP-network QAOA approach [95]. This may potentially lead to faster convergence and better QPU performance for architectures with limited connectivity.

## 4 Conclusion

In this work, we described the implementation of a binary optimization problem, relevant to hard problems in linear algebra, on a gate-based quantum computer via a QAOA approach suitable for NISQ devices. In our simulations, we show how simulated annealing with small number of steps can outperform QAOA for  $p \leq 3$  in terms of sampling the ground state for BLLS type problems. We show that the ImFil optimizer backend performs well under shot-noise for this problem type. From our experiments on IBM Q cloud-based quantum processors and simulations with noise models, we conclude that although machines with a Quantum Volume of 32 are starting to better approximate the noiseless simulation results, it is still challenging to implement linear-depth, high-connectivity circuits on the latest hardware available when it comes to variational optimization for this type of a problem. We expect a future experimental implementation would benefit greatly from gate-error mitigation techniques and post-processing readout errors. It would furthermore be very interesting to see what other hard problems in linear algebra may be implemented using the QAOA ansatz and what their expected performance would be.

## Acknowledgments

The authors would like to thank the people who made this collaboration possible. From UMBC: Dean Drake, Wendy Martin and Cameron McAdams from the office of the vice president for research (OVPR), the office of technology development (OTD) and the office of sponsored programs (OSP) respectively. From Qu & Co: we'd like to thank Benno Broer, CEO and co-

founder of Qu & Co. This work was performed in part using IBM Quantum systems as part of the IBM Q Network.

## A Detailed QUBO formulation for binary linear least squares

In this section of the Appendix, we describe the method by which O'Malley and Vesselinov [27] formulated the binary linear least squares (BLLS) problem. This QUBO formulation will be converted into its equivalent Ising objective function and used in QAOA. Let us begin by writing out  $Ax - b$  which would help us in minimizing  $x$  and thereby solve Eqn(1)

$$Ax - b = \begin{pmatrix} A_{11} & A_{12} & \dots & A_{1n} \\ A_{21} & A_{22} & \dots & A_{2n} \\ \vdots & \vdots & \vdots & \vdots \\ A_{m1} & A_{m2} & \dots & A_{mn} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} - \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_m \end{pmatrix}, \quad (\text{A.1})$$

$$Ax - b = \begin{pmatrix} A_{11}x_1 + A_{12}x_2 + \dots + A_{1n}x_n - b_1 \\ A_{21}x_1 + A_{22}x_2 + \dots + A_{2n}x_n - b_2 \\ \vdots \\ A_{m1}x_1 + A_{m2}x_2 + \dots + A_{mn}x_n - b_m \end{pmatrix}. \quad (\text{A.2})$$

Taking the 2 norm square of the resultant vector of Eqn(A.2), we get

$$\|Ax - b\|_2^2 = \sum_{i=1}^m (|A_{i1}x_1 + A_{i2}x_2 + \dots + A_{in}x_n - b_i|)^2. \quad (\text{A.3})$$

Because we are dealing with real numbers for  $A$  and  $b$ ,  $(|.|)^2 = (.)^2$

$$\|Ax - b\|_2^2 = \sum_{i=1}^m (A_{i1}x_1 + A_{i2}x_2 + \dots + A_{in}x_n - b_i)^2. \quad (\text{A.4})$$

And thus, the coefficients in Eqn(2) are found by expanding Eqn(A.4), which gives us Eqn(13) and Eqn(14). You will notice that there is a constant value from Eqn(A.4) that we leave out of Eqn(13) and Eqn(14). Because this value is not a coefficient for any of the variables, we can't optimize over it and it's left as is, which is  $\|b\|_2^2$ . Also, the ground state energy (QUBO) for when  $\|Ax^* - b\|_2 = 0$  where  $x^*$  is the best solution, is  $-\|b\|_2^2$ .

## B Example QAOA circuit for BLLS

Let us consider a simple problem, without loss of generality, to demonstrate how quantum circuits for BLLS can be designed. Consider the following problem, Find  $x \in \{0, 1\}^3$  such that  $\|Ax - b\|$  is minimized, where

$$A = \begin{pmatrix} 2 & 1 & 1 \\ -1 & 1 & -1 \\ 1 & 2 & 3 \end{pmatrix} \quad (\text{B.1})$$

$$\text{and } b = \begin{pmatrix} 3 \\ 0 \\ 3 \end{pmatrix}. \quad (\text{B.2})$$

This particular problem is more appropriately categorized as a linear system of equations ( $A \in \mathbb{R}^{n^2}$ ) and has a solution  $x^* = (1, 1, 0)^T$ , such that  $Ax^* = b$  or  $\|Ax^* - b\|_2 = 0$ . However, our problem formulation does not change.

