

PAPER • OPEN ACCESS

Unbalanced penalization: a new approach to encode inequality constraints of combinatorial problems for quantum optimization algorithms

To cite this article: J A Montañez-Barrera *et al* 2024 *Quantum Sci. Technol.* **9** 025022

View the [article online](#) for updates and enhancements.

You may also like

- [Quantum annealing for industry applications: introduction and review](#)
Sheir Yarkoni, Elena Raponi, Thomas Bäck et al.
- [An integrated programming and development environment for adiabatic quantum optimization](#)
T S Humble, A J McCaskey, R S Bennink et al.
- [Entropic property of randomized QAOA circuits](#)
A Yu Chernyavskiy, B I Bantysh and Yu I Bogdanov



Easy-to-use and Helium-3 free
cryogenics solutions



LEARN MORE

Quantum Science and Technology



PAPER

OPEN ACCESS

RECEIVED
20 June 2023

REVISED
25 January 2024

ACCEPTED FOR PUBLICATION
20 March 2024

PUBLISHED
2 April 2024

Original Content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



Unbalanced penalization: a new approach to encode inequality constraints of combinatorial problems for quantum optimization algorithms

J A Montañez-Barrera^{1,*} , Dennis Willsch¹ , A Maldonado-Romo^{2,3} and Kristel Michielsen^{1,4,5}

¹ Jülich Supercomputing Centre, Institute for Advanced Simulation, Forschungszentrum Jülich, 52425 Jülich, Germany

² Centro de investigación en Computación, Instituto Politécnico Nacional, Mexico City, Mexico

³ Entropica Labs, 186b Telok Ayer Street, Singapore 068632, Singapore

⁴ AIDAS, 52425 Jülich, Germany

⁵ RWTH Aachen University, 52056 Aachen, Germany

* Author to whom any correspondence should be addressed.

E-mail: j.montanez-barrera@fz-juelich.de

Keywords: combinatorial optimization, QUBO, inequality constraints, quantum optimization, QAOA, quantum annealing, D-Wave

Abstract

Solving combinatorial optimization problems of the kind that can be codified by quadratic unconstrained binary optimization (QUBO) is a promising application of quantum computation. Some problems of this class suitable for practical applications such as the traveling salesman problem (TSP), the bin packing problem (BPP), or the knapsack problem (KP) have inequality constraints that require a particular cost function encoding. The common approach is the use of slack variables to represent the inequality constraints in the cost function. However, the use of slack variables considerably increases the number of qubits and operations required to solve these problems using quantum devices. In this work, we present an alternative method that does not require extra slack variables and consists of using an unbalanced penalization function to represent the inequality constraints in the QUBO. This function is characterized by larger penalization when the inequality constraint is not achieved than when it is. We evaluate our approach on the TSP, BPP, and KP, successfully encoding the optimal solution of the original optimization problem near the ground state cost Hamiltonian. Additionally, we employ *D*-Wave Advantage and *D*-Wave hybrid solvers to solve the BPP, surpassing the performance of the slack variables approach by achieving solutions for up to 29 items, whereas the slack variables approach only handles up to 11 items. This new approach can be used to solve combinatorial problems with inequality constraints with a reduced number of resources compared to the slack variables approach using quantum annealing or variational quantum algorithms.

1. Introduction

The exploration of quantum enhancement in combinatorial optimization problems has been extended over several years. There are three main reasons for this; first, many combinatorial optimization problems can be encoded in Hamiltonians, the ground state being the optimal solution of the problem [1, 2]; second, they are commonly hard to solve and have practical applications [3]; and third, the quantum algorithms to solve these problems need few resources and can be tested on current state-of-the-art quantum hardware [4, 5].

The usual approach for encoding combinatorial optimization problems on quantum processing units (QPUs) is to transform them in their quadratic unconstrained binary optimization (QUBO) representation and obtain the cost Hamiltonian after a change of variables. In the QUBO encoding, the constraints of the combinatorial optimization problems are added to the cost function as penalization terms along with the objective function.

However, we are currently in a noisy intermediate-scale quantum stage and QPUs are restricted in terms of qubits number, noise level, and circuit depth [6]. The main consequence is that numerous problems solvable by quantum algorithms such as Shor's algorithm [7, 8] remain unpractical to test on real quantum hardware. Despite this fact, solving combinatorial optimization problems is possible on current hardware. For example, on gate-based QPUs, variational quantum algorithms (VQA) [9] are applicable even with the above limitations thanks to their short depth and resistance to noise [10]. Besides, quantum annealers, using the quantum annealing (QA) principle [11–13], are another quantum technology used to solve combinatorial optimization problems [14–16]. In this respect, QA shows some advantages under certain conditions compared to classical annealing [17–19].

Of the aforementioned VQA algorithms, the most studied for combinatorial optimization problems is the quantum approximate Optimization algorithm (QAOA) [20]. Even if it is still unclear whether QAOA will give any advantage compared with other algorithms [21], its simplicity makes it of great interest for analytical and practical purposes [4]. In the simplest QAOA version, the cost Hamiltonian of a combinatorial optimization problem is encoded in a parametric unitary gate along with a 'mixer', a second parametric unitary gate that does not commute with the first unitary gate. In this context, the parameters are adjusted to minimize the expectation value of the cost Hamiltonian using a classical optimizer. Multiple approaches for solving combinatorial optimization problems using QAOA or QA can be found in the literature, for example, in logistics [22], finance [5, 23–25], energy [26], communications [27], automotive industry [28], traffic signaling [29], among others.

From the different sets of problems that can be solved using quantum hardware, the ones with inequality constraints require an extra number of variables to get their QUBO representation. For example, the bin packing problem (BPP) and the traveling salesman problem (TSP) have many inequality constraints that increase with the number of items and cities, respectively. The usual approach implemented in software development kits such as Qiskit [30] or D-Wave Ocean [31] is to use slack variables to encode the inequality constraints [32]. Still, such an approach increments the number of qubits and connections needed to solve these problems. We remark that an alternative approach to handle inequality constraints with penalization would be to develop direct algorithms that respect the constraints by construction, as illustrated in [33] for the asset exchange problem.

In this paper, we propose an alternative method to encode inequality constraints in the QUBO formulation of combinatorial optimization problems. In this new heuristic encoding method, inequality constraints are encoded using an unbalanced penalization formula. This formulation adds larger penalization when the inequality constraint is not fulfilled than when it is. We test our method on the TSP, the BPP, and the knapsack problem (KP) using OpenQAOA [34], a python-based library developed by Entropica Labs, and the Jülich universal quantum computer simulator (JUQCS) [35–37] for problems with up to 43 qubits. In the majority of the cases, the optimal solution of the original combinatorial optimization problem is encoded in the ground state of the cost Hamiltonian of their respective QUBO. However, in some cases, the optimal solution is in the vicinity of the ground state. For such cases, our method still provides a good approximation given the whole set of eigenvalues.

Additionally, we test the performance for finding solutions for the BPP using the D-Wave Hybrid solver and the quantum annealer D-Wave Advantage 5.3 System JUPSI, a quantum computing system developed by D-Wave Systems Inc that has more than 5000 qubits located in Jülich, Germany. Our findings demonstrate that our method outperforms existing approach in terms of both the quantity and quality of solutions obtained for the BPP. Specifically, using the unbalanced penalization approach, we achieve remarkable results on both systems. With the D-Wave Advantage system, we obtain solutions for up to 7 nodes, while with the D-Wave Hybrid solver, we achieve solutions for up to 29 nodes. In contrast, the traditional slack variables approach only handles solutions for up to 7 and 11 nodes for the D-Wave Advantage and D-Wave Hybrid, respectively.

The rest of the paper is organized as follows. Section 2 provides a description of the implementation of the combinatorial optimization problems using QUBOs and an overview of the TSP, BPP, and KP. Section 3 presents a description of the *unbalanced penalization* approach and a new metric called coefficient of performance (CoP) to study the efficiency of QAOA to solve combinatorial optimization problems. Section 4 presents the results and discussion. Finally, section 5 provides some conclusions. The source code for the new method setup and implementation on the TSP, BPP, and KP can be found at <https://jugit.fz-juelich.de/qip/unbalanced-penalizations-qubo>.

