

Quantum optimization with linear Ising penalty functions for customer data science

Puya Mirkarimi^{1,*}, Ishaan Shukla¹, David C. Hoyle², Ross Williams² and Nicholas Chancellor^{1,3,†}

¹*Department of Physics, Durham University, Durham DH1 3LE, United Kingdom*

²*dunnhumby, 184 Shepherds Bush Road, London W6 7NL, United Kingdom*

³*School of Computing, Newcastle University, 1 Science Square, Newcastle upon Tyne NE4 5TG, United Kingdom*



(Received 17 May 2024; accepted 21 October 2024; published 5 December 2024)

Constrained combinatorial optimization problems, which are ubiquitous in industry, can be solved by quantum algorithms such as quantum annealing (QA) and the quantum approximate optimization algorithm (QAOA). In these quantum algorithms, constraints are typically implemented with quadratic penalty functions. This penalty method can introduce large energy scales and make interaction graphs much more dense. These effects can result in worse performance of quantum optimization, particularly on near-term devices that have sparse hardware graphs and other physical limitations. In this work, we consider linear Ising penalty functions, which are applied with local fields in the Ising model, as an alternative method for implementing constraints that makes more efficient use of physical resources. We study the behavior of the penalty method in the context of quantum optimization for customer data science problems. Our theoretical analysis and numerical simulations of QA and the QAOA indicate that this penalty method can lead to better performance in quantum optimization than the quadratic method. However, the linear Ising penalty method is not suitable for all problems as it cannot always exactly implement the desired constraint. In cases where the linear method is not successful in implementing all constraints, we propose that schemes involving both quadratic and linear Ising penalties can be effective.

DOI: [10.1103/PhysRevResearch.6.043241](https://doi.org/10.1103/PhysRevResearch.6.043241)

I. INTRODUCTION

Quantum optimization algorithms, such as quantum annealing (QA) [1–3] and the quantum approximate optimization algorithm (QAOA) [4], are being explored as a potential path to quantum advantage in the near future [5,6]. There is increasing interest in applying these quantum algorithms to industry problems [7–12], including commerce and retail problems [13,14]. A common feature of optimization problems found in industry is that they are often highly constrained [12]. Problems in commerce and retail settings can involve tens to hundreds of constraints [15,16], which result from both strategic and operational considerations. The large number of constraints influences the approach a researcher might take in tackling an optimization problem. Therefore, if quantum optimization is to find real value in a commercial setting, it must be able to handle and thrive in problems significantly affected by or even dominated by constraints. This motivates the research of quantum optimization on problems that include multiple constraints.

In QA and the QAOA, the standard approach to encoding constraints is to add quadratic penalty functions to a

problem's objective function [17,18]. Quadratic penalty functions can make problems much more dense and introduce large energy scales to energy landscapes. This can negatively impact the performance of a quantum optimizer, especially on near-term devices that have physical limitations [19]. Overcoming these physical limitations requires new strategies that address issues that are not typically a concern in optimization problems. A particular issue that needs to be overcome in quantum annealing is that problems often have to be mapped to quasiparallel hardware graphs, using strategies such as minor embedding [20] or parity encoding [21,22]. Some success has been realized with new encoding strategies, such as domain-wall encoding [23–25]. Domain-wall encoding works well to reduce connectivity when applied to one-hot encodings, but is not efficient for other encodings, while the linear penalty strategy we consider here is.

We consider a penalty method involving only linear Ising terms and study its suitability in quantum optimization for customer data science problems. This penalty method has previously been suggested in other contexts [8,26,27]. Unlike the quadratic method, the linear Ising penalty method does not change a problem's connectivity. Furthermore, the energy scales introduced by the linear method are often smaller than for the quadratic method. Because of these reductions in resource usage, theoretical arguments can be made for better performance of quantum optimization with the linear method than the quadratic method. However, while the quadratic penalty method can always exactly implement a desired constraint, there is no such guarantee for the linear Ising penalty method. Thus, the linear Ising penalty method is more applicable to certain problems than others. We find

*Contact author: puya.mirkarimi@durham.ac.uk

†Contact author: nicholas.chancellor@gmail.com

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

that the customer data science problems we consider are well suited for the linear method.

In this work, we perform a numerical analysis that pursues two objectives: first, to understand the behavior of linear Ising penalties, and second, to assess their impact on the performance of quantum optimizer. We study how changing the value of the penalty strength parameter affects the constraint that a linear Ising penalty implements. Strategies for finding a penalty strength that implements the desired constraint are discussed. For problems with multiple constraints, we find that using a combination of quadratic and linear Ising penalties can be beneficial in cases where the linear method is not successful for all constraints. To compare the performance of the penalty methods, we have run simulations of a customer data science problem on QA and the QAOA using both penalty methods. These simulations were limited to fully connected problems with up to 18 variables and did not include the effects of physical limitations that near-term quantum devices typically have. Therefore, many of the theoretical benefits of the linear Ising penalty method do not apply to the setting represented by these simulations. Nevertheless, we observe modest improvements in performance when switching from the quadratic penalty method to the linear Ising penalty method in simulations of both algorithms. These findings are complementary to the experimental results in Ref. [28], where more substantial improvements are observed when solving larger customer data science problems on a real quantum device.

This paper is organized as follows. In Sec. II, we outline the relevant prior work, the promotion cannibalization problems we consider, and the linear Ising penalty method. In Sec. III, we describe our numerical methods. The results of this study are presented in Sec. IV, where we study the behavior of linear Ising penalties through a numerical analysis, and Sec. V, where we compare the performance of quantum optimization using quadratic and linear Ising penalties in simulation. Finally, we summarize our findings in Sec. VI.

II. BACKGROUND

In this section, we introduce some key concepts and provide a summary of prior works that underpin our study. Section II A gives a description of QA and the QAOA, which are the two quantum optimization algorithms we consider. Section II B describes the computational problems that we base our numerical analysis on. Section II C reviews the quadratic penalty method for encoding constraints and introduces the linear Ising penalty method.

A. Quantum optimization

QA and the QAOA are quantum algorithms that are particularly suitable for solving combinatorial optimization problems that can be expressed as the quadratic unconstrained binary optimization (QUBO) problem

$$\text{find } \arg \min_{\mathbf{x}} f(\mathbf{x}) = \sum_{i=1}^n a_i x_i + \sum_{i=1}^{n-1} \sum_{j=i+1}^n b_{i,j} x_i x_j. \quad (1)$$

The real-valued linear term coefficients \mathbf{a} and quadratic term coefficients \mathbf{b} characterize an instance of the problem. An assignment of values to the input vector $\mathbf{x} \in \{0, 1\}^n$ is called a solution, and we are tasked with finding an optimal solution, which is a solution that minimizes the objective function $f(\mathbf{x})$. We note that elsewhere in the literature, the QUBO problem is typically expressed in terms of a single upper triangular matrix $Q \in \mathbb{R}^{n \times n}$, where $Q_{i,j} = b_{i,j} \forall j \neq i$ and $Q_{i,i} = a_i$.

Through the mapping $x_i \mapsto (1 - \sigma_i^z)/2$, the QUBO problem is equivalent to the problem of finding the ground state of the Ising Hamiltonian

$$H_P = \sum_{i=1}^n h_i \sigma_i^z + \sum_{i=1}^{n-1} \sum_{j=i+1}^n J_{i,j} \sigma_i^z \sigma_j^z. \quad (2)$$

Here, $\sigma_i^z = \mathbb{1}^{\otimes i-1} \otimes \sigma_z \otimes \mathbb{1}^{\otimes n-i}$ is the Pauli operator σ_z acting on qubit i and identities acting on all other qubits. The couplings \mathbf{J} and local fields \mathbf{h} can be derived from the QUBO coefficients \mathbf{a} and \mathbf{b} through the relations

$$J_{i,j} = \frac{b_{i,j}}{4}, \quad (3)$$

and

$$h_i = -\frac{a_i}{2} - \frac{1}{4} \sum_{j=1, j \neq i}^n b_{i,j}. \quad (4)$$

QA is a heuristic that resembles simulated annealing [29], which is a classical heuristic. It differs to simulated annealing in its use of quantum fluctuations, which are employed through the transverse-field driver Hamiltonian

$$H_D = -\sum_{i=1}^n \sigma_i^x, \quad (5)$$

where σ_i^x is defined in terms of the Pauli operator σ_x as $\sigma_i^x = \mathbb{1}^{\otimes i-1} \otimes \sigma_x \otimes \mathbb{1}^{\otimes n-i}$. QA operates by initializing a system in the state $\sum_{j=0}^{2^n-1} |j\rangle/2^{\frac{n}{2}}$, which is the ground state of H_D , and evolving the system according to the Hamiltonian

$$H(t) = A(t)H_D + B(t)H_P. \quad (6)$$

$A(t)$ and $B(t)$ are control functions, which are time-dependent real numbers that satisfy $A(0) \gg B(0)$ at the beginning and $A(t_f) \ll B(t_f)$ at the end of an anneal of duration t_f . After this quantum evolution, the system is measured in the computational basis. Our choice of control functions for simulations of QA are $A(t) = 1 - t/t_f$ and $B(t) = t/t_f$, which correspond to a linear schedule. The adiabatic theorem implies that the probability of exciting the system and subsequently measuring an excited state of H_P can be made arbitrarily small by making the anneal time t_f sufficiently large [30]. In this work, t_f is not necessarily in the adiabatic regime, so there is some non-negligible probability of sampling a suboptimal solution.