In order to solve this problem using QAOA, we require 3 qubits. Using the formulation process detailed in the Appendix A and Eqn(12), the QUBO formulation we get is


$$F(x) = -12x_1 - 12x_2 - 13x_3 + 6x_1x_2 + 12x_1x_3 + 12x_2x_3. \tag{B.3}$$

The constant value that didn't make it to the QUBO here is  $\|b\|_2^2 = 18$ . Converting the QUBO into Ising using Eqn(16), we get

$$F(\sigma) = -1.5\sigma_1 - 1.5\sigma_2 - 0.5\sigma_3 + 1.5\sigma_1\sigma_2 + 3\sigma_1\sigma_3 + 3\sigma_2\sigma_3. \tag{B.4}$$

The offset when going from QUBO to Ising is -11. Therefore, the Ising ground state energy for this problem is  $-18 - (-11) = -7$ . Now let us assume that we are designing a circuit for QAOA where  $p = 1$ . So for a given pair of angles  $(\beta, \gamma)$  a circuit would look like the one shown down below in Eqn(B.5).

The results of this circuit will be used to calculate the expectation value  $C(\gamma, \beta)$ . Based on the expectation value, a new pair of  $\beta$  and  $\gamma$  will be calculated using a classical black box optimization algorithm (like ImFil). These new angles will be fed into another such circuit till the optimization loop converges.



The results from Eqn (B.5) would be classical bitstrings when measured in the standard basis. In order to calculate the energy or cost of a particular bitstring with respect to the Ising Cost function Eqn(3), we would first need to substitute a 1 for each 0 and -1 for each 1 in the bitstring. In short, this is because  $\hat{\sigma}^{(z)}$  describes a quantum state to have an energy of +1 for  $|0\rangle$  and -1 for  $|1\rangle$  (in arbitrary units). For our example, if we measure a bitstring  $\xi$  to be  $\{\xi_1 = 0, \xi_2 = 0, \xi_3 = 1\}$  the equivalent Ising set would be  $\{\sigma_1 = 1, \sigma_2 = 1, \sigma_3 = -1\}$ . Using Eqn(B.4), we get back an energy of -7. This way, we can calculate  $C(\gamma, \beta)$  for diagonal  $\hat{C}$  Hamiltonians by averaging the energies of all measured bitstrings.

### C Implementing two-qubit interactions on a QPU

As mentioned in Section 3.2.1, most practical quantum computers would not have all to all qubit connectivity. But if the problem that we need to solve on a quantum device requires dense connectivity (such as the BLLS), we need SWAP gates for allowing distant qubits to interact with one another. Let us first describe the SWAP gate in its matrix form

$$\text{SWAP} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \tag{C.1}$$

Diagrammatically, we can decompose it with CNOT gates as

$$\text{SWAP} = \begin{array}{c} \times \\ | \\ \times \end{array} = \begin{array}{c} \bullet \oplus \bullet \\ | \quad | \\ \oplus \bullet \oplus \end{array} \quad (\text{C.2})$$

To illustrate how this takes place, consider a hypothetical device where every qubit is only connected to the adjacent qubit in a line. Thus in order to realize the gates described on the LHS of Eqn(C.3), one way is to SWAP between the top and the middle qubit so that it could interact with the bottom qubit.

$$\begin{array}{c} \bullet \quad \bullet \\ | \quad | \\ \oplus \quad \oplus \end{array} R_z(\omega) \oplus = \begin{array}{c} \bullet \oplus \bullet \\ | \quad | \\ \oplus \bullet \oplus \end{array} \begin{array}{c} \bullet \oplus \bullet \\ | \quad | \\ \oplus \bullet \oplus \end{array} \begin{array}{c} \bullet \oplus \bullet \\ | \quad | \\ \oplus \bullet \oplus \end{array} \quad (\text{C.3})$$

$$\begin{array}{c} \oplus \\ | \\ \oplus \end{array} \begin{array}{c} U_1(\frac{\omega}{2}) \\ | \\ U_3(\pi, 0, \pi) \\ | \\ U_1(-\frac{\omega}{2}) \\ | \\ U_3(\pi, 0, \pi) \\ | \\ \oplus \end{array}$$

After our desired two-qubit interaction takes place, the top and middle qubits are swapped again, returning the logical qubits to their original place. Of course, there are other methods (involving SWAP) that the compiler may take to realize the original unitary operations to be performed. In Eqn(C.3), we also show the decomposition of the  $R_z$  gate as defined in Eqn(28).




Figure 8: A typical instance of QAOA-circuit optimized probability distributions of the measured qubit bitstring output. We compare the result for IBM Q devices: `ibmq_london` (TOP) and `ibmq_athens` (MID) versus the qiskit shot-based simulator (BOTTOM), for the same problem instance for a total of 10240 shots. This figure represents data from an experiment done in Section 3.3.3.

## D Probability distribution of running a QAOA circuit for BLLS on a real QPU

Figure 8 shows an example of the distribution we get from running a QAOA circuit with optimized (and fixed)  $\beta$  and  $\gamma$  angles. We run the circuit for a total of 10240 shots on a noiseless QVM (1024 shots over 10 queries to the QPU), the `ibmq_london` and `ibmq_athens` QPUs.

The simulation suggests a few bitstrings with a high probability of being measured (with the ground state having the highest) With `ibmq_athens` we see the QPU noise damping the highest probability strings, whereas the distribution from the `ibmq_london` quantum processor being more evenly spread out, with the ground state not having a significantly high probability of being measured. This can be attributed partially to `ibmq_athens` having a quantum volume (QV) [91] of 32 as opposed to `ibmq_london`'s QV of 16. Other factors that may have made a difference are the coupling map of the devices, and a shorter gate depth for the former.