2. QUBO

The set of combinatorial problems that can be represented by the QUBO formulation are characterized by functions of the form

$$f(\mathbf{x}) = \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n q_{ij} x_i x_j, \quad (1)$$

where n is the number of variables, $q_{ij} \in \mathbb{R}$ are coefficients associated to the specific problem, and $x_i \in \{0, 1\}$ are the binary variables of the problem. Note that $x_i x_i \equiv x_i$ and $q_{ij} = q_{ji}$ in this formulation. Therefore, the general form of a combinatorial optimization problem solvable by QPUs is given by the cost function

$$f(\mathbf{x}) = \sum_{i=1}^{n-1} \sum_{j>i}^n q_{ij} x_i x_j + \sum_{i=1}^n q_{ii} x_i, \quad (2)$$

and equality constraints given by

$$\sum_{i=1}^n c_i x_i = C, \quad c_i \in \mathbb{Z}, \quad (3)$$

and inequality constraints given by

$$\sum_{i=1}^n l_i x_i \leq B, \quad l_i \in \mathbb{Z}. \quad (4)$$

To transform these problems into the QUBO formulation the constraints are added as penalization terms. In this respect, the equality constraints are included in the cost function using the following penalization term

$$\lambda_0 \left(\sum_{i=1}^n c_i x_i - C \right)^2, \quad (5)$$

where λ_0 is a penalization coefficient that should be chosen to guarantee that the equality constraint is fulfilled and C is a constant value. In the case of inequality constraint, the common approach is to use a slack variable [32, 38]. The slack variable, S , is an auxiliary variable that makes a penalization term vanish when the inequality constraint is achieved.

$$B - \sum_{i=1}^n l_i x_i - S = 0. \quad (6)$$

Therefore, when equation (4) is satisfied, equation (6) is already zero. This means the slack variable, S , must be in the range $0 \leq S \leq B - \min_x \sum_{i=1}^n l_i x_i$. To represent the *slack* variable in binary form, the slack is decomposed in $N = \lceil \log_2(\max_x B - \sum_{i=1}^n l_i x_i) + 1 \rceil$ binary variables:

$$S = \sum_{k=0}^{N-1} 2^k s_k, \quad (7)$$

where s_k are the slack binary variables. Then, the inequality constraints are added as penalization terms by

$$\lambda_1 \left(B - \sum_{i=1}^n l_i x_i - \sum_{k=0}^{N-1} 2^k s_k \right)^2. \quad (8)$$

Combining equation (2) and the two kinds of constraints equations (5) and (8), the general QUBO representation of a given combinatorial optimization problem is given by

$$\min_{\mathbf{x}} \left(\sum_{i=1}^{n-1} \sum_{j>i}^n c_{ij} x_i x_j + \sum_{i=1}^n h_i x_i + \lambda_0 \left(\sum_{i=1}^n q_i x_i - C \right)^2 + \lambda_1 \left(B - \sum_{i=1}^n l_i x_i - \sum_{k=0}^{N-1} 2^k s_k \right)^2 \right). \quad (9)$$

Following the same principle, more constraints can be added and note that after some manipulations, equation (9) can be rewritten in the form of equation (2). The last step to represent the QUBO problem on QPUs is to change the x_i variables to spin variables $z_i \in \{1, -1\}$ by the transformation $x_i = (1 - z_i)/2$. Hence, equation (2) represented in terms of the cost Hamiltonian model reads

$$H_c(z) = \sum_{i=1}^{n-1} \sum_{j>i}^n q_{ij} (1 - z_i) (1 - z_j) / 4 + \sum_{i=1}^n q_{ii} (1 - z_i) / 2. \quad (10)$$

2.1. TSP

In the TSP, a traveler has to stop by a set of cities, finding the shortest route to visit every city once and return to the starting point. This problem is normally solved classically for thousands to millions of variables using heuristic techniques [39, 40]. For instance, the Concorde TSP solver [41] can solve all 110 problems from TSPLIB [42] to optimality, the largest having 85 000 cities.

Different formulations of the TSP exist, but in this work, we focus on the Dantzig–Fulkerson–Johnson (DFJ) formulation [43], which is the linear programming version of the TSP. Note that while there are other formulations of this problem that do not require inequality constraints, e.g. [1], the DFJ formulation is used here to test the capabilities of the unbalanced penalization method to encode multiple inequality constraints. In this formulation, the problem is given by

$$\min \sum_{i=1}^n \sum_{j \neq i, j=1}^n c_{ij} x_{ij}, \quad (11)$$

subject to the set of constraints,

$$\sum_{i=1, i \neq j}^n x_{ij} = 1 \quad \forall j = 1, \dots, n, \quad (12)$$

$$\sum_{j=1, j \neq i}^n x_{ij} = 1 \quad \forall i = 1, \dots, n, \quad (13)$$

$$\sum_{i \in Q} \sum_{j \neq i, j \in Q} x_{ij} \leq |Q| - 1 \quad \forall Q \subsetneq \{1, \dots, n\}, |Q| \geq 2, \quad (14)$$

$$x_{ij} \in \{0, 1\} \quad \forall i = 1, \dots, n \quad \forall j = 1, \dots, n, \quad (15)$$

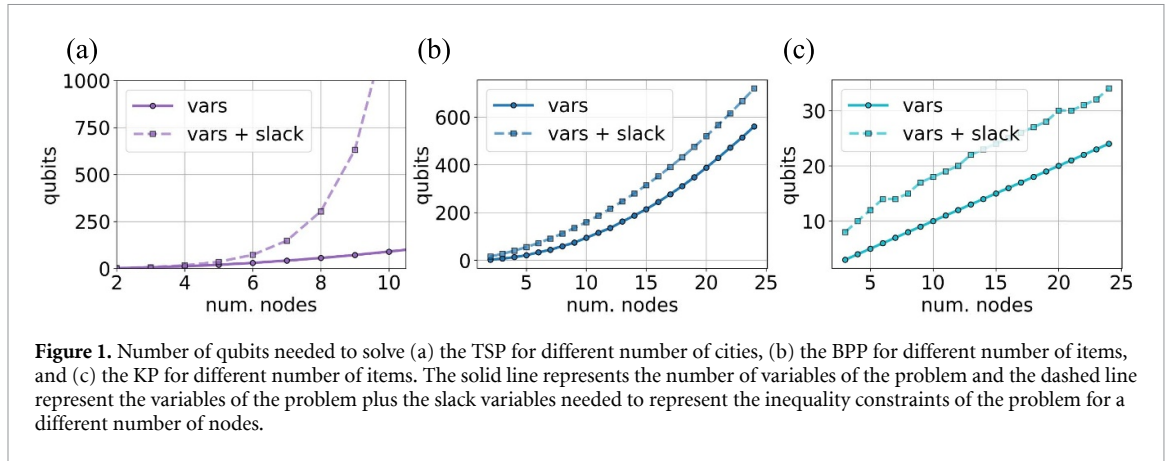
where n is the number of cities (nodes), c_{ij} is the distance between the cities i and j , x_{ij} is a binary variable that represents if the path from the city i to the city j is used or not, and Q represents a sub-tour on a set of cities. From the above equations, equation (11) is the cost function for the distance traveled, equations (12) and (13) restrict the traveler to take only one path to enter and one to leave a city, and equation (14) are the inequality constraints that prohibit sub-tours in the solution. Of all the inequality constraints analyzed in this work, constraint equation (14) is the most expensive in terms of the number of slack variables. Figure 1(a) shows the number of qubits needed to represent the TSP for a different number of cities.

For this work, we randomly place n cities on a 50×50 grid with c_{ij} being the Euclidian distance between cities. We vary from 2 to 7 cities and for problems up to 5 cities, we generate 10 random problems to test the generalization of our method.

2.2. KP

In the KP, a set of items with associated weights and values should be stored in a knapsack. The problem is to maximize the value of the items (nodes) transported in the knapsack. The KP is restricted by the maximum weight the knapsack can carry. Typically, this problem is solved classically for thousands of variables and accepts pseudo-polynomial algorithms, i.e. an algorithm with complexity bounded by a polynomial in the number of items and the maximum weight [44].

