

University of Plymouth

PEARL

<https://pearl.plymouth.ac.uk>

Faculty of Science and Engineering

School of Engineering, Computing and Mathematics

2023-12-20

Investigating techniques to optimise the layout of turbines in a windfarm using a quantum computer

James Hancock, Matthew J. Craven, Craig McNeile, and Davide Vadicchino

Centre for Mathematical Sciences, University of Plymouth

We study the optimal placement of wind turbines within a windfarm to maximize the power produced by mapping the system to a Quadratic Unconstrained Binary Optimisation (QUBO) problem. We investigate solving the resulting QUBO problem using the Variational Quantum Eigensolver (VQE) on a quantum computer simulator and compare the results to those from two classical optimisation methods: simulated annealing and the Gurobi solver. The maximum grid size we study is 4×4 , which requires 16 qubits.

1 Introduction

With the climate crisis and the continuous goal to reduce our emissions, there is a clear motivation for extracting the maximum amount of energy from renewable sources. The use of renewable energy sources will minimize our negative effect on the environment and lead to a more sustainable future.

One key renewable energy source is wind. There are many challenges to making it an economically viable and reliable energy source [1]. In this paper, we focus on using quantum computers to solve one specific problem towards maximizing the energy extracted from the wind. There are currently limitations on the available quantum computing hardware that we will discuss later. There have been a number of reviews of the use of quantum computing to solve problems in renewable energy [2, 3, 4]. There has been a recent review on using quantum computers to solve optimisation problems [5].

Energy from the wind is extracted using wind turbines. Typically, wind turbines are combined together to form windfarms. The spatial layout

of the turbines in the windfarm will change the maximum amount of energy that the windfarm can produce. There is a long history of using optimization techniques [6, 7] to find the best locations for the turbines in the windfarm to produce the maximum amount of power.

The number of turbines in the windfarms increases to maximize the power produced. For example, China is building a windfarm with 3500 turbines [8]. As the number of turbines increases, the complexity of the optimization problem grows exponentially. This motivates the study of improved optimization algorithms for finding the best layout of turbines in a windfarm. Recently, one form of the Windfarm Layout Optimization (WFLO) problem has been rewritten as a Quadratic Unconstrained Binary Optimization (QUBO) problem [9] that can be solved on a quantum computer.

There are currently two dominant types of quantum computers. One is based on quantum circuits. For example, IBM and Google have constructed quantum computers based on this paradigm. The current largest real-world quantum computer is IBM's Osprey [10], which has 433 qubits. One way to solve QUBO problems using a quantum circuit computer is to use a Variational Quantum Eigensolver (VQE).

Other types of quantum computers are quantum annealers (QA) [11]. Quantum annealers are built to solve QUBO problems. D-WAVE's Advantage is the largest real-world QA, having 5000+ qubits (although not all connected) [12]. A new system in the coming years which promises to have 7000+ qubits [13] is currently under development. Fujitsu's Digital Annealer can also solve QUBO problems using specialized hardware [14].

Another important practical issue is that quantum error correction is not implemented on existing quantum computers. We are currently in the noisy intermediate-scale quantum (NISQ) [6] era.

James Hancock: james.hancock@plymouth.ac.uk

Any algorithm should be resilient to decoherence. Practical algorithms in the NISQ era typically run on hybrid systems made from a combination of classical and quantum computers. We will postpone the study of the effect of the additional problems, such as decoherence and quantum errors which occur when running on real quantum hardware, to later work.

In this paper, we investigate using quantum optimisation algorithms to solve the WFLO problem mapped to a QUBO problem using the `qiskit` package [15] from IBM, running on a classical simulator. The solution method uses the Variational Quantum Eigensolver (VQE), so it is an appropriate method for the NISQ era of quantum computing [16]. This work is a necessary first step to solving the WFLO problem using quantum computers. The next steps are studying the introduction of quantum noise and errors in the simulator, and finally using a real quantum computer.

Senderovich et al. [9] developed the formalism to convert a WFLO problem into a QUBO problem. They compared solving the resulting QUBO problem using classical methods and the Fujitsu Digital Annealer. Our contribution is to investigate solving the QUBO problem mapped from a WFLO using a simulator of a circuit-based quantum computer.

Apart from the importance of solving the WFLO problem in order to produce more electricity from a windfarm, there is another motivation for this study. QUBO problems are one of the main problems that can be solved by quantum annealers such as those sold by D-Wave and Fujitsu [17]. Once our project is running on quantum computing hardware, the solution of WFLO problems may be an interesting test case for comparing the performance of adiabatic and circuit-based quantum computers. Previous work has shown the difficulty in comparing the performance of classical computer and adiabatic annealers in solving QUBO problems [18, 19, 20, 21].

In this work we have used the circuit model of quantum computing as the basis of the calculation. The next step in the project will be to run the code on quantum computers and we will then be able to compare the performance on circuit and quantum-based systems.

2 Windfarm Layout Optimisation

Choosing the placement of turbines within a windfarm, or WFLO [6] is selecting the locations of the turbines in order to maximize the power output, with respect to certain constraints. These constraints include the maximum number of turbines we can use, as well as the minimum distance that we must have between them due to their rotors. We use the simplifying assumption that the turbines can be placed only on a discrete grid.

When a turbine is downstream (meaning that it is behind the other turbine with respect to the wind direction) from another within its cone, it will have reduced power output compared to if it is out in the open. This effect is called the wake of a turbine. The key issue in the placement of the wind turbines in the windfarm is to minimize the lost energy due to the wakes from the turbines.

The Sum-of-Squares (SS) model for wake speeds [22] best captures the effect of these wakes. However, the formulation leads to intractable optimisation problems. To combat this, we use a simplified model, the Linear Superposition of Wakes (LS) [23]. This allows us to map the problem to a Quadratic Unconstrained Binary Optimisation (QUBO) problem, which can then in turn be optimized using various quantum (and classical) methods.

3 Windfarm layout as a Quadratic Unconstrained Binary Optimisation problem

A QUBO problem is defined as:

$$\operatorname{argmin}_x f_Q(x), \quad (1)$$

where,

$$f_Q(x) := x^T Q x = \sum_{i,j=0}^q Q_{ij} x^i x^j \quad (2)$$

$$x \in \mathbb{B}^q = \{0, 1\}^q.$$

Many combinatorial optimization problems can be formulated as QUBO problems. For example, the reconstruction of the tracks of charged particles can be computed by solving a QUBO problem [24]. QUBO is a class of NP-hard problems [25]. It is not expected that the solution

of QUBO problems will be in the BQP complexity class (problems that can be solved in polynomial time on quantum computers). QUBO problems are normally formulated in an unconstrained way. However, constraints can be included by introducing penalties in the form of large additive constants when the conditions are violated. We follow the mapping of the WFLO problem to a QUBO problem that was developed by Senderovich et al in Ref. [9].