The QAOA is a quantum algorithm that resembles a discretized version of QA. In the QAOA, p layers of gates are applied to an initial state $|\psi_0\rangle = \sum_{j=0}^{2^n-1} |j\rangle/2^{\frac{n}{2}}$, which is the same initial state as in QA. In the k th layer of gates, the system evolves according to the Ising Hamiltonian H_P for a duration specified by an angle γ_k and then by a mixer Hamiltonian H_M for a duration specified by an angle β_k . The mixer Hamiltonian H_M is a sum of the Pauli operators σ_x acting on each qubit,

which is the same as the QA driver Hamiltonian H_D in Eq. (5). After the application of all p layers of gates, the system is prepared in the state

$$|\boldsymbol{\gamma}, \boldsymbol{\beta}\rangle = e^{-i\beta_p H_M} e^{-i\gamma_p H_P} \dots e^{-i\beta_1 H_M} e^{-i\gamma_1 H_P} |\psi_0\rangle, \quad (7)$$

which is dependent on $2p$ angles $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$. The final state $|\boldsymbol{\gamma}, \boldsymbol{\beta}\rangle$ is measured and produces some solution \mathbf{x} with probability $|\langle \mathbf{x} | \boldsymbol{\gamma}, \boldsymbol{\beta} \rangle|^2$. This evolution is a discretized approximation of QA for some choices of the angles $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$ that can be inferred from the control functions $A(t)$ and $B(t)$. Therefore, the adiabatic theorem also applies to the QAOA in the limit of large p , which implies that $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$ can be chosen such that measuring $|\boldsymbol{\gamma}, \boldsymbol{\beta}\rangle$ always produces an optimal solution. In practice, p is typically too small to be able to apply the adiabatic theorem. Instead, $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$ are chosen by a classical optimizer. The optimizer runs the QAOA circuit some number of times and calculates the average objective value $\langle f(\mathbf{x}) \rangle$ of the sampled solutions. It then repeats this process with updated parameters $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$ in each iteration, using classical techniques to minimize $\langle f(\mathbf{x}) \rangle$.

For a quantum algorithm that produces the state $|\psi\rangle$ before measurement, the success probability

$$P_S = \sum_{\mathbf{x}^* \in X^*} |\langle \mathbf{x}^* | \psi \rangle|^2 \quad (8)$$

indicates the probability of measuring an optimal solution, where X^* is the set of optimal solutions. Similarly, the feasible probability

$$P_F = \sum_{\tilde{\mathbf{x}} \in \tilde{X}} |\langle \tilde{\mathbf{x}} | \psi \rangle|^2 \quad (9)$$

is the probability of measuring a solution from the set \tilde{X} of solutions that satisfy all constraints.

B. Promotion cannibalization problem

In this work, we apply QA and the QAOA to two simplified forms of a problem that arises in customer data science when retailers promote products using price reductions. Promoting a product generates additional revenue from sales of that product; however, it can also have the undesired effect of reducing the sales of other similar products. For example, a promotion of one brand of breakfast cereals may generate new sales for that brand at the expense of sales of other brands. This effect is called cannibalization [31,32], and we refer to cannibalization arising from product promotions as *promotion cannibalization*. The overall bilateral cannibalization when concurrently promoting two similar products can result in minimal new sales and possibly even a net reduction in revenue. Therefore, retailers often look to minimize the total revenue loss due to clashing concurrent promotions. We can model this by only considering promotion cannibalization that occurs between pairs of products being promoted at the same time and using a matrix C with matrix elements $C_{i,j}$ that represent the average amount of loss of revenue from sales of product i due to a promotion of product j when both products are promoted concurrently. In this study, we make the assumption that C is non-negative; that is, $C_{i,j} \geq 0 \forall i, j$.

The primary example problem that we base this work on is to choose a selection of A products to promote out of n_p

possible choices that minimizes the total amount of promotion cannibalization between pairs of promoted products. This problem can be expressed as the constrained binary quadratic programming problem [33,34]:

$$\text{find } \arg \min_{\mathbf{x}} f(\mathbf{x}) = \sum_{j=1}^{n_p} \sum_{i=1}^{n_p} C_{i,j} x_i x_j, \quad (10)$$

$$\text{s.t. } \sum_{i=1}^{n_p} x_i = A, \quad (11)$$

where the binary variable $x_i = 1$ if product i is promoted and $x_i = 0$ if not. Equation (11) is an expression of the constraint that A products are promoted in total. Note that when expressed in the form of Eq. (1), the objective function only has quadratic terms with coefficients $b_{i,j}$ equal to $C_{i,j} + C_{j,i}$. Hence, we can assume that C is symmetric in the context of this problem without loss of generality. Conversions from QUBO to Ising form are performed using Eqs. (3) and (4). For example, a C matrix with a single pair of nonzero elements $C_{i,j} = C_{j,i} = 1$ would produce a QUBO objective function containing a single quadratic term with coefficient $b_{i,j} = C_{i,j} + C_{j,i} = 2$, which would correspond to $J_{i,j} = 1/2$ and $h_i = h_j = -1/2$ in Ising form.

Retailers change which products are promoted periodically, such as once every fiscal quarter. The above problem can be viewed as finding the promotion plan for a single quarter of the year. Another problem that we consider is to find an optimal promotion plan for two consecutive quarters, subject to the additional set of constraints that each product is promoted in one of the quarters at most. The optimal promotion plan should minimize the total amount of promotion cannibalization between pairs of promotions in the same quarter, and each quarter should have A promotions. This can be expressed as

$$\text{find } \arg \min_{\mathbf{x}} f(\mathbf{x}) = \sum_{q=1}^2 \sum_{j=1}^{n_p} \sum_{i=1}^{n_p} \lambda_q C_{i,j} x_{i,q} x_{j,q}, \quad (12)$$

$$\text{s.t. } \sum_{i=1}^{n_p} x_{i,q} = A \forall q, \quad (13)$$

$$\text{and } x_{i,1} + x_{i,2} \leq 1 \forall i. \quad (14)$$

Here, the binary variable $x_{i,q}$ represents a promotion of product i in quarter q and λ_q is a seasonal scale factor that is derived from the expected total revenue of each quarter. In total, this problem has $n_p + 2$ constraints to satisfy, which allows us to study the behavior of interacting constraints. For clarity, we refer to Eqs. (10) and (11) as the *single-quarter* promotion cannibalization problem and Eqs. (12)–(14) as the *two-quarter* problem.

C. Penalty methods for encoding constraints

The promotion cannibalization problems, as expressed in Eqs. (10) and (11) and Eqs. (12)–(14), are not QUBO problems because they include constraints. Equations (11) and (13) are examples of linear equality constraints, which take the

general form

$$\sum_{i=1}^n \mu_i x_i = c, \quad (15)$$

for some coefficients $\mu \in \mathbb{R}^n$ and constraint value $c \in \mathbb{R}$. The most common method to encode constraints in quantum optimization is to incorporate them into the objective function through the addition of penalty functions $P(\mathbf{x})$ that raise the objective value of solutions that are infeasible (i.e., do not satisfy every constraint) enough that the solution with the lowest objective value is feasible.

Quantum optimizers and classical quadratic programming solvers commonly use the quadratic penalty function

$$P(\mathbf{x}) = \alpha_2 \left(\sum_{i=1}^n \mu_i x_i - c \right)^2 \quad (16)$$

to encode equality constraints of the form given in Eq. (15). Provided that the penalty strength α_2 is large enough, the optimal solution of an objective function will satisfy Eq. (15) after adding $P(\mathbf{x})$ to it. Aside from the ability to scale the function by changing the penalty strength α_2 , this penalty function has two desirable properties that are satisfied for all $\alpha_2 > 0$:

- (1) $P(\mathbf{x}) = 0$ if \mathbf{x} is feasible.
- (2) $P(\mathbf{x}) > 0$ if \mathbf{x} is infeasible.

The quadratic penalty method suffers from severe drawbacks. Expanding out the brackets in Eq. (16) for a quadratic penalty, we find that there are quadratic terms with nonzero coefficients for all pairs of variables in the constraint. This corresponds to all-to-all couplings of the associated qubits in the Ising Hamiltonian H_P . Many quantum devices that are currently available do not support all-to-all couplings between the physical qubits, meaning that most Ising Hamiltonians of interest cannot be directly mapped to the hardware. To resolve this issue, various methods can be used. In the gate-based setting of QAOA, the quantum states of two qubits can be swapped with a SWAP gate, allowing for any two logical qubit states to be routed through the hardware so they can be physically coupled [35–38]. In QA, a common approach is to map each logical qubit to a chain of ferromagnetically coupled physical qubits in a process called minor embedding [12,20] such that the chains of qubits support the necessary couplings. In both cases, a significant physical resource cost is incurred when new couplings are introduced to H_P , either in the form of a larger circuit depth in the case of the QAOA or a larger qubit count in the case of QA.