## E Curve fit and simulated annealing details

### E.1 Curve fit

In this part, we describe the details about the curve fits that appear in Figure 2 and 4. We found that for Figure 2, the polynomial curve of type  $a \times n^b$  for  $b = 0.85$  fits the data moderately well, Table 1 shows further information.

We can see from Table 1, that the fit error (a relative error value) doubles when we go from  $p = 2$  to  $p = 3$ . However, when we explored other values for the exponent, we discovered that  $b \approx 1.26$  reduces the fit error of QAOA depth  $p = 3$  the most to about 0.13, but increases the error fit of  $p = 1$  and 2 to about 0.22 and 0.17 respectively. On the other end, a value of  $b \approx 0.65$  reduces the fit error for  $p = 1$  the most to about 0.04 but raises  $p = 2$  and  $p = 3$  to around 0.12 and 0.26 respectively. Thus the  $b = 0.85$  value is a compromise that works the best for all QAOA depths we consider in this paper. Further research on more problem sizes  $n$  may help to fit better growth curves and get a more definitive value for the exponent  $b$ .

Table 1: Curve fit details for various QAOA depths from Figure 2 where the curve modeled is of the type:  $a \times n^b$  for  $b = 0.85$ . All numerical values are rounded to a precision of 4 decimal places.

QAOA Depth	Coeff. a	Fit Error
$p = 1$	0.0419	0.0852
$p = 2$	0.0226	0.0911
$p = 3$	0.0187	0.1911

Table 2: Curve fit details for various QAOA depths from Figure 4 where the curve modeled is of the type:  $1 - (1 - a/2^{bn})^k$  at  $k = 1$  for QAOA and  $k = 10$  for simulated annealing (and extrapolated to  $k = 10$  and  $k = 1$  respectively for plotting purposes). All numerical values are rounded to a precision of 4 decimal places.

Optimization Method	Coeff. a	Coeff. b	Fit Error
QAOA, $p = 1$	1.5968	0.3967	0.0655
QAOA, $p = 2$	1.7127	0.3090	0.0247
QAOA, $p = 3$	1.8926	0.3013	0.0380
Sim. Anneal (1st attempt)	0.6359	0.1397	0.0217
Sim. Anneal (selected for plot)	1	0.1786	0.0312

The per query and cumulative success probability curves for QAOA and simulated annealing as seen in Figure 4 fits the curve type  $1 - (1 - a/2^{bn})^k$  well (which turns out to be just  $a/2^{bn}$  for  $k = 1$ ). For simulated annealing, we fix  $a = 1$  such that  $\mathcal{P}_{\text{eff}} \rightarrow 1$  as  $n \rightarrow 0$ , this is because  $a \approx 0.63$  when its value is calculated in the curve fitting process, which prevents  $\mathcal{P}_{\text{eff}}$

being 1 at  $n = 0$ . The increase in error when  $a = 1$  is negligible for the simulated annealing curve. However, it is not so negligible in the case of QAOA, and therefore we don't fix the value of  $a$  to 1. This results in the per-query curves to go above 1 and cumulative curves to be less than 1 (after reaching 1) for  $n < 3$ . We consider this to be a modeling artifact and fix the success probability to 1 for  $n < 3$  for both cases.

## E.2 Simulated annealing experiments

The simulated annealing results referenced in Section 3.4.4, Figure 4 and Table 2 were conducted for BLLS problems of sizes 3 to 20. An initial and final temperature of  $T_0 = 100$  and  $T_f = 0.01$  (in arbitrary units) were chosen respectively. We define and use an exponential decay schedule with temperature at iteration  $i$  as

$$T_i = T_0 \left( \frac{T_f}{T_0} \right)^{i/k}. \quad (\text{E.1})$$

Each problem was run 1000 times with a set of randomization seeds for initialization of the Monte Carlo process. Each simulated annealing run was done with  $k = 10$  steps. Our (cumulative) success probability is the count of all the runs where we end up with the ground state divided by 1000.

## References

- [1] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information*, Cambridge University Press, Cambridge, ISBN 9780511976667 (2009), doi:[10.1017/CBO9780511976667](https://doi.org/10.1017/CBO9780511976667).
- [2] E. G. Rieffel and W. H. Polak, *Quantum computing: A gentle introduction*, MIT Press, ISBN 9780262015066 (2011).
- [3] E. Farhi, J. Goldstone, S. Gutmann and M. Sipser, *Quantum computation by adiabatic evolution* [arXiv:quant-ph/0001106](https://arxiv.org/abs/quant-ph/0001106).
- [4] E. Farhi, J. Goldstone and S. Gutmann, *A quantum approximate optimization algorithm*, [arXiv:1411.4028](https://arxiv.org/abs/1411.4028).
- [5] J. Preskill, *Quantum computing in the NISQ era and beyond*, *Quantum* **2**, 79 (2018), doi:[10.22331/q-2018-08-06-79](https://doi.org/10.22331/q-2018-08-06-79).
- [6] T. Kadowaki and H. Nishimori, *Quantum annealing in the transverse Ising model*, *Phys. Rev. E* **58**, 5355 (1998), doi:[10.1103/physreve.58.5355](https://doi.org/10.1103/physreve.58.5355).
- [7] C. C. McGeoch, R. Harris, S. P. Reinhardt and P. I. Bunyk, *Practical annealing-based quantum computing*, *Computer* **52**, 38 (2019), doi:[10.1109/mc.2019.2908836](https://doi.org/10.1109/mc.2019.2908836).
- [8] A. Cross, *The IBM Q experience and QISKit open-source quantum computing software*, *Bull. Am. Phys. Soc.* **63**, L58.00003 (2018).
- [9] J. S. Otterbach et al., *Unsupervised machine learning on a hybrid quantum computer*, [arXiv:1712.05771](https://arxiv.org/abs/1712.05771).
- [10] K. Wright et al., *Benchmarking an 11-qubit quantum computer*, *Nat. Commun.* **10**, 5464 (2019), doi:[10.1038/s41467-019-13534-2](https://doi.org/10.1038/s41467-019-13534-2).