KP is the simplest nontrivial integer programming model with binary variables, only one constraint, and positive coefficients. It is formally defined by



$$\max \sum_{i=1}^n p_i x_i, \quad (16)$$

$$\sum_{i=1}^n w_i x_i \leq W, \quad (17)$$

where n is the number of items, p_i and w_i is the value and weight of the i th item, respectively, x_i is the binary variable that represents whether the i th item is in the knapsack or not, and W is the maximum weight that the knapsack can transport. Figure 1(c) shows the number of qubits needed to solve the KP for a different number of items. The set of problems is created with weights and values selected randomly ranging between 1 and 63 for the values, between 1 and 127 for the weights, and knapsack maximum weight is equal to 70% of the sum of the weights of all the items.

2.3. BPP

In the BPP, a collection of items must be stored in the minimum possible number of bins. In this case, the items have a weight associated and the bins are restricted to carry a maximum weight. Commonly, this problem accepts hundreds of bins, but it has been shown to be NP-Hard in the strongest sense, i.e. there is no pseudo-polynomial algorithm to solve this problem [44]. The problem has many real-world applications such as loading trucks with a weight restriction [45], container scheduling [46], or design of FPGA chips [47]. Its formulation is given by

$$\min \sum_{j=1}^m y_j, \quad (18)$$

subject to:

$$\sum_{i=1}^n w_i x_{ij} \leq B y_j \quad \forall j = 1, \dots, m, \quad (19)$$

$$\sum_{j=1}^m x_{ij} = 1 \quad \forall i = 1, \dots, n, \quad (20)$$

$$x_{ij} \in \{0, 1\} \quad \forall i = 1, \dots, n \quad \forall j = 1, \dots, m, \quad (21)$$

$$y_j \in \{0, 1\} \quad \forall j = 1, \dots, m, \quad (22)$$

where n is the number of items (nodes), m is the number of bins, w_i is the i th item weight, B is the maximum weight of each bin, x_{ij} and y_j are binary variables that represent if the item i is in the bin j , and whether bin j is used or not, respectively. From the above equations, equation (18) is the cost function to minimize the number of bins, equation (19) is the inequality constraint for the maximum weight of a bin, equation (20) is the equality constraint to restrict that an item is only in one of the bins, and equations (21) and (22) means that y_j and x_{ij} are binary variables.

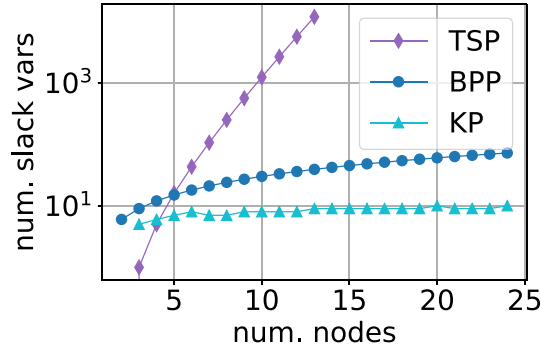


Figure 2. The number of slack variables needed to solve the TSP (diamonds), the BPP (circles), and the KP (triangles) based on the number of nodes. The solid lines are guides to the eye.

Figure 1(b) shows the number of variables needed to solve the BPP when the number of bins is equal to the number of items. The weights are chosen randomly with values between 4 and 20 and the maximum bin weight equal to 20. Without losing generality, two further simplifications are made, the first assigns the first item to the first bin $x_{11} = 1$ and therefore $x_{1j} = 0 \forall j \in \{2, \dots, m\}$. Second replaces $y_j = 1 \forall j \in \{1, 2, \dots, N_{\min}^{\text{bin}}\}$ with $N_{\min}^{\text{bin}} = \lceil (\sum_{i=1}^n s(i)) \rceil / B$ because the minimum of bins required is known.

An estimate of the number of slack variables needed to represent the QUBO of the TSP, the BPP, and the KP is shown in figure 2. In the TSP, the number of slack variables increases exponentially with the number of cities added. In the BPP, it requires many slack binary variables increasing proportionally to the number of bins of the problem. Lastly, in the KP, the number of slack variables needed is constant and depends on the maximum weight allowed in the knapsack.

3. Unbalanced penalization

The unbalanced penalization approach we propose is an approximation for including inequality constraints of combinatorial optimization problems in the QUBO representation of the cost function. This method does not require additional slack variables, and the inequality constraint of equation (4) given by

$$h(x) = B - \sum_i l_i x_i \geq 0, \quad (23)$$

is replaced by a penalization term in the QUBO formulation that makes larger penalization for negative terms and smaller for positive. The idea of the unbalanced penalization function comes from the shape of the exponential decay curve, $f(x) = e^{-h(x)}$. In this function, positive values of $h(x)$, given by equation (23), make $f(x) \approx 0$ while negative values make it grow exponentially, as shown in figure 3. However, the exponential function cannot be encoded as a QUBO penalization term. Therefore, we consider an expansion of the exponential function to quadratic order. The penalization term is given by

$$e^{-h(x)} \approx 1 - h(x) + \frac{1}{2} h(x)^2. \quad (24)$$

In general, we modify equation (24) to include free parameters to be adjusted for the different kinds of problems. Therefore, equation (24) is rewritten by

$$\zeta(x) = -\lambda_1 h(x) + \lambda_2 h(x)^2 \quad (25)$$

$$= -\lambda_1 \left(B - \sum_i l_i x_i \right) + \lambda_2 \left(B - \sum_i l_i x_i \right)^2, \quad (26)$$

where $\lambda_{1,2}$ are a set of multipliers that can be tuned for the specific problem and the constant term is removed because the position of the cost function's minimum is independent of the constant term. Note that this approach can be extended to $g(x) = B - \sum_i l_i x_i \leq 0$, in that case $f(x) = e^{g(x)}$. The new cost function based on equation (9) and the unbalanced penalization approach is given by

$$\min_x \left(\sum_{i=1}^n \sum_{j \neq i}^n c_{ij} x_i x_j + \sum_{i=1}^n h_i x_i + \lambda_0 \left(\sum_{i=1}^n q_i x_i - C \right)^2 - \lambda_1 \left(B - \sum_i l_i x_i \right) + \lambda_2 \left(B - \sum_i l_i x_i \right)^2 \right), \quad (27)$$

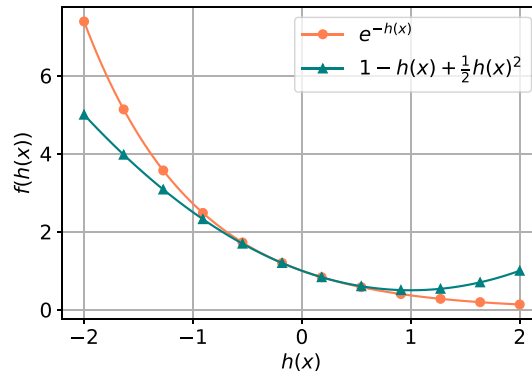


Figure 3. Comparison between $e^{-h(x)}$ and the unbalanced function $1 - h(x) + \frac{1}{2}h(x)^2$.

Table 1. Parameters $\lambda_{0,1,2}$ for the TSP, BPP, and KP used to translate the combinatorial optimization problems into the QUBO representation using the unbalanced penalization approach (equation (27)). For all the equality constraints of each problem, we use the same λ_0 and for the inequality constraints the same $\lambda_{1,2}$.

	λ_0	λ_1	λ_2
TSP	38.2584	18.2838	57.0375
BPP	20.5198	7.2949	0.8583
KP	—	0.9603	0.0371

To tune the $\lambda_{0,1,2}$ parameters of equation (27) for the TSP, the BPP, and the KP, we use the Nelder-Mead optimization method to explore the set of $\lambda_{0,1,2}$ values that brings the optimal solution of the original combinatorial optimization problems as close as possible to the ground state of the cost Hamiltonian, equation (10), of their QUBO representation. For the TSP we use a random problem with 4 cities, for the BPP we use a random problem with 4 items, and for the KP we use a problem with 13 items. The set of values $\lambda_{0,1,2}$ found by the Nelder-Mead optimization method for each problem are presented in table 1. After we find the best set of values for the specific problem, we test if those values generalize to other random cases. In the random cases used, we find that a set of $\lambda_{0,1,2}$ of one problem works well for different and/or larger problems. Therefore, we use a unique set of these parameters for each problem. However, this is not necessarily the case for other instances of the same problem with a different distribution of parameters. For example, the TSP has some instances that are harder to solve [48]. In general, one can repeat the procedure mentioned above with multiple cases of the specific distribution to estimate a good set of $\lambda_{0,1,2}$ parameters on the new distribution. The results on these aspects of the unbalanced penalization method are presented in section 4.1.