As well as the cost of requiring all-to-all couplings, another drawback of the quadratic penalty method is that it will often reduce the effective dynamic range of qubit interactions. The quadratic penalty for the constraint in Eq. (11) is

$$P(\mathbf{x}) = \alpha_2 \left(\sum_{i=1}^{n_p} x_{i,q} - A \right)^2. \quad (17)$$

Expanding out the brackets, we get

$$P(\mathbf{x}) = \alpha_2 \left(\sum_{i=1}^{n_p} (1 - 2A)x_i + \sum_{i=1}^{n_p-1} \sum_{j=i+1}^{n_p} 2x_i x_j + A^2 \right). \quad (18)$$

Mapping this to an Ising Hamiltonian with $x_i \mapsto (1 - \sigma_i^z)/2$ gives

$$P = \alpha_2 \left(\sum_{i=1}^{n_p-1} \sum_{j=i+1}^{n_p} \frac{\sigma_i^z \sigma_j^z}{2} + \sum_{i=1}^{n_p} \left(\frac{n_p}{2} - A \right) \sigma_i^z + \frac{n_p(n_p + 1)}{4} - n_p A + A^2 \right). \quad (19)$$

Now, we can read off the couplings and local fields as $J_{i,j} = \alpha_2/2$ and $h_i = \alpha_2(n_p/2 - A)$. As we would expect, the magnitudes of \mathbf{J} and \mathbf{h} increase with the magnitude of α_2 . As well as this, the magnitude of \mathbf{h} is proportional to the absolute difference $|n_p/2 - A|$. In other words, this penalty introduces strong local fields if the desired constraint value (number of promotions A) is far from half of the number of variables in the constraint (number of products n_p).

In QA, large-magnitude couplings or local fields are undesirable because there are physical limitations on the range of \mathbf{J} or \mathbf{h} values that can be implemented. The Ising Hamiltonian \tilde{H}_P implemented by a quantum annealer is normalized by a factor \mathcal{N} using

$$\tilde{H}_P = \frac{1}{\mathcal{N}} H_P, \quad (20)$$

where \mathcal{N} is typically chosen to be the minimum value such that all physical constraints on \mathbf{J} and \mathbf{h} are satisfied. A penalty that increases the maximum magnitude values in \mathbf{J} or \mathbf{h} will often result in a larger \mathcal{N} , which reduces the effective dynamic range of qubit interactions for the unconstrained part of the problem. Large energies associated with penalties would also be detrimental in the QAOA because the application of H_P would begin rotating phases by angles larger than 2π unless the rotation angles γ are reduced or H_P is normalized. This restricts the effective range of rotation angles for the interactions in H_P that are not associated with the penalties.

Various alternatives to the quadratic penalty method have also been explored in the context of quantum optimization. In QA, one approach is to start in a superposition of feasible states and use an alternative driver Hamiltonian H_D that only drives transitions between feasible states [39,40]. This can also be applied to a generalization of the QAOA by replacing H_M with unitaries that implement the constraints [41]. While this method is appealing as it naturally limits the quantum evolution to the subspace of feasible states, it requires multiqubit interactions that are more difficult to physically implement than the transverse-field Hamiltonian. This method can be combined with a parity-based encoding scheme [21,42], which presents an alternative to minor embedding in QA for compatible hardware. Another approach is to define small Ising problems called *gadgets* that have certain properties that allow them to be combined to encode the original constrained problem [43–45]. These gadgets use ancillary qubits, which is the same type of resource cost as for the quadratic penalty method when considering the extra qubits introduced in minor embedding. When tailored to a device’s hardware, the gadget-based approach can be more efficient than the quadratic penalty method in terms of dynamic range and the number of physical qubits required.

In classical computing, penalty functions that are linear in \mathbf{x} are also used [46]. To satisfy property 2, non-Ising

operations are required. For example, the constraint in Eq. (15) can be implemented with the linear penalty $\alpha_1 |\sum_{i=1}^n \mu_i x_i - c|$. Computing using non-Ising operations such as $|\cdot|$ is challenging in quantum optimization. de la Grand'rive and Hullo [47] demonstrated that the QAOA can be extended to implement non-Ising linear penalties for inequality constraints by computing $\max(\cdot)$, which is a non-Ising operation. However, this approach comes at the cost of requiring ancillary qubits and a more complex circuit.

In this paper, we consider the use of linear penalty functions with non-Ising operations removed. This method would be able to avoid many of the drawbacks of the quadratic penalty method as it does not introduce any new couplings to the problem and can be more efficient with the hardware's dynamic range [26,48]. The removal of non-Ising operations requires giving up on property 2 of a conventional penalty function, which means this penalty method does not produce the desired ground state in H_p for all cases. The upside is that this penalty method can be implemented on current quantum hardware with a smaller physical overhead than other methods and may potentially improve the performance of algorithms. This type of linear penalty method has been previously suggested for QA in the contexts of portfolio optimization [8] and quantum machine learning [27], but neither of these two studies showed any results relating to this approach beyond making the observation that it can be used in these problems.

Ohzeki [26] recently developed a method for implementing constraints with linear terms by taking the partition function of a QUBO objective function and applying a Hubbard-Stratonovich transformation [49,50]. While the mathematical motivation behind Ohzeki's work is different to what we present here, the resulting algorithm is effectively the same as applying linear Ising penalties. It has been observed that the method described by Ohzeki does not have the theoretical guarantee of being able to exactly implement all hard constraints [51]. This work examines objective functions that cannot be exactly constrained by this method and considers their implications, which the prior work does not do. We discuss how the structure of a problem affects the likelihood of this occurring and identify that objective functions with non-negative quadratic term coefficients are particularly suitable for the linear Ising penalty method. One proposal for making use of infeasible solutions is to perform a postprocessing step to produce feasible solutions from infeasible solutions using the fewest number of bit flips [51]. In this work, we propose another strategy of selectively applying linear Ising penalties to specific constraints when it is not possible to apply them to all of a problem's constraints. Whereas Ohzeki and subsequent researchers performed their analyses on D-Wave quantum annealing devices, this work presents results of numerical simulations in both gate-based and annealing settings. Performing simulations allows us to study the behavior of closed-system dynamics and to study how the dynamics are influenced by the choice of penalty strength parameters.

The linear Ising penalty function for the equality constraint in Eq. (11) is

$$P(\mathbf{x}) = \alpha_1 \left(\sum_{i=1}^{n_p} x_{i,q} - A \right), \quad (21)$$

where the penalty strength α_1 can be positive or negative. In the Ising formulation, this corresponds to local fields equal to $-\alpha_1/2$ up to an unimportant constant offset. In the rest of this paper, when using the term *linear penalty*, we refer to linear Ising penalties of the form in Eq. (21) rather than linear penalties with non-Ising operations, which are used in classical computing.

Since we assume that C is non-negative, the QUBO objective function of the single-quarter promotion cannibalization problem only has quadratic terms with non-negative coefficients. To see how this affects the problem's structure, we group all possible solutions by their Hamming weight w , which represents the number of variables equal to one in a solution. The minimum objective value of the solutions for each group is monotonically increasing with w . In other words, for every solution with some Hamming weight $w \geq 1$, there exists some solution with Hamming weight $w - 1$ that has the same or a lower objective value. To see why this is true, suppose there exists a solution \mathbf{x}_a with objective value $f(\mathbf{x}_a)$ and Hamming weight w , where $f(\mathbf{x}_a)$ is strictly less than the objective values of all solutions with Hamming weight $w - 1$. We can flip one of the variables of \mathbf{x}_a from one to zero to get another solution \mathbf{x}_b with a Hamming weight $w - 1$. Since the QUBO objective function does not have any terms with negative coefficients, flipping a variable from one to zero can never increase the objective value. Hence, $f(\mathbf{x}_b) \leq f(\mathbf{x}_a)$, which contradicts the proposition that $f(\mathbf{x}_a)$ is strictly less than the objective value of all solutions with Hamming weight $w - 1$. Therefore, solutions with a Hamming weight w can never have an objective value that is strictly less than the objective value of all solutions with Hamming weight $w - 1$. The same is true for the two-quarter promotion cannibalization problem when considering the Hamming weight of the variables for one quarter while the other quarter's variables have fixed values.