- [11] F. Arute et al., *Quantum supremacy using a programmable superconducting processor*, Nature **574**, 505 (2019), doi:[10.1038/s41586-019-1666-5](https://doi.org/10.1038/s41586-019-1666-5).
- [12] J. Gaebler, B. Bjork, D. Stack, M. Swallows, M. Fabrikant, A. Reed, B. Spaun, J. Pino, J. Dreiling and C. Figgatt, *Progress toward scalable quantum computing at Honeywell Quantum Solutions*, Bull. Am. Phys. Soc. **64**, S01.00103 (2019).
- [13] T. Last et al., *Quantum Inspire: QuTech's platform for co-development and collaboration in quantum computing*, Proc. SPIE **11324**, 113240J (2020), doi:[10.1117/12.2551853](https://doi.org/10.1117/12.2551853).
- [14] S. Chrétien and F. Corset, *Least squares reconstruction of binary images using eigenvalue optimization*, in Compstat, 419, Physica, Heidelberg, ISBN 978-3-642-57489-4 (2002), doi:[10.1007/978-3-642-57489-4\\_63](https://doi.org/10.1007/978-3-642-57489-4_63).
- [15] S. Chrétien and F. Corset, *Using the eigenvalue relaxation for binary least-squares estimation problems*, Signal Process. **89**, 2079 (2009), doi:[10.1016/j.sigpro.2009.04.025](https://doi.org/10.1016/j.sigpro.2009.04.025).
- [16] E. Tsakonas, J. Jalden and B. Ottersten, *Robust binary least squares: Relaxations and algorithms*, in 2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 3780 (2011), doi:[10.1109/icassp.2011.5947174](https://doi.org/10.1109/icassp.2011.5947174).
- [17] D. O'Malley, V. V. Vesselinov, B. S. Alexandrov and L. B. Alexandrov, *Nonnegative/binary matrix factorization with a D-Wave quantum annealer*, PLoS ONE **13**, e0206653 (2018), doi:[10.1371/journal.pone.0206653](https://doi.org/10.1371/journal.pone.0206653).
- [18] D. Ottaviani and A. Amendola, *Low rank non-negative matrix factorization with D-Wave 2000Q*, arXiv:[1808.08721](https://arxiv.org/abs/1808.08721).
- [19] J. Golden and D. O'Malley, *Reverse annealing for nonnegative/binary matrix factorization*, PLoS ONE **16**, e0244026 (2021), doi:[10.1371/journal.pone.0244026](https://doi.org/10.1371/journal.pone.0244026).
- [20] N. Gillis, *Nonnegative matrix factorization*, Society for Industrial and Applied Mathematics, Philadelphia, PA, ISBN 9781611976403 (2020), doi:[10.1137/1.9781611976410](https://doi.org/10.1137/1.9781611976410).
- [21] D. Guillamet and J. Vitrià, *Non-negative matrix factorization for face recognition*, Lect. Notes Comput. Sci. **2504**, 336 (2002), doi:[10.1007/3-540-36079-4\\_29](https://doi.org/10.1007/3-540-36079-4_29).
- [22] D. Kuang, J. Choo and H. Park, *Nonnegative matrix factorization for interactive topic modeling and document clustering*, in Partitional clustering algorithms, Springer International Publishing, Cham (2014), doi:[10.1007/978-3-319-09259-1\\_7](https://doi.org/10.1007/978-3-319-09259-1_7).
- [23] M. R. Blanton and S. Roweis, *K-corrections and filter transformations in the ultraviolet, optical, and near-infrared*, Astron. J. **133**, 734 (2007), doi:[10.1086/510127](https://doi.org/10.1086/510127).
- [24] B. Ren, L. Pueyo, G. Ben Zhu, J. Debes and G. Duchêne, *Non-negative matrix factorization: Robust extraction of extended structures*, Astrophys. J. **852**, 104 (2018), doi:[10.3847/1538-4357/aaa1f2](https://doi.org/10.3847/1538-4357/aaa1f2).
- [25] L. Jacques, J. N. Laska, P. T. Boufounos and R. G. Baraniuk, *Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors*, IEEE Trans. Inform. Theory **59**, 2082 (2013), doi:[10.1109/tit.2012.2234823](https://doi.org/10.1109/tit.2012.2234823).
- [26] A. Bourquard and M. Unser, *Binary compressed imaging*, IEEE Trans. Image Process. **22**, 1042 (2013), doi:[10.1109/tip.2012.2226900](https://doi.org/10.1109/tip.2012.2226900).