We will simplify the land that the farm is on to a square grid of length l_{grid} . There are in total l_{grid}^2 spots on which a turbine could be placed. These will be labelled as they are in Fig. 1. The indicator of the presence of a turbine on site q is x^q where $q \in \{1, \dots, l_{grid}^2\}$. Cartesian coordinates of the grid are shown in Fig. 2.

There are two constraints that we impose on this simplified model of a windfarm, which are:

1. There can be no more than m turbines on the grid.
2. Each turbine must be at least distance ξ apart from one another.

1	2	3	4
5	6	7	8
9	10	11	12
13	14	15	16

Figure 1: Labelling of sites on a $l_{grid} = 4$ windfarm grid.

(1,1)	(1,2)	(1,3)	(1,4)
(2,1)	(2,2)	(2,3)	(2,4)
(3,1)	(3,2)	(3,3)	(3,4)
(4,1)	(4,2)	(4,3)	(4,4)

Figure 2: Coordinates of sites on a $l_{grid} = 4$ windfarm grid.

A wake can be defined by three parameters: the angle α (degrees) that the wind is coming

from with respect to the west counted clockwise, x , the maximum distance that is affected by the wake and r is the radius with which it spreads out per unit distance, away from the turbine causing the wake. These physical factors are problem-specific. For our model, we say that if a grid location is at all within the wake, it is entirely within its wake. We denote a wake starting at position i as $w(i, x, r; \alpha)$.




Figure 3: Turbine on location 15 or (1,4), $D = \{0, 12ms^{-1}, 1\}$.

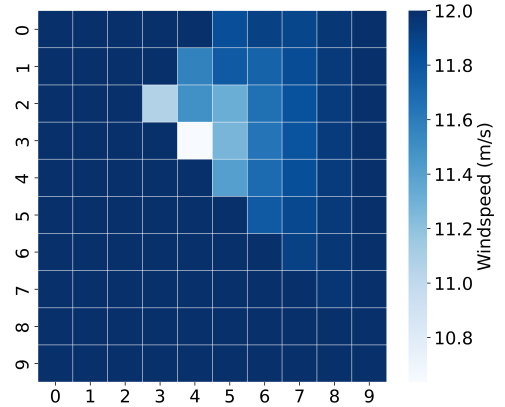


Figure 4: Turbine on location 23 or (2,2) and 34 or (3,3), $D = \{0, 12ms^{-1}, 1\}$.




Figure 5: Turbine on location 27 or (2,6), $D = \text{Second Mosetti}$.




Figure 6: Turbine on location 35 or (3,4) and 76 or (7,5), $D = \text{Second Mosetti}$.

To visualize the wakes from some turbines in the wind, they are plotted on a $l_{grid} = 10$ grid (that is bigger than used in the simulations) with the parameters: $x = 3$, $r = 1$. Fig. 3 and Fig. 5 show a configuration of the wakes from one turbine. Fig. 4 and Fig. 6 show a configuration of the wakes from two turbines. If one of the turbines is in the wake of another turbine, then it will see a reduced wind velocity and thus produce less power. The algorithm's goal is to place wind turbines in locations that have the windspeeds maximized.

3.1 Model

To ensure that the model is more realistic, the system includes several wind regimes. Each possible wind regime is defined as, $d := \{\alpha_d, v_d, p_d\}$. α_d is the angle at which the wind comes in, v_d is the *free wind speed*, i.e. the wind speed that powers a turbine not in the wake of another. p_d is the probability that we would expect to encounter

this wind regime. We denote the set of possible wind regimes in our system as D . To clarify this, D is a collection of d 's, with $\sum_{d \in D} p_d = 1$

The particular model that we are using is the Linear Superposition (LS) of wakes. This model is useful as it can be transformed into a QUBO problem. The power output for a system of q sites within this model is computed as:

$$E^{LS} = \sum_{d \in D} \sum_{i=1}^q p_d \left[\frac{1}{3} v_d^3 - \sum_{j \in w_i} \frac{1}{3} (v_d^3 - u_{ij}^3) \right], \quad (3)$$

where w_i is shorthand for the wake at i , and u_{ij} is the reduced windspeed at site j , caused by being in the wake of i . If j is in the wake of i , u_{ij} can be calculated as in,

$$u_{ij} = v_d \left(1 - \frac{2a}{(1 + \alpha_T (\delta/r)^2)^2} \right). \quad (4)$$

This is the Jansen wake model [26], where a is the axial induction factor, currently set to 0.1, and $\delta = \|i - j\|_2$. When we control the distance which the wake can go back (x), as well as its radius per distance (r), we can calculate α_T using Eq. 5.

$$\alpha_T = \frac{1}{x} (r - r_t), \quad (5)$$

where r_t is the turbine radius which must be set. For this work, we set $r_t = 0.33$ [27]. We note that r_t and a from Eq. 4 carry units, and change in real systems dependent on the size and energy output of the turbines. The values which have been chosen are arbitrary; they do not relate to any real-life system - this is a simplified model.

The QUBO formulation for this problem is then:

$$\operatorname{argmin}_x (-f(x)), \quad (6)$$

$$f(x) = \sum_{d \in D} \sum_{i=1}^m p_d \left[\frac{1}{3} v_d^3 x^i x^i - \sum_{j \in w_i} \frac{1}{3} (v_d^3 - u_{ij}^3) x^i x^j \right]. \quad (7)$$

This is the unconstrained version of the problem, i.e. the optimal solution will always be a turbine on every site. We now must add the constraints in the form of an energetic penalty.

$$g(x; \lambda_1, \lambda_2) = \lambda_1 \left(\sum_{i=1}^{l_{grid}^2} x^i - m \right)^2 + \lambda_2 \sum_{\|i-j\|_2 < \xi} x^i x^j \quad (8)$$

The term which is multiplied by λ_1 in Eq. 8 limits the number of total turbines and the second term pertains to the minimum distance. The values of λ_1, λ_2 have to be large enough so that the constraints are met. The full QUBO problem is:

$$\operatorname{argmin}_x (-f(x) + g(x; \lambda_1, \lambda_2)) \quad (9)$$

The problem from Eq. 9 can then be written in the form of a weight matrix Q from Eq. 2:

$$Q_{ij} = \begin{cases} -\frac{1}{3} \sum_d p_d v_d^3 + \lambda_1(1 - 2nm) & \text{if } i = j \\ -\frac{1}{3} \sum_d p_d (v_d^3 - u_{ij}^3) + 2\lambda_1 & \text{if } i \neq j \end{cases} \quad (10)$$

Here, we have neglected the proximity constraints term as we do not include this in our simulations. This can often lead to dense matrices, with all entries nonzero. We note here that $q = l_{grid}^2$ is the number of variables/qubits.

4 Classical methods

As a baseline comparison to the use of quantum computers to solve QUBO problems we also investigated two classical optimization techniques.

Gurobi is a set of solver packages that can be used in many different programming languages to solve linear and quadratic optimization problems. The package uses several different optimization techniques, choosing the best depending on the problem at hand. Examples of the algorithms that are made use of are: simplex and parallel barrier, further details can be found in Ref. [28]. In this work we used Gurobi to solve our QUBO problem, which is a special case of quadratic optimization. Gurobi has been used as a comparison to using quantum and digital annealers to solve different types of QUBO problems [29, 30].