The monotonic relationship between Hamming weight and minimum objective value makes this problem particularly amenable to the linear Ising penalty method. To demonstrate this, we plot $f(\mathbf{x})$ against the Hamming weight of \mathbf{x} for a promotion cannibalization problem instance with six products in Fig. 1(c) and for some other six-variable QUBO instance with both positive and negative quadratic term coefficients in Fig. 1(a). If we suppose we want to implement the constraint $\sum_{i=1}^6 x_i = 2$ for both of these problems, the penalized objective functions must have minimum values at a solution Hamming weight of two. For the promotion cannibalization problem, the minimum value of $f(\mathbf{x})$ is monotonically increasing with the Hamming weight of \mathbf{x} , so only solutions with a Hamming weight of less than two need to be penalized to implement the desired constraint. A linear penalty [shown in the inset of Fig. 1(c)] is able to do this, but it also has the undesired effect of lowering the objective value of solutions with Hamming weights greater than two. However, this is often compensated by the large existing object values of these higher Hamming weight solutions due to the monotonic structure we described. Therefore, the linear penalty method is more likely to be able to produce the correct optimal solution than for a random QUBO problem without the same structure. Indeed, we find that for this instance of the promotion cannibalization problem, the linear penalty method is successful in

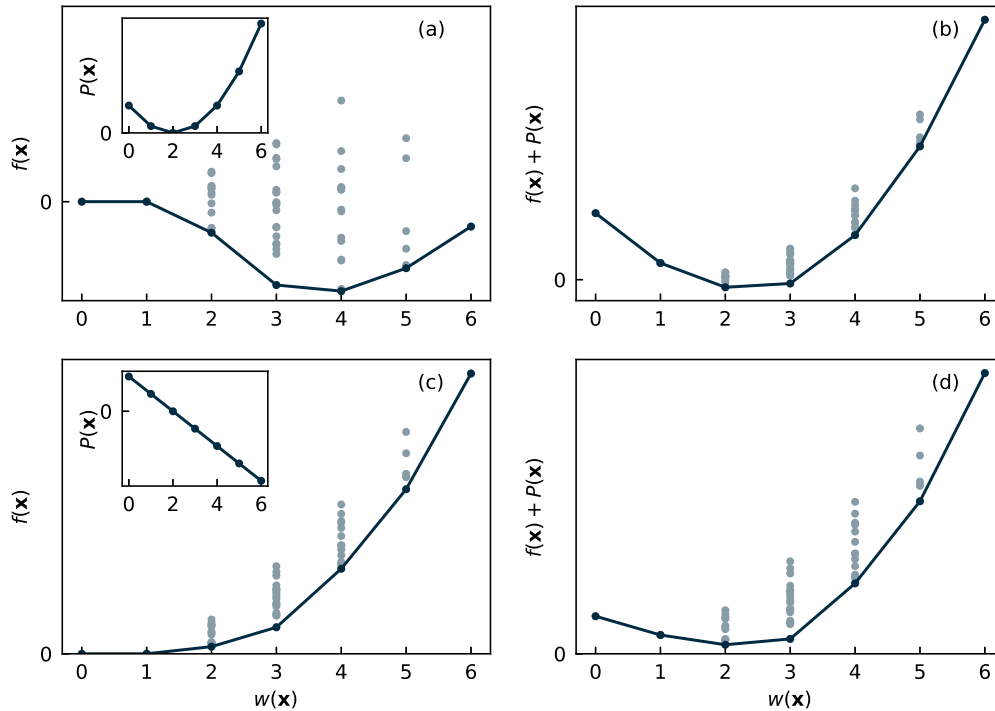


FIG. 1. Objective values $f(\mathbf{x})$ plotted against the Hamming weight $w(\mathbf{x})$ of each possible solution \mathbf{x} for (a) an example QUBO instance with both positive and negative quadratic term coefficients and (c) an instance of the single-quarter promotion cannibalization problem. For each Hamming weight, the minimum value of $f(\mathbf{x})$ is plotted in a darker color, and these points are connected to guide the eye. We apply two different penalty functions $P(\mathbf{x})$ for the equality constraint $\sum_{i=1}^6 x_i = 2$ to these problem instances, which are shown in the insets of panels (a) and (c). The penalty function is quadratic in panel (a) and linear in panel (c). Note that the gradient of the line in the inset of panel (c) is equal to α_1 . (b) The constrained objective function after applying a quadratic penalty to the QUBO instance plotted in panel (a). (d) The constrained objective function after applying a linear penalty to the promotion cannibalization problem instance plotted in panel (c).

producing a constrained objective function, which is shown in Fig. 1(d). For the other QUBO problem, there does not exist any value of α_1 that produces an optimal solution with a Hamming weight of two, so the quadratic penalty method is used instead [Figs. 1(a) and 1(b)]. We note that the monotonic structure of the promotion cannibalization problem is not sufficient to guarantee that the linear penalty method is successful every time. For that, the gradient of the line in Fig. 1(c) also needs to be monotonically increasing, which is not the case for all problem instances.

The constraints given by Eq. (14) are examples of linear inequality constraints. We encode these with quadratic penalty functions of the form

$$P(\mathbf{x}) = \alpha_2 x_{i,1} x_{i,2} \tag{22}$$

throughout our analysis of the two-quarter promotion cannibalization problem.

III. NUMERICAL METHODS

The numerical analysis in this work was performed using the Python programming language [52]. We used the libraries NumPy [53] and SciPy [54] for computationally intensive calculations. All linear fits were obtained using the weighted least-squares method implementation in `scipy.optimize.curve_fit`. PyQUBO [55] was used to formulate QUBO and Ising problems. Plots were produced with Matplotlib [56].

Problem instances were exactly solved using Gurobi Optimizer [57] through the GurobiPy Python interface. We note that Gurobi operates at an adjustable numerical precision, which can lead to minor differences in results depending on which software version and solver parameters are used. Throughout this work, we used Gurobi version 10.0.2 with a single thread and the default values of all other solver parameters.

The C matrices used in our numerical analysis of the single-quarter promotion cannibalization problem were generated by selecting symmetric off-diagonal matrix elements $C_{i,j} = C_{j,i}$ uniformly at random from the interval $[0.1, 1.0)$. All matrix elements on the main diagonal were set to zero. With this method, we generated 10 000 C matrices corresponding to different problem instances for each number of products between 6 and 18. The C matrices of the instance with ID 8_0 and the instance with ID 8_19 are also used for our analysis of the two-quarter promotion cannibalization problem.

All simulations in this work were run on the Hamilton high-performance computing cluster at Durham University. The SciPy function `expm_multiply` was used to simulate QA by discretizing the time evolution in the same manner as described in Ref. [58]. For the QAOA, the `qasm_simulator` backend in the Qiskit SDK [59] was used to simulate quantum circuits without noise. Each of the QAOA angles γ and β were initialized randomly within the interval $[0, 1)$ before being optimized by the implementation of the COBYLA

algorithm [60] in `scipy.optimize.minimize` using a maximum of 100 minimization iterations. While optimizing γ and β , the quantum circuit was run 1000 times in order to estimate the average objective value. After optimizing γ and β , 1 000 000 runs of the quantum circuit were performed, from which the probabilities P_S and P_F were inferred. To reduce the effect of getting stuck in local minima when optimizing γ and β , the entire algorithm was repeated 80 times before selecting the attempt that produced the highest success probability.

IV. BEHAVIOUR OF PENALTY FUNCTIONS

In this section, we study the behavior of linear and quadratic penalty functions as their penalty strength parameters are changed. We first look at the behavior of a single constraint using the single-quarter promotion cannibalization problems as an example. Then, we use the two-quarter promotion cannibalization problem to consider the effects of having multiple constraints.

A. Single constraint

While both the quadratic and linear penalty methods introduce a single penalty strength parameter per constraint, these two types of penalties respond differently to changes in their penalty strengths. Here, we study these differences in the context of the single-quarter promotion cannibalization problem, which has a single constraint [Eq. (11)]. For the quadratic method, if the penalty strength α_2 is too small, it is not successful in producing a feasible ground state in H_P . On the other hand, if α_2 is much larger than necessary, H_P will have a feasible ground state, but the performance of a quantum optimizer in finding the ground state will be hindered. Figure 2 shows an example of this behavior for simulations of closed-system QA. The feasible probability P_F increases with α_2 and stays elevated at large values, whereas the success probability P_S rises, peaks, and then falls. For a given problem instance, there is a value of α_2 that will produce the maximum probability P_S^{\max} of measuring an optimal solution. In practice, there is typically a broad range of values for α_2 that will produce good performance, which can be seen in Fig. 2(a).

For the linear penalty method, the parameter α_1 directly relates to the value of the constraint that is implemented, i.e., the value of A in Eq. (11). As shown in Fig. 3(a), the Hamming weight w of the ground state of H_P increases as α_1 is decreased. All instances of the Ising problem exhibit this monotonic relation between α_1 and w because negative local fields lower the energy of states in proportion to their Hamming weights. For a desired value of A , the penalty strength α_1 should be tuned such that $w = A$. There is a lower and upper limit on the range of values of α_1 that accomplish this, and this range is different for each problem instance. In comparison, there is only a lower limit for α_2 that produces a feasible ground state when using the quadratic penalty method. Therefore, the range of values of α_1 that produce a significant probability of finding an optimal solution is often smaller than the equivalent range of values of α_2 . Furthermore, for some problem instances, there does not exist any value of α_1 that produces the desired ground state in H_P with the linear penalty method, which is not the case for the quadratic method.