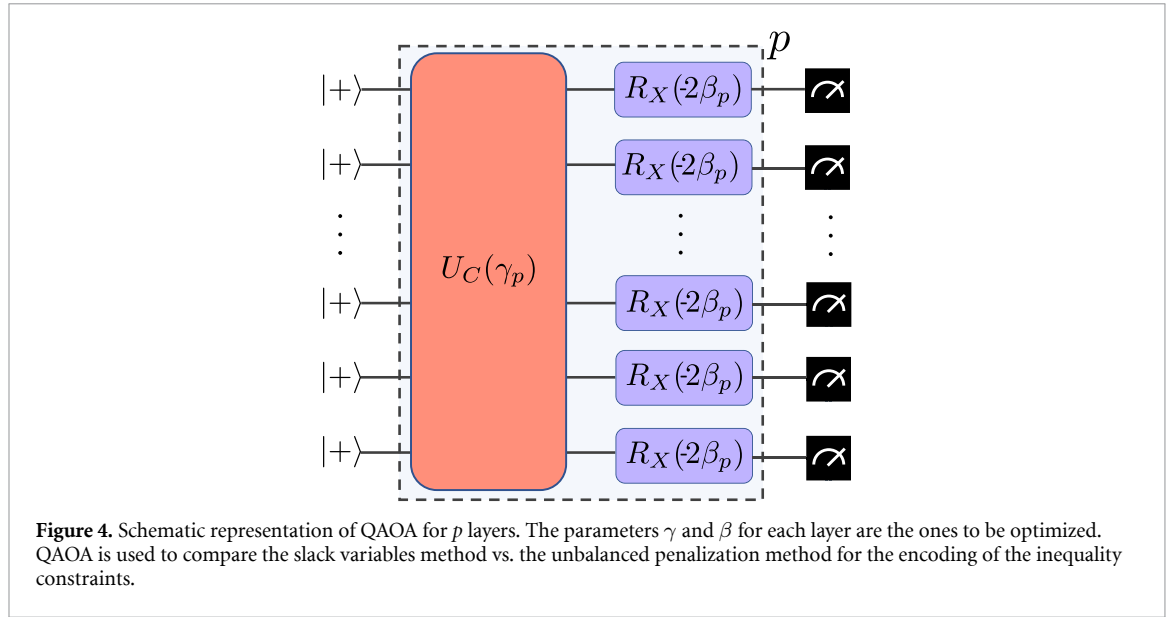
3.1. Quantum annealing

Quantum annealing is a computational technique designed to tackle optimization problems by leveraging quantum mechanical effects in the search for low-energy states of Ising Hamiltonians [13, 49]. By employing this method, the system can explore a broader solution space and potentially find good quality solution more efficiently than classical optimization approaches.

In order to validate the effectiveness of our proposed method on real hardware, we conducted experiments using the *D-Wave Advantage* and the *D-Wave Hybrid Quantum-Classical* solver to obtain results for the BPP. To ensure a fair comparison, we selected random BPP problems with the same settings as described in section 2.3, without any simplifications.

For the experiments using the *D-Wave Advantage*, we conducted 20 trials on random problem instances ranging from 3 to 8 nodes (items). Each trial consisted of 5000 samples. For the hybrid solver, we performed 20 experiments with a single run, allowing a maximum execution time of 3 s. The problem instances for the hybrid solver ranged from 5 to 31 nodes (items).

In both cases, we utilized the $\lambda_{0,1,2}$ parameters as presented in table 1. It is important to note that our objective was not to fine-tune the hyperparameters for either solver, but rather to demonstrate that under similar conditions, unbalanced penalization yields significantly better solutions in terms of both quantity and quality. Results are presented in section 4.2.



3.2. QAOA

To test the difference between the slack variables approach vs. the unbalanced penalization encoding on different random problems of the TSP, the BPP, and the KP, we use the QAOA with one layer. It is important to acknowledge that, while the impact of penalization terms of the constraints remains uncertain in terms of their influence on the performance of the QAOA [50], we employ QAOA to demonstrate the contrasting probability distributions resulting from the two approaches. In this case, the cost Hamiltonian, H_c , obtained from the QUBO formulation, is translated into a parametric unitary gate given by

$$U(H_c, \gamma) = e^{-i\gamma H_c}, \quad (28)$$

where γ is a parameter to be optimized. A second unitary operator applied is

$$U(B, \beta) = e^{i\beta X}, \quad (29)$$

where β is the second parameter that must be optimized and $X = \sum_{i=1}^n \sigma_i^x$ with σ_i^x the Pauli-x quantum gate applied to qubit i . The general QAOA circuit is shown in figure 4. Here, $R_X(\theta) = e^{-i\frac{\theta}{2}\sigma_x}$, p represents the number of repetitions of the unitary gates equations (28) and (29) with each repetition having separate values for γ_p and β_p , and the initial state is a superposition state $|+\rangle^{\otimes n}$. The case of QAOA with $p = 1$ is suitable for visualization of the landscape and in section 4.3, we show a comparison between both encodings and even though they are similar, the addition of extra qubits for the slack variables approach has a great impact on the probability of finding the optimal solution of the original combinatorial optimization problem.

3.3. CoP

Additionally, we introduce the CoP, a quantifier that expresses how good a quantum algorithm is to give a specific solution compared to random guessing over a set of solutions. In this perspective, we use the CoP to determine how likely it is to find the optimal solution for the TSP, the BPP, and the KP on the local minima of the QAOA algorithm for $p = 1$. The CoP is given by

$$\text{CoP} = \frac{P(*x)}{P_R}, \quad (30)$$

where $P(*x)$ is the probability of getting a bit string $*x$ using a quantum algorithm, e.g, QAOA and $P_R = 1/2^n$ is the probability of a random guessing with n equal to the number of qubits involved in the problem.

4. Results

4.1. Optimal solution distribution

To determine the applicability of our method, we sorted the energy eigenvalues of the cost Hamiltonian for the different combinatorial optimization problems. The parameter of interest is the position of the energy eigenvalue that describes the optimal solution of the original combinatorial optimization problems within

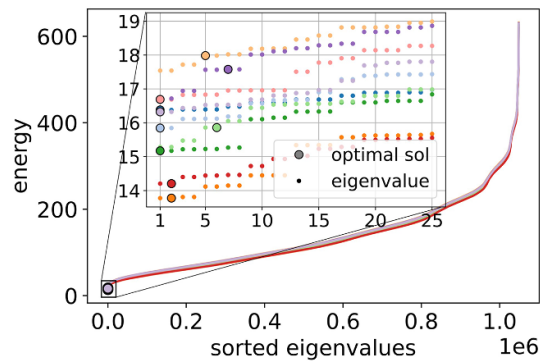


Figure 5. Eigenvalues distribution for the TSP with 5 cities (20 qubits) for 10 randomly generated problems using the unbalanced penalization method. The inset shows the lowest 25 energy eigenvalues. The big circles are the optimal solutions for the random problem and the small circles are the different eigenvalues of the cost Hamiltonian of its QUBO representation. Note that each eigenvalue is degenerate with multiplicity two, because the problem is symmetric, e.g. one clockwise solution and another anti-clockwise.

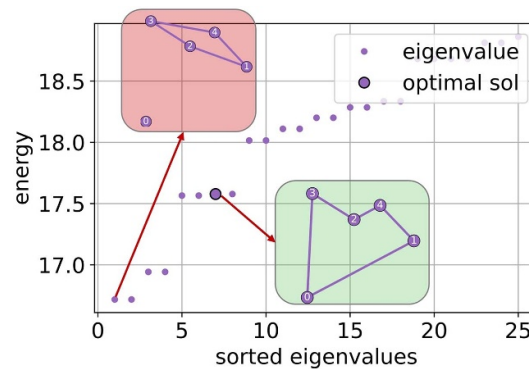


Figure 6. The first 25 sorted eigenvalues of the Hamiltonian H_c corresponding to the worst case shown in figure 5. The big circle corresponds to the optimal solution of this problem and the small circles are the other eigenvalues. The solution shown in the red inset is the ground state, which is an invalid solution of the TSP. The solution in the green inset is the optimal solution.

the list of all sorted energy eigenvalues of the cost Hamiltonian of their QUBO representation. If the position of the optimal solution is far from the ground energy of the cost Hamiltonian, the unbalanced penalization encoding is not working as expected. Otherwise, if the energy eigenvalue corresponding to the optimal solution is close to the ground state energy given the total number of eigenvalues, the unbalanced encoding is advantageous.

