

Resource-Efficient Quantum Optimization via Higher-Order Encoding

Frederik Koch¹, Shahram Panahiyan^{1,2}, Rick Mukherjee^{3,4,5}, Joseph Doetsch⁶, and Dieter Jaksch^{1,7,8}

¹Institute for Quantum Physics, University of Hamburg, Luruper Chaussee 149, 22761 Hamburg, Germany

²Max Planck Institute for the Structure and Dynamics of Matter, Luruper Chaussee 149, 22761 Hamburg, Germany

³Zentrum für Optische Quantentechnologien, University of Hamburg, Luruper Chaussee 149, 22761 Hamburg, Germany

⁴Department of Physics & Astronomy, University of Tennessee, Chattanooga, TN 37403, USA

⁵UTC Quantum Center, University of Tennessee, Chattanooga, TN37403, USA

⁶Lufthansa Industry Solutions, Südportal 7, 22848 Norderstedt, Germany

⁷The Hamburg Centre for Ultrafast Imaging, Luruper Chaussee 149, 22761 Hamburg, Germany

⁸Clarendon Laboratory, University of Oxford, Parks Road, Oxford OX1 3PU, United Kingdom

Quantum approaches to combinatorial optimization problems (COPs) are often limited by the resource demands of Quadratic Unconstrained Binary Optimization (QUBO) encodings, which enlarge circuits through penalty terms and increase qubit and gate counts. We show that Higher-Order Unconstrained Binary Optimization (HUBO) enables a more resource-efficient formulation. Our method systematically constructs HUBO Hamiltonians and, compared to QUBO in benchmarks on Gate Assignment (GAP), Maximum k-Colorable Subgraph (MkCS), and Integer Programming (IP) problems, exponentially reduces qubit requirements and decreases CNOT gate counts by at least 89.6% after compilation to single- and two-qubit gates for all tested instances. These results highlight HUBO as a practical alternative for current and near-term devices. To promote adoption, we release an open-source Python library that automates HUBO model construction, broadening access to resource-efficient quantum optimization.

Combinatorial optimization problems are prominent in numerous scientific and industrial domains, ranging from logistics and finance [1] to computational biology [2] and operations research [3]. Challenging problems require finding




Figure 1: Number of CNOT gates required to reach solution threshold for three COPs, optimized by QAOA with QUBO (yellow marker) and HUBO (blue marker) encodings as a function of problem size $n \times m$. The GAP optimizes the assignment of n aircraft to m airport gates. In the MkCS problem, n vertices are colored with m colors, IP optimizes a polynomial over n variables, each taking one of m possible values. Further details and quantum encodings are provided in the main text.

(near)-optimal solutions within vast discrete configuration spaces, where the search complexity often grows exponentially with problem size [4]. Theoretical insights from statistical physics have illuminated these challenges through their connection to disordered systems, revealing that most industrial optimization problems belong to the NP-hard complexity class [5].

Classical methods, including metaheuristics such as simulated annealing and tabu search [6, 7], and commercial solvers like CPLEX and Gurobi [8, 9] face significant limitations when confronting non-convex energy landscapes. Even

relatively small problems (sometimes below 100 variables) remain computationally intractable for them [10–15].

Quantum algorithms, such as adiabatic quantum optimization [16, 17] and the Quantum Approximate Optimization Algorithm (QAOA) [18] are poised to address these challenges. This is predominantly done by exploiting mappings of COPs to QUBO models [19–21] which is further facilitated by modeling tools such as Pyqubo [22]. Previous studies [23, 24] have identified several practical limitations of QUBO formulations, including large qubit requirements, costly penalty terms for constraint enforcement, and consequently increased circuit depth with large CNOT gates, which motivate the search for more resource-efficient alternatives.

In contrast, HUBO encodings directly incorporate such constraints and, therefore, can reduce resource demands. While demonstrated in specific settings, including protein folding and the traveling salesperson problem [25–28], prior work has lacked a general framework. Moreover, practical concerns remain regarding hardware implementation: higher-order terms formally require multiqubit interactions, whereas present quantum processors primarily support operations involving just one or two qubits at a time [29, 30]. In this work, we show how such circuits can be implemented entirely with single- and two-qubit gates.