FIG. 2. (a) Probability P_S after an anneal time $t_f = 10$ of measuring the optimal state normalized by an estimate of the maximum success probability P_S^{\max} for closed-system QA as a function of the penalty parameter α_2 . Each point is an average over 100 instances of the promotion cannibalization problem with $n_p = 12$ products, with error bars representing the standard error in the mean. The blue shaded region contains the 5th to 95th percentile values. Note that P_S^{\max} is calculated for each instance separately and is estimated by taking the maximum P_S over the plotted values of α_2 . (b) Same as panel (a) with the y axis instead showing the probability P_F of measuring a feasible state.

When choosing α_1 , while there may exist a value that implements the desired constraint in the ground state of H_P for most instances of a given problem, there typically will not be a single value that works for all of these instances, as can be seen in Fig. 3 for the single-quarter promotion cannibalization problem. However, the monotonic relationship between w and α_1 means that we can iteratively search for a good value of the parameter by decreasing (increasing) α_1 when sampled solutions have too few (many) ones. This has previously been suggested in Refs. [8,26]. For the single-quarter promotion cannibalization problem, an effective strategy is to start with a guess of $\alpha_1 = -1$, run the optimizer, and double α_1 until the returned solution has the correct Hamming weight $w = A$ or one that is too large. In the case where w is larger than A , a binary search can be performed to find a value of α_1 that produces the desired value of w . Assuming the optimizer always returns the optimal solution, this method can find α_1 in a number of calls to the optimizer that scales logarithmically with the value of α_1 that correctly implements the constraint and is closest to the initial guess of -1 . The monotonic

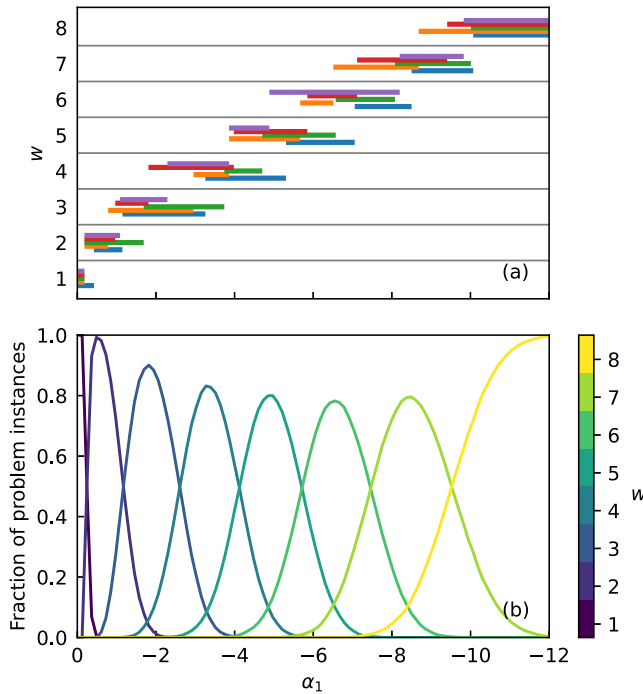


FIG. 3. (a) The Hamming weight w of the ground state of H_P plotted against the linear penalty strength α_1 for five random instances of the promotion cannibalization problem with eight products. Each instance corresponds to a different colored bar. (b) The fraction of 10 000 single-quarter problem instances that have a ground-state Hamming weight of w plotted against α_1 . Note that α_1 is decreasing along the x axis.

relationship between α_1 and w is the condition that makes this true. There is a possibility of the binary search terminating before it finds a correct value of α_1 if the interval of correct values of α_1 is smaller than the precision of the search. The search can be made exponentially more precise by increasing the number of binary search iterations. See Ref. [28] for an example algorithm that iteratively searches for good values of α_1 for problems with one or more linear penalties. The iterative nature of the application of linear penalties makes this method technically a hybrid quantum-classical technique [61].

Not only does α_1 need to be within a finite interval to implement the desired constraint, but the specific choice of α_1 within that interval impacts the performance of the quantum algorithm. Figure 4 demonstrates this for one particular promotion cannibalization problem instance. While all choices of α_1 in the shaded regions of the plot produce the correct ground state in H_P , some choices result in significantly better success probabilities than others. Since in practice the optimal choice of α_1 is not known in advance, our analysis of the performance of the linear penalty method in Sec. V uses values that have been selected uniformly at random from the interval that produces the correct ground state in H_P . Figure 4 also shows that it is possible for QA to have a significant probability of measuring the optimal feasible solution even when the choice of α_1 does not make the ground state feasible (i.e., outside the shaded regions). This probability will tend to zero as the anneal time is increased towards the adiabatic limit.




FIG. 4. Behavior of the linear penalty method for a random single-quarter problem instance with ten products and a constraint of $A = 3$ promotions. In orange, we plot the success probability P_S of finding the optimal feasible solution \mathbf{x}^* after an anneal time $t_f = 200$ for closed-system QA. Note that the penalty strength α_1 is decreasing along the x axis. In blue, we plot the energy difference between the desired state $|\mathbf{x}^*\rangle$ and the minimum-energy eigenstate of H_P that is not $|\mathbf{x}^*\rangle$, where an energy difference that is negative (highlighted by the yellow and green shaded regions) indicates that $|\mathbf{x}^*\rangle$ is the nondegenerate ground state of H_P . The green shaded region indicates where $|\mathbf{x}^*\rangle$ is the ground state and the energy separation to the first excited state is maximized. The α_1 value that maximizes P_S is indicated by the green dashed line.

B. Multiple constraints

In applied optimization, it is common for problems to have many constraints. Therefore, it is important to consider how linear penalty functions interact with each other. We use the two-quarter promotion cannibalization problem to study this. The seasonal scale factors in Eq. (12) are set to $\lambda = (1.5, 1.0)^T$ in the problem instances we consider. For the ground state of H_P , we denote the Hamming weights of the variables associated with the first and second fiscal quarters as w_1 and w_2 , respectively, which correspond to the number of promotions in each quarter. The constraints in Eq. (13) require that $w_1 = w_2 = A$.

We first consider the case where linear penalty functions are used for both constraints in Eq. (13), with the penalty strengths for the constraints on the first and second quarters denoted as $\alpha_1^{(1)}$ and $\alpha_1^{(2)}$, respectively. Figure 5 shows how w_1 and w_2 change with $\alpha_1^{(1)}$ and $\alpha_1^{(2)}$ for two problem instances. In Fig. 5(a), there are regions where $w_1 = w_2 = A$ is satisfied for each possible value of A that also satisfies Eq. (14). However, finding these regions is not as simple as using the same strategy that can be used for the single-quarter problem to search for $\alpha_1^{(1)}$ and $\alpha_1^{(2)}$ independently. This is because changes in $\alpha_1^{(1)}$ impact both w_1 and w_2 , and similarly for $\alpha_1^{(2)}$. Therefore, $\alpha_1^{(1)}$ and $\alpha_1^{(2)}$ cannot be found independently of each other. The monotonic relationship between w_i and $\alpha_1^{(i)}$ can still be used to develop a search strategy that finds $\alpha_1^{(1)}$ and $\alpha_1^{(2)}$ at the same time. An example of a search strategy for problems with multiple linear penalties is given in Ref. [28].

In situations where applying multiple linear penalties is not able to implement the desired constraints, it is often possible




FIG. 5. For the two-quarter promotion cannibalization problem with $n_p = 8$ products, linear penalties are applied to the first and second quarters with penalty strengths $\alpha_1^{(1)}$ and $\alpha_1^{(2)}$, respectively. This produces a ground state of H_P that has a first-quarter Hamming weight w_1 and second-quarter Hamming weight w_2 . Heat maps show the sum $w_1 + w_2$ as a function of the penalty strengths for the C matrices of (a) the instance with ID 8_0 and (b) the instance with ID 8_19. Two different color bars are used depending on whether w_1 and w_2 are equal or not. More saturated colors are used where $w_1 = w_2$ and less saturated colors are used where $w_1 \neq w_2$. Note that $\alpha_1^{(1)}$ and $\alpha_1^{(2)}$ are decreasing along the axes of the plots.

to resolve this issue by switching some of the linear penalty functions to quadratic penalties. In many cases, this maintains some of the advantages of using the linear penalty method. An example where this approach can be used is the instance shown in Fig. 5(b), for which we did not find any combination of $\alpha_1^{(1)}$ and $\alpha_1^{(2)}$ values that satisfies both constraints in Eq. (13) with $A = 3$. Since our search was conducted at a finite precision, this implies that the region in which these constraints are satisfied either does not exist or is very small. In Fig. 6, we consider the same problem instance and instead apply a linear penalty to the first quarter and a quadratic penalty to the second quarter. We find that with this scheme, there exists an interval of values of $\alpha_1^{(1)}$ for which all of the problem's constraints are satisfied. In Ref. [28], there is a more in-depth analysis of combining linear and quadratic penalties using a four-quarter variation of the promotion cannibalization problem. An alternative approach for cases where linear penalties are not successful in implementing all constraints is proposed in Ref. [51], where postprocessing of infeasible solutions is used to obtain feasible solutions.