Figure 5 shows the energy eigenvalues of 10 random cases for the TSP with 5 nodes using the unbalanced penalization encoding approach. The eigenvalues are sorted in terms of their energy from minimum to maximum with the x -axis representing the eigenvalue position (number) of eigenvalues, and the y -axis its energy. In the inset, showing the lowest 25 energy eigenvalues, the big circles represent where the optimal result is positioned compared to all the other eigenvalues (small circles) of the cost Hamiltonian. Here, the worst eigenvalue position is 7 out of $2^{20} = 1048576$ eigenvalues. This means that in the worst scenario the optimal solution is among the lowest 0.00057% of all eigenvalues.

From figure 5 it is seen that the unbalanced penalization approach does not ensure that the optimal solution is the lowest eigenvalue of the cost Hamiltonian, but in all the cases analyzed, it is very close to it. For example, in figure 6, this feature is illustrated for one of the random problems of figure 5. Here, there are a couple of eigenvalues that do not fulfill all the restrictions of the problem but have lower energy than the optimal solution. In the end, the unbalanced penalization encoding is a trade-off between adding slack variables and ensuring the ground state of the cost Hamiltonian corresponds to the optimal solution of the combinatorial optimization problem, or reducing the number of variables but encoding the optimal solution in the vicinity of the ground state. In many cases, there is no preference for encoding the optimal solution in the lowest eigenvalue of the cost Hamiltonian. For instance, QAOA brings the expectation value of a cost Hamiltonian to a region of overall minimum energy, so we can expect that the optimal and sub-optimal solutions occur with reasonable probability in regions of minimal energy.

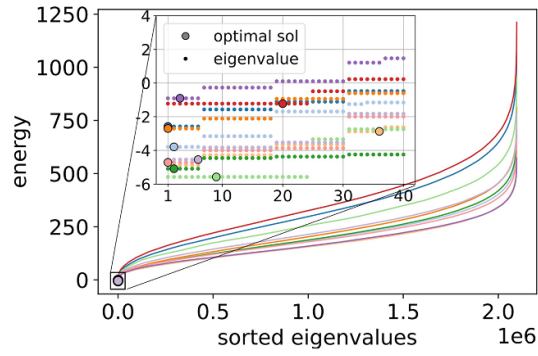


Figure 7. Eigenvalue distribution for the BPP with 5 items (21 qubits) for 10 randomly generated problems using the unbalanced inequality-constrained penalization method. The inset shows the 40 lowest energy eigenvalues. The big circles are the optimal solutions for the random problem and the small circles are the different eigenvalues for the cost Hamiltonian of its QUBO representation. Note that each eigenvalue has a different degeneracy, because the problem has some symmetries.

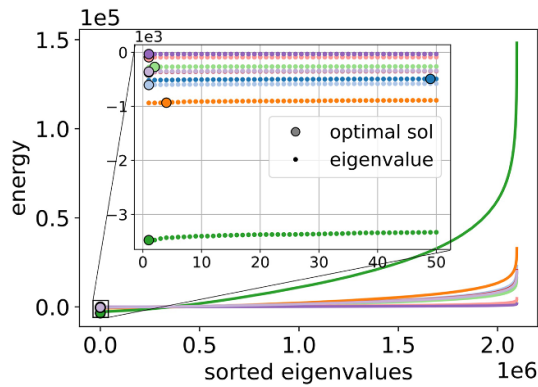


Figure 8. Eigenvalue distribution for the KP with 21 items (21 qubits) for 10 randomly generated problems using the unbalanced penalization method. The inset shows the 50 lowest eigenvalues. The big circles give the position of the optimal solutions for the random problems and the small circles depict the position of the different eigenvalues of the cost Hamiltonian. In this case, there are no degeneracies but the energies are close to each other.

Figure 7 shows the energy eigenvalues sorted for the cost Hamiltonian of the BPP for 10 random problems with 5 items. In the inset, showing the 40 lowest eigenvalues, the optimal solutions of the cost Hamiltonians are depicted as big circles and the small circles are the other eigenvalues of the cost Hamiltonian. In this case, there is an increasing number of degeneracies explained by the symmetries of the problem (exchange items of one bin with the others). To keep the same number of qubits (21) for all the problems, we select 10 random cases that require more than 3 bins to store the items. Here, the worst case is in the position 36 out of $2^{21} = 2097152$ eigenvalues which means it is within 0.00124%, of all eigenvalues.

Finally, figure 8 shows the eigenvalue distribution for 10 random cases of the KP with 21 items (21 qubits). In this case, there are no degeneracies, but the possible solutions are close to each other. The inset shows the first 50 eigenvalues. In this case, for one of the random cases the optimal solution is located at position 49 out of $2^{21} = 2097152$ eigenvalues, which means it is within 0.0023% of the first eigenvalues. In this case, the parameters $\lambda_{0,1,2}$ were optimized for a random case with 13 items, which suggests great generalization capabilities of the current method.

4.2. Quantum annealing

Figure 9 illustrates the results obtained using the D-Wave Advantage. The upper x-axis in this figure, as well as in subsequent plots, represents the logical variables required to represent each problem size, as per the formulation provided in section 2.3. Across all cases, the unbalanced penalization approach consistently generates more optimal solutions compared to the slack variables approach. Although the optimal solution cannot be guaranteed to be the ground state of the Hamiltonian, more optimal solutions are found for the unbalanced penalization method.

For the scenario involving 8 nodes, neither of the two methods were able to find valid solutions. However, it is important to acknowledge that the number of statistical samples used in these experiments

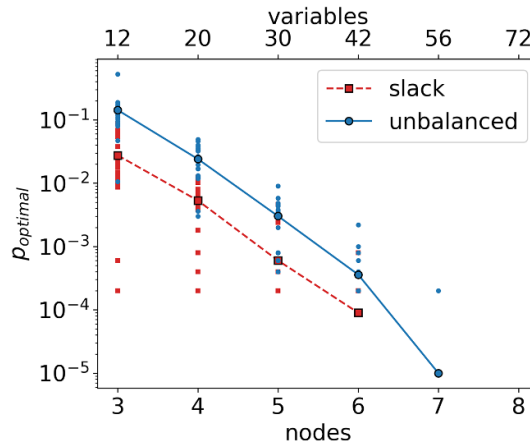


Figure 9. D-Wave Advantage success rate of finding optimal solutions for different cases of the BPP ranging from 3 to 8 nodes (items). The 20 small circles (squares) for each problem size consist of 5000 samples for the unbalanced (slack) method.

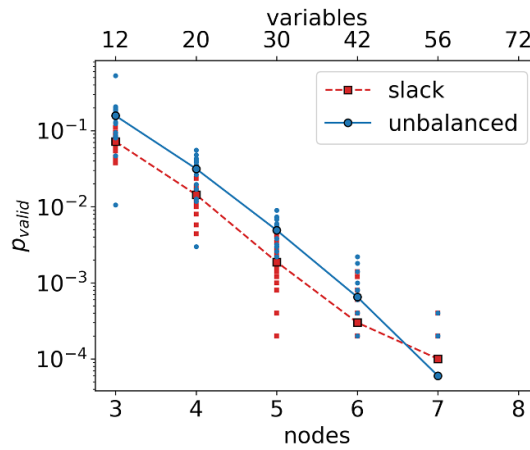


Figure 10. D-Wave Advantage success rate of finding valid solutions for different cases of the BPP ranging from 3 to 8 nodes (items).

might not have been sufficient to draw conclusions about the limit beyond which no further solutions are found on the quantum annealer.