Simulated Annealing [31] (SA) is a standard approach to finding the optimum of a function which has many local minima in a huge search space. It is particularly useful for searching in discrete spaces [32]. The algorithm searches the space via random fluctuations, controlled by a probability to transfer between states which is a function of temperature. SA has been used to find solutions to various NP-complete combinatorial problems such as the Travelling Salesman Problem (TSP) [33], Minimum Linear Arrangement (MLA) [34] and instances of the Packing Problem (PP) [35].

5 Variational Quantum Eigensolver

The Variational Quantum Eigensolver (VQE) is a hybrid quantum-classical algorithm for finding the lowest eigenvalue of a Hamiltonian that is in the form of a Pauli string [36]. For further details about the VQE see Ref. [37]. The VQE is an algorithm that has shown great promise in making use of current NISQ computers. It uses a small parameterized quantum circuit to calculate the expected value of the Hamiltonian with respect to the parameters (θ). A classical optimization routine is then used in order to best choose θ so that we minimize this value.

The VQE uses the inequality,

$$E_{min}^H \leq \langle 0|U^\dagger(\theta)HU(\theta)|0\rangle, \quad (11)$$

and the variational principle in order to find a tight upper bound on the lowest eigenvalue, E_{min}^H , of a Pauli string matrix, H which are defined as follows,

$$H = \sum_{\gamma} h_{\gamma} P_{\gamma} \quad (12)$$

$$P_{\gamma} = \bigotimes_{i=0}^q \sigma_{m_i} \quad (13)$$

Where $m_i \in \{0, 3\}$ tells us the Pauli matrix and q is the number of qubits. The two Pauli matrices that we make use of in this problem are:

$$\sigma_0 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} = I_2, \quad (14)$$

and

$$\sigma_3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \quad (15)$$

The parameter values are then chosen by a classical optimization routine to minimize this value.

$$\min_{\theta} \langle 0|U^\dagger(\theta)HU(\theta)|0\rangle \quad (16)$$

There are two main features of the VQE that we must design and control in order to make use of it efficiently. These are: the design of the parameterized circuit (the *ansatz*) and the choice of optimization routine. Our choice of ansatz must create states which have sufficient overlap with solution space in order to be able to find the minimum. When using the VQE to solve QUBO problems this is fairly trivial as one parameter on each qubit would be sufficient (we are only looking for

basis states). However, the choice of optimization routine is certainly non-trivial and is vital in order to best make use of the algorithm. In this work we examine the effectiveness of three different optimization routines: Powell’s gradient-free conjugate search method (PO), Constrained Optimization by Linear Approximation (COByLA) and Bayesian Optimization (BO).

5.0.1 Powell optimization

PO is a widely known standard optimization routine for finding the global minimum of a function where we do not know/cannot calculate the derivatives. It works well for the VQE (in small parameter spaces) but requires the system to use a high shot¹ count in order to not be too greatly affected by the quantum measurement noise.

For this work, we make use of `scipy.optimize.minimize`’s option ‘`powell`’.

5.0.2 COByLA

COByLA is a trust region based, surrogate-assisted method for finding the global minimum of a function. Similar to PO, COByLA requires sufficient shots in order to not be too greatly affected by the noise; however, it is more resilient than PO. COByLA has been used for the VQE several times within the literature. Liu, et al. [38] used it as a test for their Layer-VQE (L-VQE) approach to generating ansatz, and Tim Schwägerl, et al. [39] used it to minimize a QUBO problem for reconstructing particle track results from a Large Hadron Collider (LHC). COByLA has shown good promise for this black-box optimization.

For this work, we make use of `scipy.optimize.minimize`’s option ‘`COBYLA`’.

5.0.3 Bayesian Optimization

BO is a surrogate-assisted method for finding the global minimum of a function, which models the underlying function as a Gaussian Process Regression (GPR) and then updates this model based on samples, according to Bayes’ theorem [40, 41, 42, 43] We follow the work using BO for the VQE [43], where the authors made use of the Noisy Expected Improvement acquisition

¹Number of calls of the quantum circuit to gain one expected value evaluation.

function. BO has also been used for other Variational Quantum Algorithms (VQAs), such as solving a MAXCUT problem with the Quantum Approximate Optimization Algorithm (QAOA) [44]. It has also been used for generative modelling, where combinations of Pauli gates are used rather than parameterized gates, to implement a training algorithm on the canonical Bars-and-Stripes data set. Instead of iterations, unique circuits at every step are used, see Ref. [45]. As well as these quantum examples, BO has been extensively studied in the field of hyperparameter optimization for machine learning (e.g. Ref. [46], [47], [48]).

For this work, we use a BO that we have coded. The definitions that we use are the same as presented in Ref. [43]. It uses the periodic kernel, for K parameters, defined as,

$$k^P(\theta^1, \theta^2) = \sigma^2 \prod_{i=1}^K \exp \left[\frac{-2}{l^2} \sin^2 \left(\pi \left| \frac{\theta_i^1 - \theta_i^2}{p} \right|^2 \right) \right], \quad (17)$$

where σ is the expected variance of the GP, l is the length scale over which we would expect data points to be correlated and p is the period of the underlying function.

We make use of Expected Improvement (EI) as the *acquisition function*. If we define our best guess so far to be $E_{min} = \min(E_1, \dots, E_n)$, EI is defined as,

$$a_{EI}(\theta) = \mathbb{E}_\theta[\max(0, E_{min} - E(\theta))] \quad (18)$$

Where $\mathbb{E}_\theta[\cdot]$ is evaluated over all the possible $E(\theta)$ from the surrogate model.

Expected Improvement can be used in its explicit closed form which is,

$$a_{EI}(\theta) = (E_{min} - E(\theta)) \Phi \left(\frac{E_{min} - E(\theta)}{\Delta E(\theta)} \right) + \Delta E(\theta) \phi \left(\frac{E_{min} - E(\theta)}{\Delta E(\theta)} \right). \quad (19)$$

Where Φ is the Normal cumulative distribution function and ϕ is the corresponding probability density function. We find the minimum of a_{EI} using Powell’s algorithm.

Bayesian optimization uses very few shots when compared to the other methods, as it takes into account the noise on the measurement. It is, however, extremely computationally expensive.

5.1 Mapping QUBO to the VQE

We can use the VQE to solve QUBO problems using the transformation

$$x^i \mapsto \frac{1}{2} (\sigma_0^i + \sigma_3^i), \quad (20)$$

where σ_3^i denotes the Pauli-Z spin matrix acting on the i^{th} qubit. We then find the groundstate eigenvalue of the Hamiltonian

$$H = \sum_{i,j=1}^q Q_{ij} \frac{1}{4} ((\sigma_0^i + \sigma_3^i)(\sigma_0^j + \sigma_3^j)). \quad (21)$$

We then take the parameters that gave us this groundstate, measure then the produced state in the $\sigma_3^{\otimes q}$ basis, and the state with the highest likelihood of being measured, which also meets our constraints, is our solution.