V. PERFORMANCE IN SIMULATION

We have simulated QA and the QAOA in solving the single-quarter promotion cannibalization problem. Both sets of simulations are of closed-system dynamics and they assume all-to-all connectivity of the physical qubits. The constraint that we consider in these simulations is to select $A = 3$ products to promote. We used the quadratic penalty strength $\alpha_2 = 2$ for all instances. We chose this value as it produced both a good success probability and feasible probability for most instances at the range of problem sizes we have




FIG. 6. Ground-state Hamming weight w_1 of the first quarter of a two-quarter promotion cannibalization problem instance plotted against the linear penalty strength $\alpha_1^{(1)}$ that is applied to the first quarter. A quadratic penalty is applied to the second quarter that constrains its ground-state Hamming weight to $w_2 = 3$. The C matrix is of the instance with ID 8_19, which we did not find any combination of linear penalty strengths for which applying linear penalties to both quarters was successful in creating a ground state in H_P that satisfies $w_1 = w_2 = 3$. Note that the penalty strength $\alpha_1^{(1)}$ is decreasing along the x axis.

considered. This is supported by Fig. 2 and similar figures in the Appendix. The linear penalty strengths α_1 were selected uniformly at random from the interval of values that implement the $A = 3$ constraint in the ground state of H_P for each instance. In some cases, no such interval could be found up to a precision of 10^{-5} , either because the interval does not exist or because it is smaller than the precision of the search. For these instances, a simulation using the linear penalty method was not performed. We have not performed an analysis of the time required to find good values of α_1 in this section. Therefore, it is important to note that this would result in an overhead when using the linear penalty method in practice.

In real-world scenarios, C matrices are sparse. This would result in H_P having fewer nonzero couplings when using linear penalties than when using quadratic penalties, which could contribute to differences in performance, as discussed in Sec. II. For example, a real-world problem might be concerned with 1000 products that each have nonzero cannibalization interactions with five other products on average. In this example, switching from using the quadratic penalty method to the linear penalty method would result in a $\approx 99.5\%$ reduction in the number of nonzero couplings in the objective function. This is because the linear penalty method maintains the sparsity of the C matrix in the objective function, whereas the quadratic penalty method introduces nonzero couplings between all pairs of variables in the single-quarter problem. Due to the small problem sizes that we are limited to in simulation, all C instances used in this work are fully connected. Therefore, our numerical results do not reflect the effect that the sparsity of C matrices would have on the performance of the linear penalty method at large problem sizes.




FIG. 7. Ratio between the average normalization factor of the Ising Hamiltonian when using the quadratic penalty method, \mathcal{N}_Q , and when using the linear penalty method, \mathcal{N}_L , plotted against the number of products n_p for instances of the single-quarter promotion cannibalization problem.




FIG. 8. (a) The success probability P_S of coherent quantum annealing using the quadratic (orange) and linear (blue) penalty methods after an anneal time of $t = 10$ against the number of products n_p in a single-quarter promotion cannibalization problem. The distribution of P_S for the different problem instances is plotted alongside a bar representing the median value of the distribution. Note that the y-axis scale and histogram bin widths are logarithmic. (b) Similar plot for the probability P_F of measuring a feasible state on axes with linear scales.

A. Quantum annealing simulation

In our simulations of QA, the Ising Hamiltonians were normalized by a factor \mathcal{N} , as in Eq. (20), such that the maximum coupling strength is one and the maximum local field strength is three. This reflects the fact that the energy scales that are realizable on physical devices are finite. When using the linear penalty method, the normalization factor \mathcal{N}_L of a given problem instance is different to the normalization factor \mathcal{N}_Q when using the quadratic penalty method. As mentioned in Sec. II, the linear penalty method often produces a smaller normalization factor, which benefits the dynamic range of qubit interactions. In Fig. 7, we see that for all of the problem sizes we have considered, \mathcal{N}_Q is significantly larger than \mathcal{N}_L on average. The larger effective dynamic range when using the linear penalty method can lead to a more efficient exploration of the search space in QA.

For the range of number of products we have considered in the QA simulations, the linear penalty method outperforms the quadratic penalty method on average in finding the optimal solution, as shown in Fig. 8(a). Both methods are more successful than random guessing. As the number of products is increased, the variation in performance across different




FIG. 9. (a) The inferred success probability P_S of the QAOA with $p = 8$ layers using the quadratic (orange) and linear (blue) penalty methods against the number of products n_p in a single-quarter promotion cannibalization problem. The distribution of P_S for the different problem instances is plotted alongside a bar representing the median value of the distribution. Note that the y axis scale and histogram bin widths are logarithmic. (b) Similar plot for the probability P_F of measuring a feasible state on axes with linear scales.

instances shrinks when using the quadratic method, whereas for the linear method it remains large.

In Fig. 8(b), we find that when the number of products is greater than eight, the quadratic penalty method is on average more likely to sample feasible states. The width of the distribution of feasible probabilities is much larger when using the linear method than the quadratic method. This indicates that with our parameter choices, the ability for the linear penalty method to be effective in implementing the constraint is highly dependent on the problem instance, and it explains why the success probability distributions have similarly large widths for the linear method. The fact that the linear penalty method produces a lower feasible probability than the quadratic method on some instances and a higher feasible probability on others suggests that hybrid approaches that attempt using both penalty methods would be effective in practice.

These simulations are of the ideal regime where the effects of limited qubit connectivity and noise do not play a role. Thus, while the advantage of using fewer couplings with the linear penalty method is more clear when the hardware does not support all-to-all connectivity, these simulation results indicate that the linear penalty method may also benefit algorithms on fault-tolerant quantum devices with full qubit connectivity. This is complementary to promising experimental results obtained on D-Wave annealers [26,28,48,51], where the effects of hardware limitations are prominent.

B. Quantum approximate optimization algorithm simulation

We performed another set of simulations that are of the QAOA with $p = 8$ layers of gates. The same problem instances were used as in the QA simulations. Due to time and computational resource limitations, we simulated problems with up to $n_p = 14$ products. Figure 9 shows the inferred success probabilities and feasible probabilities for the QAOA plotted against the number of products in the problem. We find that the QAOA performs better than random guessing on all of the problem instances, whether the linear or quadratic penalty method is used. For all values of n_p we considered, the linear penalty method produces more favorable success and feasible probabilities on average compared with the quadratic penalty method. Since H_p does not need to be normalized in the QAOA, these results are encouraging as they suggest that the linear penalty method continues to have advantages even when there are no strict limitations on the energy scales of interactions.

Comparing the results from the two algorithms, we find that the difference in average success probabilities between the two penalty methods is smaller for the QAOA simulations than the QA simulations. We note that the QAOA and QA results are dependent on the choices of p and t_f , respectively, so a fair comparison of the two algorithms would require these two parameters to be chosen in such a way that the effective anneal times are the same. Since the purpose of this study is not to compare these two algorithms, we have