Figure 10 presents the results of the same experiments depicted in figure 9, focusing on the quantity of valid solutions obtained. Consistent with the observations in figure 9, the employment of the unbalanced penalization method enhances the likelihood of discovering valid solutions compared to the slack approach. However, it is important to note that in the case of 7 nodes, where the slack approach appears to exhibit superior performance, more samples of valid solutions (than could be produced by the machine) would be needed to draw a solid conclusion.

In order to investigate the viability of finding valid solutions for larger problem sizes, we employed the D-Wave Hybrid solver. We conducted experiments using problem sizes ranging from 5 to 31 nodes (items), with intervals of 2, and executed 20 random problem instances for each case. Each problem instance was solved using a single run on the D-Wave Hybrid solver, with a maximum time limit of 3 s.

Figure 11 depicts the probability of finding optimal solutions across the 20 random problems for each node case. Similar to the observations with the D-Wave Advantage, the unbalanced penalization method consistently outperforms the slack approach. It successfully discovers optimal solutions for nearly all cases, with the exception of 19, 23, 27, and 31 nodes. Conversely, the slack approach fails to find optimal solutions after 7 nodes.

Figure 12 illustrates the probability of finding valid solutions using both the unbalanced penalization method and the slack variables approach on the D-Wave Advantage. Consistent with the findings in figure 11, the unbalanced penalization method demonstrates superior performance compared to the slack variables approach.

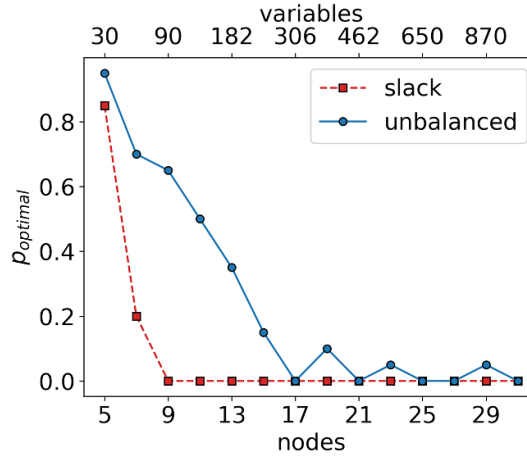


Figure 11. D-Wave Hybrid solver success rate of finding optimal solutions for different cases of the BPP ranging from 5 to 31 nodes (items). Each point consists of a single execution of 20 random problems with a maximum time for the hybrid solver of 3 s.

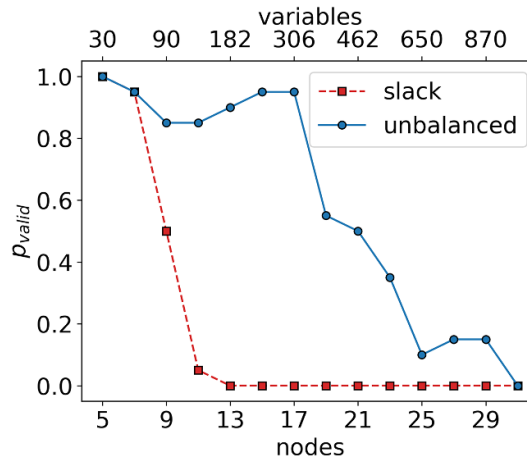


Figure 12. D-Wave Hybrid solver success rate of finding valid solutions for different cases of the BPP ranging from 5 to 31 nodes (items). Each circle (square) represents the mean value of 20 random problems with a maximum time for the hybrid solver of 3 s for the unbalanced (slack) method.

Up to 17 nodes, the unbalanced penalization method consistently discovers valid solutions, whereas the slack variables approach only manages to find valid solutions up to 11 nodes. Beyond 17 nodes, the probability of finding valid solutions gradually declines for the unbalanced penalization method, with no solutions being found for the 31-node case. This behavior can be attributed to the fact that the maximum time limit remains constant throughout the experiments.

Even though increasing the time will impact positively on improving the performance of both methods, it is important to reiterate that the primary objective of this section is not to provide high-quality solutions, but rather to evaluate the performance of both methods under similar conditions.

4.3. QAOA landscape

In this section, the results of the cost Hamiltonian energy landscape using the QAOA algorithm with one layer are presented for the TSP, the BPP, and the KP. The slack variables method and the unbalanced penalization method are used for the cost Hamiltonian encoding. For visualization purposes, the energy landscape for all cases is limited to the region of $\gamma \in (-\pi/\max\{q_{ij}\}, 0)$ where q_{ij} is given in equation (10) and $\beta \in (-\pi/2, 0)$.

Figure 13 shows the energy landscape of the QAOA algorithm for a TSP with 4 cities. This is equivalent to a 17 and 12 qubits problem for the slack variables and the unbalanced penalization encodings, respectively. The probabilities for finding the optimal solutions of the TSP are 0.02% and 0.227% for the slack variables and the unbalanced encodings, respectively. This shows a clear advantage of the unbalanced encoding method that increases the probability more than 10 times compared to the probability obtained with the slack variables encoding. This characteristic is mainly attributed to the difference in the number of qubits needed to represent the problem and thus the size of the set of possible solutions. For the slack variables

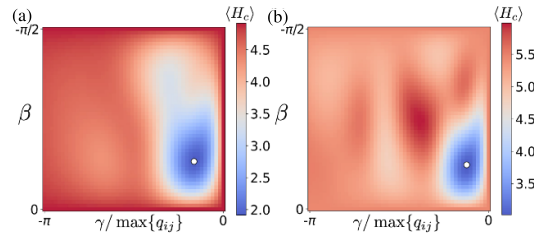


Figure 13. QAOA energy for $p = 1$ for the cost Hamiltonian $\langle H_c \rangle$ of the TSP with 4 cities with (a) slack variables encoding and (b) unbalanced penalization encoding. The minimum energy is indicated with a white dot. The probability that this minimum corresponds to the optimal solution of the TSP problem is $P(*x) = 0.02\%$ and $P(*x) = 0.227\%$ for encoding (a) and (b), respectively.

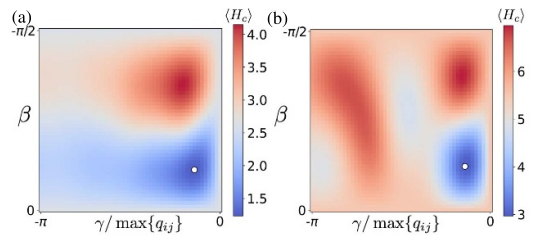


Figure 14. Same as the figure 13 for the BPP with 3 items. The probability that this minimum corresponds to the optimal solution of the BPP problem is $P(*x) = 0.001\%$ and $P(*x) = 10.661\%$ for encoding (a) and (b), respectively.

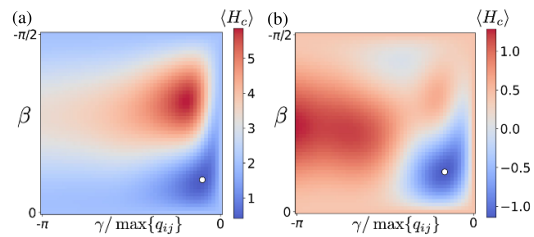


Figure 15. Same as the figure 13 for the KP with 10 items. The probability that this minimum corresponds to the optimal solution of the KP problem is $P(*x) = 0.001\%$ and $P(*x) = 0.346\%$ for encoding (a) and (b), respectively.