Since our Hamiltonians only consist of σ_0 and σ_3 terms, they are diagonal, hence we only need one parameter per qubit in order to map a sufficient amount of the space to receive the optimum. However, to aid the classical optimizer in the noisy space of quantum measurements, q layers of rotation and entangling *CNOT* gates will be used.

We note here that due to the density of the weight matrix Q defined in Eq. 10, the Hamiltonian is very dense, often containing q^2 nonzero terms. This means that this is a worst-case scenario in terms of stress on the VQE.

5.2 Dimensionality Expressivity Analysis

When using variational quantum circuits, something that we often look at is Dimensionality Expressivity Analysis (DEA) [49]. This is a method by which we can determine if the parameters of the gates are independent or redundant. The methodology behind this can be found in Ref. [50].

DEA is more applicable when we are looking at more complex systems, i.e. when we are not only searching for basis states and thus need to be able to span a subset of the whole space.

If we do carry out this analysis on our circuit, we see that all the parameters are non-redundant, given we do **not** start the optimization process with any parameter values equal to $(2k+1)/4 \cdot \pi$ and $(2k)/4 \cdot \pi$ for $k = \{0, 1\}$. In order to ensure this, the parameter values are uniformly randomly generated in the region $[0, 2\pi)$, using `numpy.random.uniform`.

Further details about DEA are included in Appendix A.

5.3 Conditional Value at Risk

A recent development in the usability of the VQE to solve combinatorial optimization problems is to use the Conditional Value at Risk (CVaR) technique from finance. Full details of CVaR-VQE and its advantages can be found in Ref. [51]. In this reference, CVaR was used with regard to the sample mean, but in this work we apply it to the expected value.

For a Hamiltonian containing K terms, we define the energy measurements to be H_γ . In this notation, the coefficients and normalisation with regards to the shot count have been included in H_γ . Our usual objective function is,

$$\text{VQE} \rightarrow \sum_{\gamma=1}^K H_\gamma \quad (22)$$

If we now order the measurements so that $H_k \leq H_{k+1}$, the CVaR adaptation is to then introduce a new parameter: $0 < \alpha \leq 1$ such that,

$$\text{CVaR-VQE} \rightarrow \sum_{\gamma=1}^{\lceil \alpha K \rceil} H_\gamma \quad (23)$$

In this way, we truncate our measurements to only include the terms which greatest contribute to our *minimization*. By changing this value we can more quickly find the state which gives us our minimum.

Note that the minimum value of CVaR-VQE is not necessarily equal to the minimum of the VQE, but the correct solution can be the same.

6 Design of the tests of the algorithms

We have tested the power produced by the different methods for grids of size $l_{grid} = 4$. This was the maximum system size that we could run on quantum simulators using the resources available to us. By exhaustive search of all the possibilities we found the maximum power for the grid with $l_{grid} = 4$. The WFLO problem has many local optima, so we run each algorithm 36 times from different initial conditions. To check the scaling of the performance of the different algorithms, we run the quantum simulators with

$l_{grid} = 2$ to $l_{grid} = 4$, and the classical optimizers with $l_{grid} = 2$ to $l_{grid} = 10$.

Justification for 36 being a sufficient sample size can be found in Appendix B.

6.1 Test model

In order to test these different algorithms, we will use the following parameter values (the Mosetti benchmark cases are from Ref. [6]),

$$D = \text{Mosetti second benchmark case}$$

$$x = 1, r = 1.5, m = 4, \xi = 0$$

where x is the maximum distance that is affected by the wake and r is the radius with which it spreads out per unit distance away from the turbine causing the wake. m is the maximum number of turbines allowed, and ξ is the minimum distance allowed between turbines. In this work we have set ξ to zero so that turbines may be on adjacent sites. Due to our limited ability to simulate large quantum systems, we will be looking at a very small grid with $l_{grid} = 4$. This means that there will be 16 possible positions for the turbines, and so will need 16 qubits.

Mosetti's second benchmark case is defined as:

$$D = \{\{10k, 12ms^{-1}, 1/36\}\}_{k=0}^{36} \quad (24)$$

$$= \{\{0, 12ms^{-1}, 0.028\}, \dots, \{350, 12ms^{-1}, 0.028\}\} \quad (25)$$

This can be visualized as the distribution shown in Fig. 7. Here, each box represents a windspeed of $12ms^{-1}$.

6.2 Computational details of simulations

The quantum simulations were run using IBM's `qiskit:0.39.2` package for python [15].

Both PO and COBYLA were utilized via `scipy.optimize.minimize` - using the methods 'powell' and 'COBYLA' respectively as part of the `scipy` library [27]. BO was coded up for this work, using a heavily modified version of the code used in [52], where we no longer make use of `sklearn.gaussian_process.GaussianProcessRegressor`. This had problems when using the periodic kernel, and so a self-made algorithm was required.

SA was implemented using the package `pyqubo` [53, 54]. This library has an inbuilt annealer by using

`neal.SimulatedAnnealingSampler()` once the system has been turned into a QUBO problem via `model.to_qubo()`. `Pyqubo` has a wrapper to the D-Wave Ocean SDK, so we are ready to run on the D-wave systems.

Gurobi was implemented using the `gurobipy` package [28]. Gurobi was used with an academic license through the University of Plymouth.

As well as the previously mentioned methods, we also investigated `qiskit`'s built-in VQE QUBO solver. This uses their `SamplingVQE` function, and then performs a `MinimumEigenOptimizer` over it in order to find the groundstate. The ansatz used is the same as the one described in Sec. 5.1. This method acts as a control, as it is the VQE method with no statistical noise present. We use COBYLA as the classical optimization routine for this method. Details about this method can be found in Ref. [55].

All of the simulations for the results provided below were run using the computational facilities of the High Performance Computing Centre located at the University of Plymouth [56]. The CPUs were dual node Intel Xeon E5-2683v4 CPU each with 16 cores. Task farming was used to run the simulations using all the cores.

7 Results

7.1 How solutions are selected

For both the Gurobi solver and quantum annealing, the selection of the solution is trivial as it will be the final binary string on which the algorithm lands.

However, for the VQE this is less obvious. Once the optimization process has finished, we are given a set of parameters which corresponds to the quantum circuit gates. We then input these values back into the circuit, and measure in the σ_3 basis on all qubits; this produces a distribution over all possible basis states. From this distribution, we then take the state which has the highest probability and only has four $|1\rangle$ terms. We only accepted solutions which have exactly the correct number of turbines, which is four for our test case.

7.2 Discussion about the degeneracy of solutions

As for many other QUBO problems, there are many degenerate optimal solutions. For the case $l_{grid} = 4$, there are 79 unique optimal solutions out of 1820 binary vectors which meet the constraints and 65536 total possible binary strings of length 16. A list of all optimal solutions can be found in Appendix C. Each of these optimal solutions gives us a power output of 2304.0 kW.