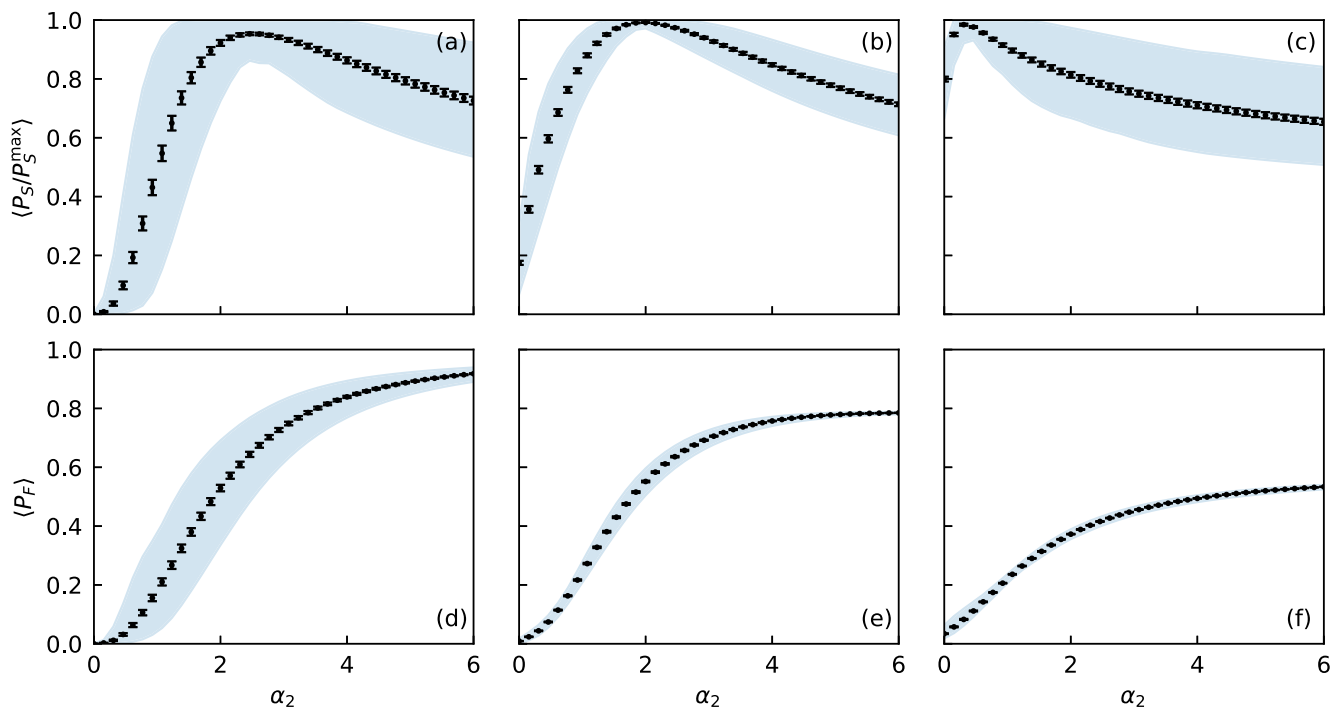


FIG. 10. We have measured the probability P_S after an anneal time $t_f = 10$ of measuring the optimal state normalized by an estimate of the maximum success probability P_S^{\max} for closed-system QA at different values of the penalty parameter α_2 . We plot the average of P_S/P_S^{\max} for 100 instances of the promotion cannibalization problem with (a) 6, (b) 12, and (c) 18 products. Error bars represent the standard error in the mean. The blue shaded regions contain the 5th to 95th percentile values of P_S/P_S^{\max} . Note that P_S^{\max} is calculated for each instance separately and is estimated by taking the maximum of P_S over the plotted values of α_2 . We similarly plot the probability P_F of measuring a feasible state averaged over the same 100 instances with (d) 6, (e) 12, and (f) 18 products.

not done this. An interesting direction for future work would be to determine how the performance difference between the linear and quadratic penalty methods depends on the number of QAOA layers p or anneal time t_f .

Something else that is not explored in our analysis of these simulations is the time required to execute the QAOA algorithm and how it differs between the two penalty methods. In real-world problems, where C matrices are typically sparse, encoding constraints with the linear penalty method would require fewer nonzero couplings in H_p . Therefore, the quantum circuit to implement the unitary $e^{-i\gamma_k H_p}$ requires fewer two-qubit gates, making it more efficient to run the QAOA circuit with the linear method. Another way in which the runtime could differ between the penalty methods would be if it were easier to classically optimize the angles β and γ using one penalty method than the other, if such a difference exists.

VI. CONCLUSIONS

In this work, we have investigated the linear penalty method for encoding constraints and made two theoretical arguments for why the method can lead to better performance in quantum optimization than the quadratic penalty method, which is currently the standard approach. The first argument is based on the fact that the linear penalty method does not introduce any additional quadratic terms to the objective function, and the second is based on considerations of the energy scales associated with the two types of penalty functions. The linear penalty method is not always successful in exactly implementing a desired constraint. However, we have identified a type of customer data science problem for which it is often successful because all quadratic terms in the problem's objective function have non-negative coefficients. There may be other structures that make problems more amenable to the linear penalty method. While this is beyond the scope of our current work, it provides an interesting direction for future research.

We have studied the behavior of the linear and quadratic penalty methods with respect to changes in their penalty strengths α_1 and α_2 . The linear penalty method is more sensitive to its penalty strength, making it necessary to search for different values of α_1 for each problem instance. While this is an extra step that often is not necessary when using the quadratic penalty method, the monotonic relationship between α_1 and Hamming weight can make this search easy to perform in practice. In problems where multiple linear penalties are applied, the different linear penalty strengths can influence each other, making the search for α_1 more difficult. For cases where one or more constraints cannot be successfully implemented with linear penalties, we have shown that it is sometimes possible to use a combination of quadratic and linear penalties to implement the desired constraints.

Our simulations of QA and the QAOA indicate that there may be a performance enhancement when using the linear penalty method over the quadratic method. We suspect that the linear penalty method would more convincingly outperform at larger problem sizes, where the dynamic range effects are more prominent and the problems may be more sparse. This is supported by the findings in Ref. [28], where larger

instances of customer data science problems are tackled on real quantum hardware.

This work only considered the application of linear penalties to equality constraints, but they can also be used for inequality constraints [48]. In this case, the linear penalty would create an approximation of the original problem as it would penalize different feasible solutions by different amounts. However, the larger range of acceptable Hamming weights may make it easier to find values of the linear penalty strengths that produce feasible solutions than for equality constraints. Further research is required to investigate this.

The data and code that support the findings of this study are openly available at Ref. [62].

ACKNOWLEDGMENTS

We thank P. van den Heuvel for providing constructive feedback. N.C. was supported by UK Engineering and Physical Sciences Research Council (EPSRC) Grant No. EP/T026715/2 (CCP-QC). P.M. was supported by UK EPSRC Doctoral Training Funds awarded to Durham University (EP/T518001/1) in partnership with dunnhumby.




FIG. 11. (a) Success probability P_S of QA plotted against the quadratic penalty strength α_2 for the single-quarter cannibalization problem with $n_p = 12$ products. This is plotted for 100 instances of the problem that are represented by different colors, with dashed lines connecting each instance's points to guide the eye. (b) P_S is averaged over the different problem instances and plotted with error bars representing the standard error in the mean. The blue shaded region shows the area that contains the 5th to 95th percentile values of P_S .

APPENDIX: QUADRATIC PENALTY STRENGTH CHOICE FOR NUMERICAL SIMULATIONS

For our simulations of QA and the QAOA, we used the quadratic penalty strength $\alpha_2 = 2$. To choose this value, we measured the success probability P_S and feasible probability P_F of QA over a range of values of α_2 . The resulting plots are shown in Fig. 10 for three different numbers of products n_p . Note that Figs. 10(b) and 10(e) show the same data as Fig. 2. As n_p is increased, we find that P_S peaks at smaller values of α_2 on average, whereas the shape of the curve for P_F does not change as significantly. There is a gradual drop in P_S as α_2 is increased beyond the value that maximizes

P_S . Across the range of values of n_p we have considered in our simulations, setting $\alpha_2 = 2$ results in a good compromise between maximizing P_S and maximizing P_F for most problem instances.

Figures 2 and 10 plot average normalized success probabilities of QA. The normalization was performed in order to avoid the averages being skewed in favor of instances for which the maximum success probability is larger. In Fig. 11, we plot unnormalized success probabilities for problem instances with $n_p = 12$ products. Comparing this with Fig. 2, we see that the average success probability peaks at roughly the same value of α_2 regardless of whether a normalization is performed.