- [27] D. O'Malley and V. V. Vesselinov, *ToQ.jl: A high-level programming language for D-Wave machines based on Julia*, in 2016 IEEE High Performance Extreme Computing Conference (HPEC) 1 (2016), doi:[10.1109/hpec.2016.7761616](https://doi.org/10.1109/hpec.2016.7761616).
- [28] C. C. Chang, A. Gambhir, T. S. Humble and S. Sota, *Quantum annealing for systems of polynomial equations*, Sci. Rep. **9**, 10258 (2019), doi:[10.1038/s41598-019-46729-0](https://doi.org/10.1038/s41598-019-46729-0).
- [29] A. Borle and S. J. Lomonaco, *Analyzing the quantum annealing approach for solving linear least squares problems*, Lect. Notes Comput. Sci. **11355**, 289 (2018), doi:[10.1007/978-3-030-10564-8\\_23](https://doi.org/10.1007/978-3-030-10564-8_23).
- [30] T. H. Chang, T. C. H. Lux and S. S. Tipirneni, *Least-squares solutions to polynomial systems of equations with quantum annealing*, Quantum Inf. Process. **18**, 374 (2019), doi:[10.1007/s11128-019-2489-x](https://doi.org/10.1007/s11128-019-2489-x).
- [31] P. Date and T. Potok, *Adiabatic quantum linear regression*, Sci. Rep. **11**, 21905 (2021), doi:[10.1038/s41598-021-01445-6](https://doi.org/10.1038/s41598-021-01445-6).
- [32] S. S. T. Desu, P. K. Srijith, M. V. P. Rao and N. Sivadasan, *Adiabatic quantum feature selection for sparse linear regression*, Lect. Notes Comput. Sci. **12747**, 98 (2021), doi:[10.1007/978-3-030-77980-1\\_8](https://doi.org/10.1007/978-3-030-77980-1_8).
- [33] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik and J. L. O'Brien, *A variational eigenvalue solver on a photonic quantum processor*, Nat. Commun. **5**, 4213 (2014), doi:[10.1038/ncomms5213](https://doi.org/10.1038/ncomms5213).
- [34] C. T. Kelley, *Implicit filtering*, Society for Industrial and Applied Mathematics, ISBN 9781611971897 (2011), doi:[10.1137/1.9781611971903](https://doi.org/10.1137/1.9781611971903).
- [35] N. Gillis, *Introduction to nonnegative matrix factorization*, SIAG/OPT Views News **25**, 7 (2017).
- [36] S. A. Vavasis, *On the complexity of nonnegative matrix factorization*, SIAM J. Optim. **20**, 1364 (2010), doi:[10.1137/070709967](https://doi.org/10.1137/070709967).
- [37] C.-J. Lin, *Projected gradient methods for nonnegative matrix factorization*, Neural Comput. **19**, 2756 (2007), doi:[10.1162/neco.2007.19.10.2756](https://doi.org/10.1162/neco.2007.19.10.2756).
- [38] D. Chen and R. J. Plemmons, *Nonnegativity constraints in numerical analysis*, in The birth of numerical analysis, World Scientific (2009), doi:[10.1142/9789812836267\\_0008](https://doi.org/10.1142/9789812836267_0008).
- [39] O. Tichy, L. Bodiova and V. Smidl, *Bayesian non-negative matrix factorization with adaptive sparsity and smoothness prior*, IEEE Signal Process. Lett. **26**, 510 (2019), doi:[10.1109/lsp.2019.2897230](https://doi.org/10.1109/lsp.2019.2897230).
- [40] F. Glover, G. Kochenberger and Y. Du, *A tutorial on formulating and using QUBO models*, arXiv:[1811.11538](https://arxiv.org/abs/1811.11538).
- [41] B. O'Gorman, R. Babbush, A. Perdomo-Ortiz, A. Aspuru-Guzik and V. Smelyanskiy, *Bayesian network structure learning using quantum annealing*, Eur. Phys. J. Spec. Top. **224**, 163 (2015), doi:[10.1140/epjst/e2015-02349-9](https://doi.org/10.1140/epjst/e2015-02349-9).
- [42] F. Neukart, G. Compostella, C. Seidel, D. von Dollen, S. Yarkoni and B. Parney, *Traffic flow optimization using a quantum annealer*, Front. ICT **4**, 29 (2017), doi:[10.3389/fict.2017.00029](https://doi.org/10.3389/fict.2017.00029).