Some of these degeneracies are due to the fact that there are optimal solutions for the conditions in smaller system sizes. The ‘smaller’ solutions can then be embedded in the larger systems. No solutions exists with output 2304.0 kW for $l_{grid} \leq 2$.




Figure 7: 2^{nd} Mosetti benchmark case wind regime distribution. Each bar represents a windspeed of 12 m.s^{-1} .




Figure 8: Heatmap of $\langle X \rangle$ values for the degenerate solutions with maximum power. All the values sum to 4 (the number of turbines).

Another way in which we could study a structure in the optimal solutions is to look at their average placement locations. If we label a solution

matrix to be X_*^i , this is a binary matrix which is representative of the solution grid. We can map from solutions by using the labelling as in Fig. 2. We can then calculate the average placement as:

$$\langle X \rangle = \frac{1}{\#\text{solutions}} \sum_i X_*^i \quad (26)$$

The results from this can be seen in Fig. 8, note here that the sum of these should equal four as this is the maximum number of turbines, but the numbers have been limited to three d.p. for ease of viewing. From this, we see that there are more optimal solutions with turbines in the outer corners of the grid. This makes sense physically with the wind regime used, as it maximizes the distance between any two turbines and thus minimizes any possible wake effects.

7.3 Discussion of the results

We found that the Powell optimization method in the simulations produced results in a reasonable amount of time for $l_{grid} = 2$ and $l_{grid} = 3$ cases. However, the Powell method failed at doing this for $l_{grid} = 4$. This was due to requiring a large shot count and also requiring many function evaluations to be effective. We only present results using the two surrogate model-based methods, COByLA and BO, and the two classical algorithms.

Fig. 9 reports the results for the power output from the different algorithms investigated for the $l_{grid} = 4$ case. Each of the methods was run 36 times and the boxplots show the spread of these results. The x -axis shows the names of the different methods, the number under the name of the VQE method is the CVaR α value.

From Fig. 9 we can draw a few conclusions. The Gurobi method with the default starting guess outperforms all the other methods, consistently reaching the optimal solution every time. This is also true for the noiseless COBYLA-VQE method from `qiskit`. We also see that the different VQE methods provide similar quality of results to one another, which are also comparable to SA.

We see that each of the COBYLA-CVaR methods is able to reach the optimal solution. We also see that using $\alpha < 1$ has little effect on the quality of results. We see that the Bayesian-CVaR has very comparable results.




Figure 9: Each value along the x -axis represents the different methods. The box plots show the power obtained from running the problems 36 times. The red dashed line is the optimal solution power output, found through an exhaustive list search. The Gurobi method and Qiskit Optimization may look empty, but there is a line at the optimal - it was able to achieve this every run. The circle points which are not within the boxes are outlier points.

Method	Average output result (% of optimal)
COByLA-1.00	95.5
COByLA-0.75	94.7
COByLA-0.50	94.5
COByLA-0.25	95.2
Bayes-1.00	93.3
Bayes-0.75	94.2
Bayes-0.50	94.6
Bayes-0.25	93.4

Table 1: Average output power from solutions of different VQE-based methods as a percentage of the optimal power output.

Tab. 1 shows the average power output of each of the methods as a percentage of the optimal power output. We see that each of these methods do perform very well on average; with the lowest average output being 93.3% of the maximum.

7.4 Timing the results

As well as looking at the quality of results, we are also interested in how long the algorithms take to achieve their results.

In order to test this, we have used Python’s `time` package to take a total time measurement from start to end of the whole process. The results from this can be shown in Tab. 2

From these results, we can easily see that the simulated quantum system takes far more time than the other two algorithms. However, what is of more interest to us is the scalability of these results, i.e. how much slower they get as we increase the system size. We can study this by looking at the coefficients of LogLog plots of the timings.

The time to solve the QUBO problem is expected to depend on the system size exponentially. However, as we only have three data points, we parameterize the time to solution as a polynomial with coefficient α .

$$\text{Time to solve} \propto (V)^\alpha$$

Where $V = l_{grid}^2$ is the system volume. Fig. 10 shows the log of time taken to solution versus the log of the system volume. The timings are in Table 3. The results from all four algorithms in Fig. 10 have similar slopes.

We find that: Gurobi $\sim O(V^{4.24})$, SA $\sim O(V^{4.51})$, COByLA $\sim O(V^{4.57})$ and Bayesian

$\sim O(V^{3.89})$. These values are calculated by linearly fitting the LogLog data from timing each algorithm twelve times and taking the mean. This is interesting as while Gurobi and SA are the true algorithms, the quantum method is a simulation, and thus we would expect that the real quantum circuit would reduce this time.

Very interestingly, using Bayesian optimization even on the simulator scales better than all the other methods, having the lowest polynomial scaling; this does come with the caveat that the pre-factor is much larger than the other systems. We can look at this more closely by searching for the intersection point of the Bayesian and Gurobi timing lines, this is shown in Fig. 11. We can estimate that the simulated Bayesian-CVaR method would be faster than the Gurobi method for a system volume of $\approx 10^{15}$, which is a square of sides length ≈ 31622776 .




Figure 11: Extrapolation of time taken for Gurobi and Bayesian Optimization in order to find the point at which Bayesian would be faster.

7.5 Conclusions

The VQE-based solver of the WFLO problem finds good solutions when a sufficient number of measurements have been performed. The Gurobi optimizer always finds the optimal solution, and it outperforms the noisy VQE-based method. This is not surprising, as Gurobi has been developed to specifically solve linear and quadratic programming problems, of which QUBO is a special case. We have only investigated small system sizes and it would be interesting to compare the results from the VQE algorithm running on a quantum computer. There are many benchmark-

l_{grid}	Gurobi [s]	Simulated Annealing [s]	COByLA-CVaR [s]	BO-CVaR [s]
2	-	-	19.74	250.11
3	0.09	0.04	561.78	5741.43
4	1.11	0.85	11546.10	55454.77
5	3.91	3.31	-	-
6	14.51	11.49	-	-
7	107.11	106.62	-	-
8	374.72	377.49	-	-
9	983.27	985.57	-	-
10	2240.78	2220.44	-	-

Table 2: Average time in seconds taken for the different algorithms.




Figure 10: LogLog plots of the time taken for the different methods over different system sizes.

Method	Gradient	Intercept
Gurobi	4.24	-5.18
Simulated Annealing	4.51	-5.67
CVaR-VQE using COByLA	4.57	-1.50
CVaR-VQE using BO	3.89	0.04

Table 3: Coefficients of linear fits for LogLog timings.

ing studies that compare different algorithms to solve the WFLO problem use a grid with 100 points, which would require 100 qubits in our approach. Solving the WFLO on a quantum computer would require many more than 100 qubits when error correction is used.

It is interesting to see that the scaling of the time taken with system size had similar exponents between the Gurobi and the VQE-based method running on the quantum simulator. From our results, the performance of the VQE method critically depends on the performance of the optimizer that runs on the classical computer and this will have to be tuned for good performance.