-
- [1] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse Ising model, *Phys. Rev. E* **58**, 5355 (1998).
 - [2] J. Brooke, D. Bitko, T. F. Rosenbaum, and G. Aeppli, Quantum annealing of a disordered magnet, *Science* **284**, 779 (1999).
 - [3] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem, *Science* **292**, 472 (2001).
 - [4] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, [arXiv:1411.4028](https://arxiv.org/abs/1411.4028).
 - [5] L. T. Brady and S. Hadfield, Iterative quantum algorithms for maximum independent set: A tale of low-depth quantum algorithms, [arXiv:2309.13110](https://arxiv.org/abs/2309.13110).
 - [6] A. D. King *et al.*, Computational supremacy in quantum simulation, [arXiv:2403.00910](https://arxiv.org/abs/2403.00910).
 - [7] R. Orús, S. Mugel, and E. Lizaso, Quantum computing for finance: Overview and prospects, *Rev. Phys.* **4**, 100028 (2019).
 - [8] D. Venturelli and A. Kondratyev, Reverse quantum annealing approach to portfolio optimization problems, *Quantum Mach. Intell.* **1**, 17 (2019).
 - [9] K. Kitai, J. Guo, S. Ju, S. Tanaka, K. Tsuda, J. Shiomi, and R. Tamura, Designing metamaterials with quantum annealing and factorization machines, *Phys. Rev. Res.* **2**, 013319 (2020).
 - [10] T. Stollenwerk, B. O’Gorman, D. Venturelli, S. Mandra, O. Rodionova, H. Ng, B. Sridhar, E. G. Rieffel, and R. Biswas, Quantum annealing applied to De-Conflicting optimal trajectories for air traffic management, *IEEE Trans. Intell. Transport. Syst.* **21**, 285 (2020).
 - [11] D. M. Fox, K. M. Branson, and R. C. Walker, mRNA codon optimization with quantum computers, *PLoS One* **16**, e0259101 (2021).
 - [12] S. Yarkoni, E. Raponi, T. Bäck, and S. Schmitt, Quantum annealing for industry applications: Introduction and review, *Rep. Prog. Phys.* **85**, 104001 (2022).
 - [13] N. Nishimura, K. Tanahashi, K. Suganuma, M. J. Miyama, and M. Ohzeki, Item listing optimization for e-commerce websites based on diversity, *Front. Comput. Sci.* **1**, 2 (2019).
 - [14] S. J. Weinberg, F. Sanches, T. Ide, K. Kamiya, and R. Correll, Supply chain logistics with quantum and classical annealing algorithms, *Sci. Rep.* **13**, 4770 (2023).
 - [15] K. Czerniachowska, A genetic algorithm for the retail shelf space allocation problem with virtual segments, *OPSEARCH* **59**, 364 (2022).
 - [16] S. Subramanian and H. D. Sherali, A fractional programming approach for retail category price optimization, *J. Global Optim.* **48**, 263 (2010).
 - [17] A. Lucas, Ising formulations of many NP problems, *Front. Phys.* **2**, 5 (2014).
 - [18] B. Lodewijks, Mapping NP-hard and NP-complete optimisation problems to quadratic unconstrained binary optimisation problems, [arXiv:1911.08043](https://arxiv.org/abs/1911.08043).
 - [19] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
 - [20] V. Choi, Minor-embedding in adiabatic quantum computation: I. The parameter setting problem, *Quantum Inf. Process.* **7**, 193 (2008).
 - [21] W. Lechner, P. Hauke, and P. Zoller, A quantum annealing architecture with all-to-all connectivity from local interactions, *Sci. Adv.* **1**, e1500838 (2015).
 - [22] A. Rocchetto, S. C. Benjamin, and Y. Li, Stabilizers as a design tool for new forms of the Lechner-Hauke-Zoller annealer, *Sci. Adv.* **2**, e1601246 (2016).
 - [23] N. Chancellor, Domain wall encoding of discrete variables for quantum annealing and QAOA, *Quantum Sci. Technol.* **4**, 045004 (2019).
 - [24] J. Chen, T. Stollenwerk, and N. Chancellor, Performance of domain-wall encoding for quantum annealing, *IEEE Trans. Quantum Eng.* **2**, 1 (2021).
 - [25] J. Berwald, N. Chancellor, and R. Dridi, Understanding domain-wall encoding theoretically and experimentally, *Philos. Trans. R. Soc. A* **381**, 20210410 (2023).
 - [26] M. Ohzeki, Breaking limitation of quantum annealer in solving optimization problems under constraints, *Sci. Rep.* **10**, 3126 (2020).
 - [27] D. Willsch, M. Willsch, H. De Raedt, and K. Michielsen, Support vector machines on the D-Wave quantum annealer, *Comput. Phys. Commun.* **248**, 107006 (2020).
 - [28] P. Mirkarimi, D. C. Hoyle, R. Williams, and N. Chancellor, Experimental demonstration of improved quantum optimization with linear Ising penalties, *New J. Phys.* **26**, 103005 (2024).
 - [29] S. Kirkpatrick, C. D. Gelatt, Jr., and M. P. Vecchi, Optimization by simulated annealing, *Science* **220**, 671 (1983).
 - [30] T. Albash and D. A. Lidar, Adiabatic quantum computation, *Rev. Mod. Phys.* **90**, 015002 (2018).
 - [31] L. Meredith and D. Maki, Product cannibalization and the role of prices, *Appl. Econ.* **33**, 1785 (2001).

- [32] C. Aguilar-Palacios, S. Munoz-Romero, and J. L. Rojo-Alvarez, Causal quantification of cannibalization during promotional sales in grocery retail, *IEEE Access* **9**, 34078 (2021).
- [33] J. Nocedal and S. J. Wright, *Numerical Optimization*, Springer Series in Operations Research and Financial Engineering (Springer, New York, 2006).
- [34] N. van Thoai, [Solution methods for general quadratic programming problem with continuous and binary variables: Overview](#), in *Advanced Computational Methods for Knowledge Engineering*, edited by N. T. Nguyen, T. van Do, and H. A. le Thi (Springer International Publishing, Heidelberg, 2013), pp. 3–17.
- [35] I. D. Kivlichan, J. McClean, N. Wiebe, C. Gidney, A. Aspuru-Guzik, G. K.-L. Chan, and R. Babbush, Quantum simulation of electronic structure with linear depth and connectivity, *Phys. Rev. Lett.* **120**, 110501 (2018).
- [36] B. O’Gorman, W. J. Huggins, E. G. Rieffel, and K. B. Whaley, Generalized swap networks for near-term quantum computing, [arXiv:1905.05118](#).
- [37] T. Hage, Optimal fermionic swap networks for Hubbard models, [arXiv:2001.08324](#).
- [38] A. Hashim, R. Rines, V. Omole, R. K. Naik, J. M. Kreikebaum, D. I. Santiago, F. T. Chong, I. Siddiqi, and P. Gokhale, Optimized SWAP networks with equivalent circuit averaging for QAOA, *Phys. Rev. Res.* **4**, 033028 (2022).
- [39] I. Hen and F. M. Spedalieri, Quantum annealing for constrained optimization, *Phys. Rev. Appl.* **5**, 034007 (2016).
- [40] I. Hen and M. S. Sarandy, Driver Hamiltonians for constrained optimization in quantum annealing, *Phys. Rev. A* **93**, 062312 (2016).
- [41] 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).
- [42] M. Drieb-Schön, K. Ender, Y. Javanmard, and W. Lechner, Parity quantum optimization: Encoding constraints, *Quantum* **7**, 951 (2023).
- [43] T. Vyskocil and H. Djidjev, Embedding equality constraints of optimization problems into a quantum annealer, *Algorithms* **12**, 77 (2019).
- [44] T. Vyskočil, S. Pakin, and H. N. Djidjev, [Embedding inequality constraints for quantum annealing optimization](#), in *Quantum Technology and Optimization Problems*, Lecture Notes in Computer Science, edited by S. Feld and C. Linnhoff-Popien (Springer International Publishing, Cham, 2019), Vol. 11413, pp. 11–22.
- [45] H. Djidjev, [Automaton-based methodology for implementing optimization constraints for quantum annealing](#), in *Proceedings of the 17th ACM International Conference on Computing Frontiers* (ACM, New York, 2020), pp. 118–125.
- [46] R. Fletcher, [Penalty Functions](#), in *Mathematical Programming the State of the Art*, edited by A. Bachem, B. Korte, and M. Grötschel (Springer, Berlin, Heidelberg, 1983), pp. 87–114.
- [47] P. D. de la Grand’rive and J.-F. Hullo, Knapsack Problem variants of QAOA for battery revenue optimisation, [arXiv:1908.02210](#).
- [48] S. Yu and T. Nabil, Applying the Hubbard-Stratonovich transformation to solve scheduling problems under inequality constraints with quantum annealing, *Front. Phys.* **9**, 730685 (2021).
- [49] R. L. Stratonovich, On a method of calculating quantum distribution functions, in *Soviet Physics Doklady* (American Institute of Physics, New York, 1957), Vol. 2, p. 416.
- [50] J. Hubbard, Calculation of partition functions, *Phys. Rev. Lett.* **3**, 77 (1959).
- [51] M. Kuramata, R. Katsuki, and K. Nakata, [Larger sparse quadratic assignment problem optimization using quantum annealing and a bit-flip heuristic algorithm](#), in *2021 IEEE 8th International Conference on Industrial Engineering and Applications (ICIEA)* (IEEE, Chengdu, China, 2021), pp. 556–565.
- [52] G. van Rossum and F. L. Drake, Jr., *Python Tutorial* (Centrum voor Wiskunde en Informatica, The Netherlands, 1995).
- [53] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith *et al.*, Array programming with NumPy, *Nature (London)* **585**, 357 (2020).
- [54] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright *et al.*, SciPy 1.0: Fundamental algorithms for scientific computing in Python, *Nat. Methods* **17**, 261 (2020).
- [55] M. Zaman, K. Tanahashi, and S. Tanaka, PyQUBO: Python library for mapping combinatorial optimization problems to QUBO form, *IEEE Trans. Comput.* **71**, 838 (2022).
- [56] J. D. Hunter, Matplotlib: A 2D graphics environment, *Comput. Sci. Eng.* **9**, 90 (2007).
- [57] Gurobi Optimization, LLC, Gurobi Optimizer Reference Manual (2023), <https://www.gurobi.com/>.
- [58] A. Callison, M. Festenstein, J. Chen, L. Nita, V. Kendon, and N. Chancellor, Energetic perspective on rapid quenches in quantum annealing, *PRX Quantum* **2**, 010338 (2021).
- [59] Qiskit contributors, Qiskit: An open-source framework for quantum computing, Version 0.44.0, Zenodo (2023), <https://doi.org/10.5281/zenodo.8190968>.
- [60] M. J. D. Powell, [A direct search optimization method that models the objective and constraint functions by linear interpolation](#), in *Advances in Optimization and Numerical Analysis*, edited by S. Gomez and J.-P. Hennart (Springer Netherlands, Dordrecht, 1994), pp. 51–67.
- [61] A. Callison and N. Chancellor, Hybrid quantum-classical algorithms in the noisy intermediate-scale quantum era and beyond, *Phys. Rev. A* **106**, 010101 (2022).
- [62] P. Mirkarimi, I. Shukla, D. C. Hoyle, R. Williams, and N. Chancellor, [Quantum Optimization with Linear Ising Penalty Functions for Customer Data Science \[dataset\]](#) (Durham University, UK, 2024).