- [43] E. Farhi and A. W. Harrow, *Quantum supremacy through the quantum approximate optimization algorithm*, [arXiv:1602.07674](https://arxiv.org/abs/1602.07674).
- [44] G. Gallavotti, *Statistical mechanics*, Springer, Berlin, Heidelberg, ISBN 9783642084386 (1999), doi:[10.1007/978-3-662-03952-6](https://doi.org/10.1007/978-3-662-03952-6).
- [45] F. Barahona, *On the computational complexity of Ising spin glass models*, *J. Phys. A: Math. Gen.* **15**, 3241 (1982), doi:[10.1088/0305-4470/15/10/028](https://doi.org/10.1088/0305-4470/15/10/028).
- [46] C. C. McGeoch and C. Wang, *Experimental evaluation of an adiabatic quantum system for combinatorial optimization*, Association for Computing Machinery, New York, NY, USA, ISBN 9781450320535 (2013), doi:[10.1145/2482767.2482797](https://doi.org/10.1145/2482767.2482797).
- [47] S. Aaronson, *Guest column: NP-complete problems and physical reality*, *SIGACT News* **36**, 30 (2005), doi:[10.1145/1052796.1052804](https://doi.org/10.1145/1052796.1052804).
- [48] S. Amaran, N. V. Sahinidis, B. Sharda and S. J. Bury, *Simulation optimization: a review of algorithms and applications*, *Ann. Oper. Res.* **240**, 351 (2015), doi:[10.1007/s10479-015-2019-x](https://doi.org/10.1007/s10479-015-2019-x).
- [49] L. M. Rios and N. V. Sahinidis, *Derivative-free optimization: a review of algorithms and comparison of software implementations*, *J. Glob. Optim.* **56**, 1247 (2012), doi:[10.1007/s10898-012-9951-y](https://doi.org/10.1007/s10898-012-9951-y).
- [50] J. Moćkus, *On Bayesian methods for seeking the extremum*, in *Optimization techniques IFIP technical conference*, Springer, Berlin, Heidelberg (1975), doi:[10.1007/978-3-662-38527-2\\_55](https://doi.org/10.1007/978-3-662-38527-2_55).
- [51] A. Mayer, T. Mora, O. Rivoire and A. M. Walczak, *Diversity of immune strategies explained by adaptation to pathogen statistics*, *Proc. Natl. Acad. Sci. USA* **113**, 8630 (2016), doi:[10.1073/pnas.1600663113](https://doi.org/10.1073/pnas.1600663113).
- [52] J. C. Spall, *Implementation of the simultaneous perturbation algorithm for stochastic optimization*, *IEEE Trans. Aerosp. Electron. Syst.* **34**, 817 (1998), doi:[10.1109/7.705889](https://doi.org/10.1109/7.705889).
- [53] A. W. Harrow, A. Hassidim and S. Lloyd, *Quantum algorithm for linear systems of equations*, *Phys. Rev. Lett.* **103**, 150502 (2009), doi:[10.1103/physrevlett.103.150502](https://doi.org/10.1103/physrevlett.103.150502).
- [54] N. Wiebe, D. Braun and S. Lloyd, *Quantum algorithm for data fitting*, *Phys. Rev. Lett.* **109**, 050505 (2012), doi:[10.1103/physrevlett.109.050505](https://doi.org/10.1103/physrevlett.109.050505).
- [55] B. D. Clader, B. C. Jacobs and C. R. Sprouse, *Preconditioned quantum linear system algorithm*, *Phys. Rev. Lett.* **110**, 250504 (2013), doi:[10.1103/physrevlett.110.250504](https://doi.org/10.1103/physrevlett.110.250504).
- [56] I. Kerenidis and A. Prakash, *Quantum recommendation systems*, [arXiv:1603.08675](https://arxiv.org/abs/1603.08675).
- [57] A. M. Childs, R. Kothari and R. D. Somma, *Quantum algorithm for systems of linear equations with exponentially improved dependence on precision*, *SIAM J. Comput.* **46**, 1920 (2017), doi:[10.1137/16m1087072](https://doi.org/10.1137/16m1087072).
- [58] H. Wang and H. Xiang, *Quantum algorithm for total least squares data fitting*, *Phys. Lett. A* **383**, 2235 (2019), doi:[10.1016/j.physleta.2019.04.037](https://doi.org/10.1016/j.physleta.2019.04.037).
- [59] L. Lin and Y. Tong, *Optimal polynomial based quantum eigenstate filtering with application to solving quantum linear systems*, *Quantum* **4**, 361 (2020), doi:[10.22331/q-2020-11-11-361](https://doi.org/10.22331/q-2020-11-11-361).