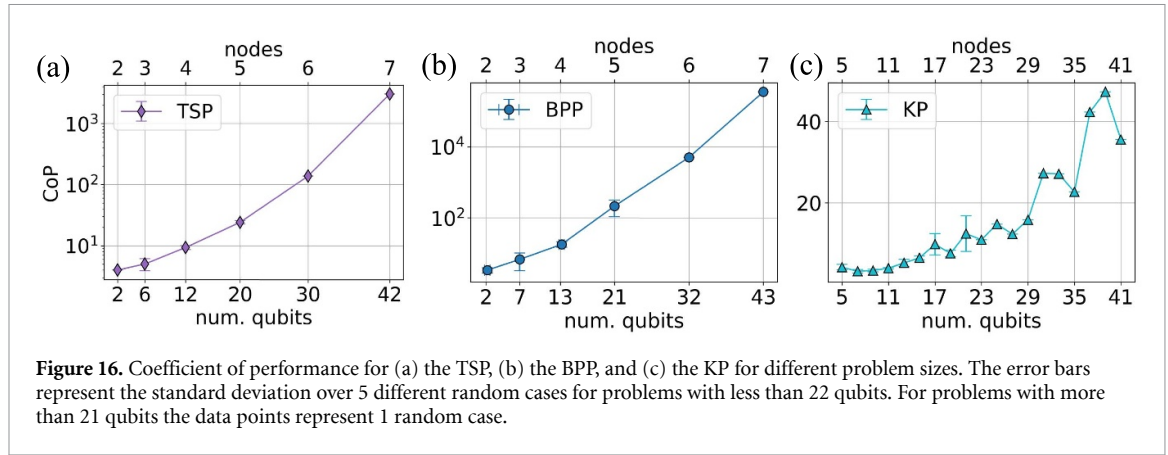
encoding there are 131 072 (2^{17}) possible solutions while for the unbalanced penalization encoding there are 4096 (2^{12}) possible solutions.

Figure 14 shows the energy landscape of the QAOA algorithm for a BPP with 3 items, which is equivalent to a 6 and 19 qubits problem for the unbalanced and slack encodings, respectively. The probability of finding the optimal solution of the BPP in the local minima of QAOA is 0.001% for the slack variables encoding and 10.66% for the unbalanced penalization method. Here, the optimal solution is 10 000 times more likely to be found using the unbalanced penalization approach than using the slack variables approach.

Finally, figure 15 shows the energy landscape of the QAOA algorithm for a KP with 10 items, which is equivalent to a 18 and 10 qubits problem for the slack variables and the unbalanced penalization encoding, respectively. Here, the probability to find the optimal solution of the KP is 0.001% and 0.346% for the slack variables and the unbalanced penalization encoding. This means finding the solution with the unbalanced penalization method is 346 times more likely than with the slack variables encoding.

4.4. The CoP

Figure 16 shows the CoP for different problem sizes for the TSP, BPP, and KP using the unbalanced penalization method. All the cases show the probability of finding the optimal solution for each problem using QAOA with 1 layer in the minimum point of the energy landscape region $\beta \in (-\pi/2, 0)$ and $\gamma \in (-\pi/\max\{q_{ij}\}, 0)$. At this point, we do not use a classical optimizer to find the optimal beta and gamma, but instead, we divide the beta and gamma region in a 50×50 grid landscape and take the probability of the optimal solution at the minimum energy. It is expected that using an optimization method for this step improves the CoP while keeping the same tendency. Figure 16(a) shows the CoP for the TSP problem with 2, 3, 4, 5, 6, and 7 cities (2, 6, 12, 20, 30, and 42 qubits). For the cases with 30 and 42 qubits, we use JUQCS [35,



36] in its GPU-accelerated version [37] to obtain the results. Note that the same TSP cases using the slack variables approach will require 2, 7, 17, 36, 73, 148, and 304 qubits. We thus explore problem sizes (5, 6, and 7 cities) with the unbalanced penalization approach that are unfeasible using the slack approach on quantum computer simulators. Interestingly, the CoP is increasing exponentially with the problem size, which means that the probability of finding the optimal solution exponentially increases when we compare it with random guessing in the set of possible configurations (see [51]).

Figure 16(b) shows the CoP for the BPP for 2, 3, 4, 5, 6, and 7 items (2, 6, 13, 21, 32, and 43 qubits). For this case, the CoP is improving even more rapidly compared to TSP and also here for the BPP with 6 and 7 items the solutions using slack variables (61 and 78 qubits) is beyond range of what can be simulated with quantum computer simulators. The largest simulations up to 43 qubits were performed on JUWELS Booster [52] as they required more than 1/8 PiB of distributed memory and took more than one million core hours. Finally, figure 16(c) shows the CoP for the KP with 5 to 41 items (5 to 41 qubits) with increments of two (same number of qubits). In this case, the CoP increases linearly with the number of qubits. This is a poor performance if we compare it with the TSP and BPP. We suspect the improved CoP is related to the number of constraints of the problem.

5. Conclusions

We have presented *unbalanced penalization*, a new method to encode inequality constraints of combinatorial optimization problems into QUBO penalizations. The method does not require extra slack variables to encode the inequality constraints. This is extremely beneficial because there is no increase in the number of variables or qubits needed to encode the problems, especially for the TSP and the BPP.

The method is suitable for VQA on gate-based quantum computers and QA on quantum annealers, which both require expressing the combinatorial optimization problems in terms of QUBOs. We have tested the method using QA for the BPP and QAOA with 1 layer for the TSP, BPP, and KP. We have shown that in 10 random cases for different problem sizes the optimal solution of the combinatorial optimization problem is located in the vicinity of the ground state energy of the cost Hamiltonian. The method is highly generalizable in the sense that the tuned parameters $\lambda_{0,1,2}$ for small problems generalize well to larger problems. For QA, we have demonstrated that the unbalanced penalization method outperforms the slack variables approach, finding better solutions in terms of quality and quantity. Additionally, the probabilities to find the optimal solutions among the local minima of the unbalanced penalization (figures 13–15) are much better than using slack variables. Qualitatively, this can be understood because the addition of slack variables increases the search space and therefore reduces the probability of finding optimal and nearly-optimal solutions.

The proposed method was evaluated on both the D-Wave Advantage and D-Wave Hybrid solvers. Our findings demonstrate that the method utilizing unbalanced penalization consistently outperforms the slack variables approach in solving the BPP. Specifically, the unbalanced penalization method successfully finds solutions for problem instances with up to 31 items using the Hybrid solver, while the slack variables approach achieves solutions for only up to 11 items. We remark that additional evaluations of this method on D-Wave systems for the TSP support the same conclusions [53].

We have tested the unbalanced penalization method close to the limits of what is possible to be simulated on supercomputers using JUQCS. For the case of the TSP, we went up to 42 qubits, for the BPP up to 43 qubits, and for the KP up to 41 qubits. Indeed, we are exploring regions unfeasible using the slack variables method. For example, for the TSP with 6 and 7 cities using the slack variables method it requires 73 and 148

qubits, and for the BPP with 6 and 7 items it requires 61 and 78 qubits. In terms of probability, the unbalanced penalization method shows large advantages for the small problems we made a comparison for. For the TSP with 4 cities, the probability of finding the optimal solution at the local minimum energy of the QAOA with $p = 1$ using the unbalanced penalization method is more than 10 times larger than the one using the slack variables approach. For the BPP with 3 items, the probability for finding the optimal solution with the unbalanced penalization method is more than 10 000 times larger, and for the KP with 10 items it is almost 346 higher.

Finally, we have presented a new metric, the CoP, which compares the probability of obtaining a specific solution using a quantum algorithm against random guessing over the set of possible solutions. The CoP is useful to characterize how an algorithm performs with increasing problem size or to compare how different combinatorial optimization problems perform for a specific algorithm. We have shown the CoP for the TSP, BPP, and KP using QAOA with $p = 1$ in the region $\gamma \in (-\pi/\max\{q_{ij}\}, 0)$ and $\beta \in (-\pi/2, 0)$. Interestingly, for all of them, the CoP increases with the problem size. For the KP the increase is linear while for the TSP and BPP the increase is exponential. Further research is worthwhile to see if this provides an implicit advantage of using QAOA to solve combinatorial optimization problems, what the characteristics of the TSP and the BPP are that make the CoP increase exponentially, and to see how increasing the number of QAOA layers improves the CoP.

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: <https://jugit.fz-juelich.de/qip/unbalanced-penalizations-qubo>.

Acknowledgments

We would like to thank Leonardo Disilvestro and Vishal Sharma from Entropica Labs for helpful comments and suggestions about the results of the present work. J A Montañez-Barrera acknowledges support by the German Federal Ministry of Education and Research (BMBF), funding program Quantum technologies—from basic research to market, project QSolid (Grant No. 13N16149). D Willsch acknowledges support from the project Jülich UNified Infrastructure for Quantum computing (JUNIQU) that has received funding from the German Federal Ministry of Education and Research (BMBF) and the Ministry of Culture and Science of the State of North Rhine-Westphalia. The authors gratefully acknowledge the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for funding this project by providing computing time on the GCS Supercomputer JUWELS at Jülich Supercomputing Centre (JSC).