For WFLO problems, it is useful to use multi-objective optimisation [7]. For example, maximize the power and minimize the cost. The QUBO formalism has recently been extended to include multiple objectives [57]. This would be an interesting area to explore in the future.

We are encouraged by the results we have obtained on the quantum computing simulator to study the algorithms on real quantum computers. Next, we will study the resilience of the implemented methods against quantum errors on the simulators. It will be interesting to compare the performance on circuit-based quantum computers to those from adiabatic quantum systems.

Recent work has studied the iteration complexity of Variational Quantum Algorithms in a noisy environment [58].

There has also been a generalization of CVaR-VQE proposed, called the Filter-VQE (F-VQE). F-VQE uses a technique based on filtering operators to achieve faster and more reliable convergence to the optimal solution[59].

When dealing with high gate count variational circuits, we encounter exponentially vanishing gradients or *barren plateaus* [60]. There has been recent work [61] in methods to overcome this problem by utilizing parallel running optimization routines (particles), and moving these particles when they encounter regions of vanishing gradient or when the gradient is dominated by noise. We plan to investigate these new techniques in the future, as well as studying the influence of quantum errors.

7.6 Acknowledgements

We thank Julian Stander for the discussions about Bayesian optimization. This work

was carried out using the computational facilities of the High Performance Computing Centre located at the University of Plymouth – <https://www.plymouth.ac.uk/about-us/university-structure/faculties/science-engineering/hpc>.

References

- [1] P. Veers *et al.*, *Science* **366**, eaau2027 (2019).
- [2] A. Giani and Z. Eldredge, *SN Computer Science* **2**, 393 (2021).
- [3] C. Berger *et al.*, arXiv preprint arXiv:2107.05362 (2021).
- [4] A. Ajagekar and F. You, *Renewable and Sustainable Energy Reviews* **165**, 112493 (2022).
- [5] A. Abbas *et al.*, arXiv preprint arXiv:2312.02279 (2023).
- [6] G. Mosetti, C. Poloni, and B. Diviacco, *Journal of Wind Engineering and Industrial Aerodynamics* **51**, 105 (1994).
- [7] P. L. Manikowski, D. J. Walker, and M. J. Craven, *Journal of Marine Science and Engineering* **9**, 1376 (2021).
- [8] Vu Phong Energy Group JSC, The world's biggest wind farms, <https://vuphong.com/the-worlds-biggest-wind-farms/>, 2022 (accessed 01/11/2023).
- [9] A. Senderovich, J. Zhang, E. Cohen, and J. C. Beck, *IEEE Access* **10**, 78044 (2022).
- [10] J. Gambetta, Quantum-centric supercomputing: The next wave of computing, <https://research.ibm.com/blog/next-wave-quantum-centric-supercomputing>, Accessed: 27-06-2023.
- [11] A. Rajak, S. Suzuki, A. Dutta, and B. K. Chakrabarti, *Philosophical Transactions of the Royal Society A* **381**, 20210417 (2023).
- [12] C. McGeoch and P. Farré, *D-Wave systems* (2022).
- [13] C. McGeoch, P. Farré, and K. Boothby, *D-Wave systems* (2022).
- [14] S. Matsubara *et al.*, Digital annealer for high-speed solving of combinatorial optimization problems and its applications, in *2020 25th Asia and South Pacific Design Automation Conference (ASP-DAC)*, pp. 667–672, IEEE, 2020.
- [15] Qiskit contributors, *Qiskit: An open-source framework for quantum computing*, 2023.
- [16] J. Preskill, *Quantum* **2**, 79 (2018).
- [17] E. Zahedinejad and A. Zaribafiyani, arXiv preprint arXiv:1708.05294 (2017).
- [18] T. Albash and D. A. Lidar, *Physical Review X* **8**, 031016 (2018).
- [19] A. Pearson, A. Mishra, I. Hen, and D. A. Lidar, *npj Quantum Information* **5**, 107 (2019).
- [20] R. Yaacoby *et al.*, *Physical Review E* **105**, 035305 (2022).
- [21] V. Martin-Mayor and I. Hen, *Scientific reports* **5**, 15324 (2015).
- [22] P. Y. Zhang, D. A. Romero, J. C. Beck, and C. H. Amon, Solving wind farm layout optimization with mixed integer programming and constraint programming, in *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*, pp. 284–299, Berlin, Heidelberg, 2013, Springer.
- [23] S. Donovan, *Proceedings of the Annual ORSNZ Conference* **40** (2005).
- [24] A. Crippa *et al.*, (2023), arXiv:2304.01690.
- [25] G. Kochenberger *et al.*, *Journal of combinatorial optimization* **28**, 58 (2014).
- [26] M. Khanali, S. Ahmadzadegan, M. Omid, F. Keyhani Nasab, and K. W. Chau, *International Journal of Energy and Environmental Engineering* **9**, 399 (2018).
- [27] T. do Cuoto, B. Farias, A. Carlos G.C. Diniz, and M. Vinicius G. de Moraes, *10th World Congress on Structural and Multidisciplinary Optimization* (2013).
- [28] Gurobi Optimization LLC, *Gurobi optimizer reference manual*, <https://www.gurobi.com>, 2023, Accessed: 27-06-2023.
- [29] O. Şeker, N. Tanoumand, and M. Bodur, *Applied Soft Computing* **127**, 109367 (2022).
- [30] K. Nakano *et al.*, arXiv preprint arXiv:2207.03069 (2022).
- [31] S. Kirkpatrick, C. D. Gelatt Jr, and M. P. Vecchi, *science* **220**, 671 (1983).
- [32] M. Pincus, *Operations Research* **18**, 1225 (1970).
- [33] Z. Wang, X. Geng, and Z. Shao, *Journal of Computational and Theoretical Nanoscience* **6**, 1680 (2009).