- [60] I. Kerenidis and A. Prakash, *Quantum gradient descent for linear systems and least squares*, Phys. Rev. A **101**, 022316 (2020), doi:[10.1103/physreva.101.022316](https://doi.org/10.1103/physreva.101.022316).
- [61] Y. Tong, D. An, N. Wiebe and L. Lin, *Fast inversion, preconditioned quantum linear system solvers, fast Green's-function computation, and fast evaluation of matrix functions*, Phys. Rev. A **104**, 032422 (2021), doi:[10.1103/PhysRevA.104.032422](https://doi.org/10.1103/PhysRevA.104.032422).
- [62] S. Aaronson, *Read the fine print*, Nat. Phys. **11**, 291 (2015), doi:[10.1038/nphys3272](https://doi.org/10.1038/nphys3272).
- [63] V. Giovannetti, S. Lloyd and L. Maccone, *Quantum random access memory*, Phys. Rev. Lett. **100**, 160501 (2008), doi:[10.1103/physrevlett.100.160501](https://doi.org/10.1103/physrevlett.100.160501).
- [64] V. Giovannetti, S. Lloyd and L. Maccone, *Architectures for a quantum random access memory*, Phys. Rev. A **78**, 052310 (2008), doi:[10.1103/physreva.78.052310](https://doi.org/10.1103/physreva.78.052310).
- [65] C. Ciliberto, M. Herbster, A. D. Ialongo, M. Pontil, A. Rocchetto, S. Severini and L. Wossnig, *Quantum machine learning: a classical perspective*, Proc. R. Soc. A: Math. Phys. Eng. Sci. **474**, 20170551 (2018), doi:[10.1098/rspa.2017.0551](https://doi.org/10.1098/rspa.2017.0551).
- [66] A. Montanaro and S. Pallister, *Quantum algorithms and the finite element method*, Phys. Rev. A **93**, 032324 (2016), doi:[10.1103/physreva.93.032324](https://doi.org/10.1103/physreva.93.032324).
- [67] R. Ayanzadeh, S. Mousavi, M. Halem and T. Finin, *Quantum annealing based binary compressive sensing with matrix uncertainty*, [arXiv:1901.00088](https://arxiv.org/abs/1901.00088).
- [68] S. W. Shin, G. Smith, J. A. Smolin and U. Vazirani, *How "quantum" is the D-Wave machine?*, [arXiv:1401.7087](https://arxiv.org/abs/1401.7087).
- [69] T. Albash and I. Hen, *Future of physical quantum annealers: impediments and hopes*, Sci. Cult. **85**, 163 (2019).
- [70] F. Glover, *Tabu search: A tutorial*, Interfaces **20**, 74 (1990), doi:[10.1287/inte.20.4.74](https://doi.org/10.1287/inte.20.4.74).
- [71] Gurobi Optimization, LLC, *Gurobi optimizer reference manual* (2021).
- [72] S. Mandrà, Z. Zhu, W. Wang, A. Perdomo-Ortiz and H. G. Katzgraber, *Strengths and weaknesses of weak-strong cluster problems: A detailed overview of state-of-the-art classical heuristics versus quantum approaches*, Phys. Rev. A **94**, 022337 (2016), doi:[10.1103/physreva.94.022337](https://doi.org/10.1103/physreva.94.022337).
- [73] D. An and L. Lin, *Quantum linear system solver based on time-optimal adiabatic quantum computing and quantum approximate optimization algorithm*, [arXiv:1909.05500](https://arxiv.org/abs/1909.05500).
- [74] C.-C. Chen, S.-Y. Shiao, M.-F. Wu and Y.-R. Wu, *Hybrid classical-quantum linear solver using noisy intermediate-scale quantum machines*, Sci. Rep. **9**, 16251 (2019), doi:[10.1038/s41598-019-52275-6](https://doi.org/10.1038/s41598-019-52275-6).
- [75] X. Xu, J. Sun, S. Endo, Y. Li, S. C. Benjamin and X. Yuan, *Variational algorithms for linear algebra*, Sci. Bull. **66**, 2181 (2021), doi:[10.1016/j.scib.2021.06.023](https://doi.org/10.1016/j.scib.2021.06.023).
- [76] C. Bravo-Prieto, R. LaRose, M. Cerezo, Y. Subasi, L. Cincio and P. Coles, *Variational quantum linear solver: A hybrid algorithm for linear systems*, Bull. Am. Phys. Soc. **65**, F07.00005 (2020).
- [77] X. Wang, Z. Song and Y. Wang, *Variational quantum singular value decomposition*, Quantum **5**, 483 (2021), doi:[10.22331/q-2021-06-29-483](https://doi.org/10.22331/q-2021-06-29-483).