ORCID iDs

J A Montañez-Barrera  <https://orcid.org/0000-0002-8103-4581>

Dennis Willsch  <https://orcid.org/0000-0003-3855-5100>

A Maldonado-Romo  <https://orcid.org/0000-0002-6342-3508>

Kristel Michielsen  <https://orcid.org/0000-0003-1444-4262>

References

- [1] Lucas A 2014 *Front. Phys.* **2** 12
- [2] Kochenberger G, Hao J K, Glover F, Lewis M, Lü Z, Wang H and Wang Y 2014 *J. Comb. Optim.* **28** 58
- [3] Ohzeki M 2020 *Sci. Rep.* **10** 1
- [4] Harrigan M P et al 2021 *Nat. Phys.* **17** 332
- [5] Niroula P, Shaydulin R, Yalovetzky R, Minssen P, Herman D, Hu S and Pistoia M 2022 arXiv:2206.06290
- [6] Preskill J 2018 *Quantum* **2** 1
- [7] Shor P W 1994 *Proc. 35th Annual Symp. on Foundations of Computer Science* pp 124–34
- [8] Shor P W 1997 *SIAM J. Comput.* **26** 1484
- [9] Cerezo M et al 2021 *Nat. Rev. Phys.* **3** 1
- [10] Khairy S, Shaydulin R, Cincio L, Alexeev Y and Balaprakash P 2020 *Proc. AAAI Conf. on Artificial Intelligence* vol 34 p 2367
- [11] Apolloni B, Carvalho C and de Falco D 1989 *Stoch. Process. Their Appl.* **33** 233
- [12] Finnila A B, Gomez M A, Sebenik C, Stenson C and Doll J D 1994 *Chem. Phys. Lett.* **219** 343
- [13] Kadowaki T and Nishimori H 1998 *Phys. Rev. E* **58** 5355
- [14] de Falco D and Tamascelli D 2011 *RAIRO Theor. Inf. Appl.* **45** 99–116
- [15] Ayanzadeh R, Dorband J, Halem M and Finin T 2021 *Sci. Rep.* **11** 1
- [16] Willsch D, Willsch M, Gonzalez Calaza C D, Jin F, De Raedt H, Svensson M and Michielsen K 2022 *Quantum Inf. Process.* **21** 141
- [17] Heim B, Rønnow T F, Isakov S V and Troyer M 2015 *Science* **348** 215
- [18] Yan B and Sinitsyn N A 2022 *Nat. Commun.* **13** 1
- [19] Tasseff B, Albash T, Morrell Z, Vuffray M, Lokhov A Y, Misra S and Coffrin C 2022 arXiv:2210.04291

- [20] Farhi E, Goldstone J and Gutmann S 2014 A Quantum approximate optimization algorithm (arXiv:1411.4028)
- [21] Willsch M, Willsch D, Jin F, De Raedt H and Michielsen K 2020 *Quantum Inf. Process.* **19** 197
- [22] Jiang H, Shen Z-J M and Liu J 2022 arXiv:2209.08246
- [23] Orús R, Mugel S and Lizaso E 2019 *Rev. Phys.* **4** 100028
- [24] Souza A M, Martins E O, Roditi I, Sá N, Sarthour R S and Oliveira I S 2022 *Front. Phys.* **9** 1
- [25] Mugel S, Kuchkovsky C, Sánchez E, Fernández-Lorenzo S, Luis-Hita J, Lizaso E and Orús R 2022 *Phys. Rev. Res.* **4** 1
- [26] Sharabiani M T A, Jakobsen V B, Jeppesen M and Mahani A S 2021 arXiv:2105.11322
- [27] Urgelles H, Picazo-martinez P, Garcia-roger D and Monserrat J F 2022 Multi-objective routing optimization for 6G communication networks using a quantum approximate optimization algorithm *Sensors* **22** 7570
- [28] Luckow A, Klepsch J and Pichlmeier J 2021 *Digitale Welt* **5** 38
- [29] Inoue D, Okada A, Matsumori T, Aihara K and Yoshida H 2021 *Sci. Rep.* **11** 1
- [30] IBM Q team 2021 Qiskit: an open-source framework for quantum computing (release 0.39.4) Zenodo <https://doi.org/10.5281/zenodo.2573505>
- [31] D-Wave Systems 2022 D-wave ocean SDK (release 6.2.0) (available at: <https://github.com/dwavesystems/dwave-ocean-sdk>)
- [32] Glover F, Kochenberger G and Du Y 2019 *4OR* **17** 335
- [33] Glover F, Kochenberger G, Ma M and Du Y 2022 *Ann. Oper. Res.* **314** 185
- [34] Sharma V, Saharan N S B, Chiew S-H, Chiacchio E I R, Disilvestro L, Demarie T F and Munro E 2022 1–10 arXiv:2210.08695
- [35] De Raedt K, Michielsen K, De Raedt H, Trieu B, Arnold G, Richter M, Lippert T, Watanabe H and Ito N 2007 *Comput. Phys. Commun.* **176** 121
- [36] De Raedt H, Jin F, Willsch D, Willsch M, Yoshioka N, Ito N, Yuan S and Michielsen K 2019 *Comput. Phys. Commun.* **237** 47
- [37] Willsch D, Willsch M, Jin F, Michielsen K and De Raedt H 2022 *Comput. Phys. Commun.* **278** 108411
- [38] Verma Amit and Lewis Mark 2022 *Disc. Optim.* **44** 100594
- [39] Helsgaun K 2006 An effective implementation of K-opt moves for the Lin-Kernighan TSP heuristic pp 1–99 (available at: <https://api.semanticscholar.org/CorpusID:12065848>)
- [40] Applegate D L, Bixby R E, Chvátal V and Cook W J 2006 *The Traveling Salesman Problem: A Computational Study* (Princeton University Press)
- [41] Applegate D L, Bixby R E, Chvátal V and Cook W J 1997 Concorde TSP solver (available at: www.math.uwaterloo.ca/tsp/concorde.html)
- [42] 1995 TSPLIB95: a library of traveling salesman and related problems (available at: <http://comopt.ifl.uni-heidelberg.de/software/TSPLIB95/>)
- [43] Grötschel M and Nemhauser G L 2008 *Disc. Optim.* **5** 168
- [44] Martello S and Toth P 1990 *Knapsack Problems: Algorithms and Computer Implementations* (Wiley, Inc.)
- [45] Heßler K, Irnich S, Kreiter T and Pferschy U 2020 Lexicographic bin-packing optimization for loading trucks in a direct-shipping system (available at: <https://ideas.repec.org/p/jgu/wpaper/2009.html>)
- [46] Yan J, Lu Y, Chen L, Qin S, Fang Y, Lin Q, Moscirobroda T, Rajmohan S and Zhang D 2022 *Proc. 28th ACM SIGKDD Conf. on Knowledge Discovery and Data Mining, 2022, (Washington, DC, USA)* vol 1 (Association for Computing Machinery)
- [47] Kroes M, Petrica L, Cotofana S and Blott M 2020 *GECCO 2020 - Proc. 2020 Genetic and Evolutionary Computation Conf.* p 1125
- [48] Cheeseman P, Kanefsky B and Taylor W M 1991 *The 12nd Int. Joint Conf. on Artificial Intelligence* p 331
- [49] Johnson M et al 2011 *Nature* **473** 194
- [50] Wang Z, Rubin N C, Dominy J M and Rieffel E G 2019 *Phys. Rev. A* **101** 012320
- [51] Willsch D, Jattana M, Willsch M, Schulz S, Jin F, De Raedt H and Michielsen K 2022 *NIC Symp. 2022 (Publication Series of the John von Neumann Institute for Computing (NIC) NIC Series vol 51)* ed M Müller, C Peter and A Trautmann (Forschungszentrum Jülich GmbH Zentralbibliothek, Verlag) pp 31–44
- [52] Kesselheim S et al 2021 JUWELS booster – a supercomputer for large-scale AI research (arXiv:2108.11976 [cs.DC])
- [53] Montanez-Barrera J A, van den Heuvel P, Willsch D and Michielsen K 2023 arXiv:2305.18757