- [34] E. Rodriguez-Tello, J.-K. Hao, and J. Torres-Jimenez, *Computers & Operations Research* **35**, 3331 (2008).
- [35] K. Tole, R. Moqa, J. Zheng, and K. He, *Computers & Industrial Engineering* **176**, 109004 (2023).
- [36] A. Peruzzo *et al.*, *Nature Communications* **5** (2014).
- [37] J. Tilly *et al.*, *Physics Reports* **986**, 1 (2022).
- [38] X. Liu *et al.*, *IEEE Transactions on Quantum Engineering* **3**, 1 (2022).
- [39] T. Schwägerl *et al.*, *Frontiers of Physics* **18**, 21308 (2023).
- [40] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, *New J. Phys.* **18**, 023023 (2016), arXiv:1509.04279.
- [41] N. Klco *et al.*, *Phys. Rev. A* **98**, 032331 (2018), arXiv:1803.03326.
- [42] R. Santagati *et al.*, *Science advances* **4**, eaap9646 (2018).
- [43] G. Iannelli and K. Jansen, arXiv preprint arXiv:2112.00426 (2021).
- [44] J. S. Otterbach *et al.*, arXiv preprint arXiv:1712.05771 (2017).
- [45] D. Zhu *et al.*, *Science advances* **5**, eaaw9918 (2019).
- [46] J. Snoek, H. Larochelle, and R. P. Adams, *Advances in neural information processing systems* **25** (2012).
- [47] V. Nguyen, Bayesian optimization for accelerating hyper-parameter tuning, in *2019 IEEE second international conference on artificial intelligence and knowledge engineering (AIKE)*, pp. 302–305, IEEE, 2019.
- [48] J. Bergstra, R. Bardenet, Y. Bengio, and B. Kégl, *Advances in neural information processing systems* **24** (2011).
- [49] L. Funcke, T. Hartung, K. Jansen, S. Kühn, and P. Stornati, *Quantum* **5**, 422 (2021).
- [50] L. Funcke *et al.*, Dimensional expressivity analysis, best-approximation errors, and automated design of parametric quantum circuits, 2021, arXiv:2111.11489.
- [51] P. K. Barkoutsos, G. Nannicini, A. Robert, I. Tavernelli, and S. Woerner, *Quantum* **4**, 256 (2020).
- [52] J. Brownlee, How to implement bayesian optimization from scratch in python, <https://machinelearningmastery.com/what-is-bayesian-optimization/>, Accessed: 27-06-2023.
- [53] M. Zaman, K. Tanahashi, and S. Tanaka, *IEEE Transactions on Computers* (2021).
- [54] K. Tanahashi, S. Takayanagi, T. Motohashi, and S. Tanaka, *Journal of the Physical Society of Japan* **88**, 061010 (2019).
- [55] Q. O. D. Team, Improving variational quantum optimization using cvar, https://qiskit.org/ecosystem/optimization/tutorials/08_cvar_optimization.html, Accessed: 27-11-2023.
- [56] HPC team UoP, High performance computing (hpc), <https://www.plymouth.ac.uk/about-us/university-structure/faculties/science-engineering/hpc>, Accessed: 06-11-2023.
- [57] M. Ayodele, R. Allmendinger, M. López-Ibáñez, and M. Parizy, Multi-objective qubo solver: Bi-objective quadratic assignment problem, in *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 467–475, 2022.
- [58] V. Kungurtsev, G. Korpas, J. Marecek, and E. Y. Zhu, (2023), arXiv:2209.10615.
- [59] D. Amaro *et al.*, *Quantum Science and Technology* **7**, 015021 (2022).
- [60] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, *Nature Communications* **12** (2021).
- [61] D. Mastropietro, G. Korpas, V. Kungurtsev, and J. Marecek, (2023), arXiv:2311.18090.

Appendix

A Dimensionality Expressivity Analysis

A way in which we can analyse the system and improve our results is to carry out expressivity analysis on the circuit matrix. For an objective function $F(\vec{\theta})$, we can then perturb θ_k by $\delta\theta_k$ to produce $F(\vec{\theta} + \hat{e}_k\delta\theta_k) = f$. If we can also have $F(\vec{\theta} + \sum_{i \neq k} \hat{e}_i\delta\theta_i) = f$, θ_k is redundant. In order to do this analysis, we are using the method that is described by Lena Funcke, et al in Ref. [50]. The method is as follows:

1. This can be checked by considering the real partial Jacobians J_k of C (C is the matrix of the circuit applied to the state $|0\rangle^{\otimes q}$):

$$J_k(\theta) = \begin{pmatrix} Re(\partial_1 C) \dots Re(\partial_k C) \\ Im(\partial_1 C) \dots Im(\partial_k C) \end{pmatrix}$$

Here it is key to understand that $\partial_k C$ is itself a $2^q \times 1$ vector, and thus J_k is a $2^{q+1} \times k$ matrix, where we start with $k = 1$ for the first parameter.

2. We then check the rank of the matrix J_k for each k , and if adding a new parameter does not increase the rank, the parameter must be redundant.

In order to efficiently check the rank of J_k , we consider the matrix $S_k = J_k^* J_k$. Thus if we check that S_k is invertible we know that all the parameters are independent (i.e., not redundant).

B Justification for thirty-six samples

When gathering results for this work, we chose 36 samples for each of the methods. In this section we will look at one particular method, being COBYLA with the CVaR $\alpha = 0.25$. We take a much larger sample (284) and compare the averages. By doing this we can see that 36 samples shows similar features and is thus sufficient to draw conclusions from.

The initial 36 samples taken were:

2136.54, 2118.96, 2202.69, 2220.27, 2220.27, 2220.27, 2268.84,
2220.27, 2286.42, 2220.27, 2136.54, 2118.96, 2118.96, 2220.27,
2101.38, 2202.69, 2220.27, 2286.42, 2202.69, 2286.42, 2304.0,
2118.96, 2286.42, 2118.96, 2118.96, 2136.54, 2136.54, 2268.84,
2118.96, 2118.96, 2101.38, 2286.42, 2220.27, 2286.42, 2202.69,
2118.96

The extra 284 samples taken were:

2251.25, 2202.69, 2185.11, 2185.11, 2136.54, 2136.54, 2286.42,
2118.96, 2268.84, 2304.0, 2202.69, 2136.54, 2220.27, 2268.84,
2101.38, 2118.96, 2268.84, 2220.27, 2220.27, 2202.69, 2202.69,
1933.92, 2118.96, 2202.69, 2268.84, 2220.27, 2220.27, 2220.27,
2268.84, 2220.27, 2185.11, 2136.54, 2118.96, 2202.69, 2251.25,
2304.0, 2052.82, 2118.96, 2286.42, 2136.54, 2118.96, 2220.27,
2118.96, 2118.96, 2268.84, 2304.0, 2101.38, 2268.84, 2286.42,
2035.23, 2304.0, 2202.69, 2118.96, 2304.0, 2202.69, 2017.65,
2268.84, 2220.27, 2136.54, 2304.0, 2202.69, 2118.96, 2268.84,
2136.54, 2136.54, 2202.69, 2118.96, 2035.23, 2202.69, 2304.0,
2118.96, 2220.27, 2136.54, 2185.11, 2118.96, 2220.27, 2251.25,