- [78] S. Hadfield, Z. Wang, B. O’Gorman, E. Rieffel, D. Venturelli and R. Biswas, *From the quantum approximate optimization algorithm to a quantum alternating operator ansatz*, *Algorithms* **12**, 34 (2019), doi:[10.3390/a12020034](https://doi.org/10.3390/a12020034).
- [79] V. Choi, *Minor-embedding in adiabatic quantum computation: I. The parameter setting problem*, *Quantum Inf. Process.* **7**, 193 (2008), doi:[10.1007/s11128-008-0082-9](https://doi.org/10.1007/s11128-008-0082-9).
- [80] V. Choi, *Minor-embedding in adiabatic quantum computation: II. Minor-universal graph design*, *Quantum Inf. Process.* **10**, 343 (2010), doi:[10.1007/s11128-010-0200-3](https://doi.org/10.1007/s11128-010-0200-3).
- [81] Z.-C. Yang, A. Rahmani, A. Shabani, H. Neven and C. Chamon, *Optimizing variational quantum algorithms using Pontryagin’s minimum principle*, *Phys. Rev. X* **7**, 021027 (2017), doi:[10.1103/physrevx.7.021027](https://doi.org/10.1103/physrevx.7.021027).
- [82] L. Zhou, S.-T. Wang, S. Choi, H. Pichler and M. D. Lukin, *Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices*, *Phys. Rev. X* **10**, 021067 (2020), doi:[10.1103/physrevx.10.021067](https://doi.org/10.1103/physrevx.10.021067).
- [83] L. T. Brady, C. L. Baldwin, A. Bapat, Y. Kharkov and A. V. Gorshkov, *Optimal protocols in quantum annealing and quantum approximate optimization algorithm problems*, *Phys. Rev. Lett.* **126**, 070505 (2021), doi:[10.1103/physrevlett.126.070505](https://doi.org/10.1103/physrevlett.126.070505).
- [84] A. J. et al., *Quantum algorithm implementations for beginners*, [arXiv:1804.03719](https://arxiv.org/abs/1804.03719).
- [85] D. J. Rose and R. A. Willoughby, *Sparse matrices and their applications*, Springer US, Boston, MA, ISBN 9781461586777 (1972), doi:[10.1007/978-1-4615-8675-3](https://doi.org/10.1007/978-1-4615-8675-3).
- [86] D. Chung and S. Keles, *Sparse partial least squares classification for high dimensional data*, *Stat. Appl. Genet. Mol. Biol.* **9**, 17 (2010), doi:[10.2202/1544-6115.1492](https://doi.org/10.2202/1544-6115.1492).
- [87] S. Venkataraman, E. Bodzsar, I. Roy, A. AuYoung and R. S. Schreiber, *Presto: Distributed machine learning and graph processing with sparse matrices*, in *EuroSys ’13: Proceedings of the 8th ACM European Conference on Computer Systems* 197 (2013), doi:[10.1145/2465351.2465371](https://doi.org/10.1145/2465351.2465371).
- [88] G. Aleksandrowicz et al., *Qiskit: An open-source framework for quantum computing*, accessed on: Mar. 16 (2019), doi:[10.5281/zenodo.2562110](https://doi.org/10.5281/zenodo.2562110).
- [89] J. Wurtz and P. Love, *MaxCut quantum approximate optimization algorithm performance guarantees for  $p > 1$* , *Phys. Rev. A* **103**, 042612 (2021), doi:[10.1103/physreva.103.042612](https://doi.org/10.1103/physreva.103.042612).
- [90] G. G. Guerreschi and A. Y. Matsuura, *QAOA for Max-Cut requires hundreds of qubits for quantum speed-up*, *Sci. Rep.* **9**, 6903 (2019), doi:[10.1038/s41598-019-43176-9](https://doi.org/10.1038/s41598-019-43176-9).
- [91] A. W. Cross, L. S. Bishop, S. Sheldon, P. D. Nation and J. M. Gambetta, *Validating quantum computers using randomized model circuits*, *Phys. Rev. A* **100**, 032328 (2019), doi:[10.1103/physreva.100.032328](https://doi.org/10.1103/physreva.100.032328).
- [92] M. Alam, A. Ash-Saki and S. Ghosh, *Analysis of quantum approximate optimization algorithm under realistic noise in superconducting qubits*, [arXiv:1907.09631](https://arxiv.org/abs/1907.09631).
- [93] Utkarsh, B. K. Behera and P. K. Panigrahi, *Solving vehicle routing problem using quantum approximate optimization algorithm*, [arXiv:2002.01351](https://arxiv.org/abs/2002.01351).

- [94] M. Willsch, D. Willsch, F. Jin, H. De Raedt and K. Michielsen, *Benchmarking the quantum approximate optimization algorithm*, Quantum Inf. Process. **19**, 197 (2020), doi:[10.1007/s11128-020-02692-8](https://doi.org/10.1007/s11128-020-02692-8).
- [95] M. P. Harrigan et al., *Quantum approximate optimization of non-planar graph problems on a planar superconducting processor*, Nat. Phys. **17**, 332 (2021), doi:[10.1038/s41567-020-01105-y](https://doi.org/10.1038/s41567-020-01105-y).
- [96] A. Kandala, K. Temme, A. D. Córcoles, A. Mezzacapo, J. M. Chow and J. M. Gambetta, *Error mitigation extends the computational reach of a noisy quantum processor*, Nature **567**, 491 (2019), doi:[10.1038/s41586-019-1040-7](https://doi.org/10.1038/s41586-019-1040-7).
- [97] V. Akshay, H. Philathong, I. Zacharov and J. Biamonte, *Reachability deficits in quantum approximate optimization of graph problems*, Quantum **5**, 532 (2021), doi:[10.22331/q-2021-08-30-532](https://doi.org/10.22331/q-2021-08-30-532).
- [98] F. Hadlock, *Finding a maximum cut of a planar graph in polynomial time*, SIAM J. Comput. **4**, 221 (1975), doi:[10.1137/0204019](https://doi.org/10.1137/0204019).
- [99] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow and J. M. Gambetta, *Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets*, Nature **549**, 242 (2017), doi:[10.1038/nature23879](https://doi.org/10.1038/nature23879).
- [100] P. Kl. Barkoutsos, G. Nannicini, A. Robert, I. Tavernelli and S. Woerner, *Improving variational quantum optimization using CVaR*, Quantum **4**, 256 (2020), doi:[10.22331/q-2020-04-20-256](https://doi.org/10.22331/q-2020-04-20-256).