2220.27, 2304.0, 2202.69, 2202.69, 2220.27, 2268.84, 2304.0,
2185.11, 2136.54, 2202.69, 2118.96, 2185.11, 2220.27, 2118.96,
2118.96, 2268.84, 2220.27, 2185.11, 2136.54, 2251.25, 2304.0,
2268.84, 2202.69, 2220.27, 2185.11, 2202.69, 2118.96, 2220.27,
2220.27, 2220.27, 2286.42, 2268.84, 2220.27, 2185.11, 2202.69,
2101.38, 2202.69, 2220.27, 2202.69, 2304.0, 2304.0, 2220.27,
2202.69, 2220.27, 2286.42, 2220.27, 2118.96, 2118.96, 2220.27,
2202.69, 2035.23, 2118.96, 2202.69, 2286.42, 2118.96, 2101.38,
2118.96, 2202.69, 2304.0, 2136.54, 2304.0, 2185.11, 2202.69,
2286.42, 2101.38, 2202.69, 2220.27, 2220.27, 2220.27, 2220.27,
2220.27, 2136.54, 2220.27, 2101.38, 2220.27, 2286.42, 2118.96,
2220.27, 2220.27, 2286.42, 2220.27, 2136.54, 2136.54, 2268.84,
2101.38, 2202.69, 2101.38, 2035.23, 2220.27, 2136.54, 2202.69,
2136.54, 2286.42, 2185.11, 2286.42, 2220.27, 2118.96, 2118.96,
2220.27, 2202.69, 2202.69, 2118.96, 2118.96, 2118.96, 2202.69,
2220.27, 2202.69, 2220.27, 2202.69, 2286.42, 2118.96, 2202.69,
2136.54, 2220.27, 2268.84, 2118.96, 2251.25, 2286.42, 2136.54,
2136.54, 2286.42, 2286.42, 2220.27, 2202.69, 2286.42, 2136.54,
2220.27, 2202.69, 2220.27, 2118.96, 2202.69, 2304.0, 2220.27,
2202.69, 2304.0, 2286.42, 2304.0, 2286.42, 2101.38, 2220.27,
2101.38, 2220.27, 2202.69, 2202.69, 2220.27, 2035.23, 2304.0,
2304.0, 2202.69, 2118.96, 2185.11, 2202.69, 2220.27, 2118.96,
2101.38, 2202.69, 2268.84, 2304.0, 2202.69, 2017.65, 2118.96,
2286.42, 1933.92, 2304.0, 2118.96, 2202.69, 2286.42, 2220.27,
2136.54, 2136.54, 2185.11, 2017.65, 2185.11, 2220.27, 2202.69,
2118.96, 2136.54, 2118.96, 2220.27, 2017.65, 2251.25, 2136.54,
2202.69, 2035.23, 1933.92, 2136.54, 2017.65, 2220.27, 2118.96,
2101.38, 2220.27, 2286.42, 2118.96, 2286.42, 2220.27, 2286.42,
2220.27, 2286.42, 2202.69, 2268.84, 2101.38, 2251.25, 2220.27,
2035.23, 2286.42, 2202.69, 2202.69

The sample mean of the 36 initial samples is: 2193.1, and the sample mean of the extra is: 2192.4. We can see that these are similar values (0.03% difference), and so we are able to draw conclusions about the method's effectiveness *on average* by only taking 36 samples. This is very useful as minimizing time and resource waste is an important factor when carrying out simulations.

C Degeneracy of solutions

Below is a list of all possible optimal solutions for the problem defined in Sec. 6.1.

(0, 0, 0, 0, 0, 1, 0, 1, 0, 0, 0, 0, 0, 1, 0, 1),
(0, 0, 0, 0, 0, 1, 0, 1, 0, 0, 0, 0, 1, 0, 0, 1),
(0, 0, 0, 0, 0, 1, 0, 1, 0, 0, 0, 0, 1, 0, 1, 0),
(0, 0, 0, 0, 1, 0, 0, 1, 0, 0, 0, 0, 0, 1, 0, 1),
(0, 0, 0, 0, 1, 0, 0, 1, 0, 0, 0, 0, 1, 0, 0, 1),
(0, 0, 0, 0, 1, 0, 0, 1, 0, 0, 0, 0, 1, 0, 1, 0),
(0, 0, 0, 0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 1, 0, 1),
(0, 0, 0, 0, 1, 0, 1, 0, 0, 0, 0, 0, 1, 0, 0, 1),
(0, 0, 0, 0, 1, 0, 1, 0, 0, 0, 0, 0, 1, 0, 1, 0),
(0, 0, 0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1),
(0, 0, 0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 1, 0, 0, 1),
(0, 0, 0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0),

(0, 0, 0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0, 0),
 (0, 0, 0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 1, 1, 0, 0, 0),
 (0, 0, 0, 1, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1),
 (0, 0, 0, 1, 1, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 1),
 (0, 0, 0, 1, 1, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0),
 (0, 0, 0, 1, 1, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0, 0),
 (0, 0, 0, 1, 1, 0, 0, 0, 0, 0, 0, 0, 1, 1, 0, 0, 0),
 (0, 0, 0, 1, 1, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0, 0, 0),
 (0, 0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1),
 (0, 0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 1),
 (0, 0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0),
 (0, 0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0, 0),
 (0, 0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 1, 1, 0, 0, 0),
 (0, 0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0, 0, 0),
 (0, 1, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 0, 1, 0, 1),
 (0, 1, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 1, 0, 0, 1),
 (0, 1, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 1, 0, 1, 0),
 (0, 1, 0, 0, 0, 0, 0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 1),
 (0, 1, 0, 0, 0, 0, 0, 1, 1, 0, 0, 0, 0, 0, 0, 0, 1),
 (0, 1, 0, 0, 0, 0, 0, 1, 1, 0, 0, 0, 0, 0, 0, 1, 0),
 (0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1),
 (0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 1),
 (0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 1),
 (0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0),
 (0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0, 0),
 (0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 1, 1, 0, 0, 0),
 (0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0, 0, 0),
 (0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 1),
 (0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0, 0, 0, 0),
 (0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0, 0, 0, 0),
 (1, 0, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 0, 1, 0, 1),
 (1, 0, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 1, 0, 0, 1),
 (1, 0, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 1, 0, 1, 0),
 (1, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 1),
 (1, 0, 0, 0, 0, 0, 0, 1, 1, 0, 0, 0, 0, 0, 0, 0, 1),
 (1, 0, 0, 0, 0, 0, 0, 1, 1, 0, 0, 0, 0, 0, 0, 1, 0),
 (1, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1),
 (1, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 0, 1, 0, 0, 1),
 (1, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0),
 (1, 0, 0, 0, 0, 0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 1),
 (1, 0, 0, 0, 0, 0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 1, 0),
 (1, 0, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1),
 (1, 0, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 1),
 (1, 0, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0),
 (1, 0, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0, 0),
 (1, 0, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 1, 1, 0, 0, 0),
 (1, 0, 0, 1, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0, 0, 0, 0),
 (1, 0, 0, 1, 0, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 1),
 (1, 0, 0, 1, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0, 0, 0, 0),

(1, 0, 0, 1, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 0, 1),
(1, 0, 0, 1, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 1, 0),
(1, 0, 0, 1, 0, 0, 0, 0, 1, 0, 0, 1, 0, 0, 0, 0),
(1, 0, 0, 1, 0, 0, 0, 0, 1, 0, 1, 0, 0, 0, 0, 0),
(1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1),
(1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0),
(1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1),
(1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0),
(1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0),
(1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 1, 0, 0),
(1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0, 0),
(1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0, 1),
(1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0, 0, 0),
(1, 0, 1, 0, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 1),
(1, 0, 1, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 1, 0),
(1, 0, 1, 0, 0, 0, 0, 0, 1, 0, 0, 1, 0, 0, 0, 0),
(1, 0, 1, 0, 0, 0, 0, 0, 1, 0, 1, 0, 0, 0, 0, 0)