By comparing to QUBO encodings, we demonstrate that HUBO encodings provide substantial and consistent resource advantages for three representative COPs: GAP, MkCS, and IP. These advantages persist after full circuit compilation, with HUBO exponentially lowering qubit counts and reducing CNOT requirements by at least 89.6% for all benchmarked problem classes and sizes (see Figure 1). To support broad adoption, we present PyHUBO, an open-source Python library [31], and introduce a systematic framework for constructing HUBO Hamiltonians, showing that for the problems considered, our method yields numerical complexity that scales polynomially with problem size. Although we focus on QAOA, the approach applies to other algorithms and problems with one-hot or related constraints. Our results demonstrate that HUBO encodings can directly mitigate resource bottlenecks in quantum combinatorial optimization.

1 Methodology

The methodology of this work is structured as follows. In Sec. (1.1), we establish a formal definition and notation for COPs and illustrate this formulation using three representative examples: the GAP, the MkCS, and IP. In Sec. (1.2), we show how the COP formulations from Sec. (1.1) can be systematically mapped to QUBO and HUBO encodings. Finally, Sec. (1.3) describes the implementation of these encodings within the QAOA algorithm, utilizing only single- and two-qubit gates.

1.1 Combinatorial Optimization Problems

A COP consists of a finite set of feasible solutions and an objective function to be optimized [32]. For the remainder of this work, we use the following definition: Consider an n variable COP, where every variable can assume one of m values. Variables are indexed by $1 \leq i \leq n$ and values by $1 \leq v \leq m$. In the following, we will use the terms i th variable and variable i interchangeably to refer to the variable indexed by i , and similarly, we will use the v th value and value v interchangeably to refer to the value indexed by v . A solution $\mathbf{s} \in \{1, \dots, m\}^n$ assigns a value to each variable. For every $\mathbf{s}$ there exists an objective function $C(\mathbf{s})$ that can be computed in polynomial time and a polynomial time oracle that verifies whether $\mathbf{s}$ is a feasible solution.

The task is to find a solution such that the objective function is minimal among all feasible solutions, i.e., $\mathbf{s}^* = \operatorname{argmin}_{\mathbf{s}} C(\mathbf{s})$. In most applications, the objective function is a polynomial function of the variables [3], and for practical purposes, we express it as

$$C(\mathbf{s}) = \sum_{1 \leq i \leq n} c_1(i, s_i) + \sum_{1 \leq i < j \leq n} c_2(i, j, s_i, s_j) + \dots, \quad (1)$$

where $c_1(i, s_i)$ is the cost of assigning the value s_i to the i th variable, $c_2(i, j, s_i, s_j)$ the cost of assigning the value s_i to the i th variable while simultaneously assigning the value s_j to the j th variable.

In principle, the objective function can be of any order. However, for most industrial applications, it is linear or quadratic [33], i.e. Eq. (1) only contains terms of the form $c_1(i, s_i)$ or $c_2(i, j, s_i, s_j)$. Therefore, without loss of generality, we only consider linear and quadratic COPs.

Often, COPs also involve constraints. For practical purposes, we formulate these constraints as inequalities, which can be converted to equality constraints, e.g., using slack variables [34]. The equality constraints are added as penalty terms to Eq. (1) with appropriately chosen Lagrange multipliers [35].

For example, a COP with a constraint that all variable pairs $(i, j) \in P$ for some set P can not assume the same value can be stated as

$$\begin{aligned} & \underset{\mathbf{s} \in \{1, \dots, m\}^n}{\text{minimize}} && C(\mathbf{s}) \\ & \text{subject to} && s_i \neq s_j. \\ & && \forall (i, j) \in P \end{aligned} \quad (2)$$

The penalty term for this constraint has the form

$$\begin{aligned} c_2(i, j, v, v) &= \lambda \\ \forall (i, j) \in P \quad \forall v \in \{1, \dots, m\}, \end{aligned} \quad (3)$$

with the Lagrange multiplier λ chosen large enough such that the optimal solution respects the constraint. By adding this penalty term to the objective function, we can transform the COP into an unconstrained optimization problem

$$\underset{\mathbf{s} \in \{1, \dots, m\}^n}{\text{minimize}} \quad C(\mathbf{s}) + \lambda \sum_{(i, j) \in P} \delta_{s_i, s_j}, \quad (4)$$

where δ_{s_i, s_j} is the Kronecker delta function, which is 1 if $s_i = s_j$ and 0 otherwise.

The above notation applies broadly to COPs. We now illustrate this on three representative classes: the GAP, the MkCS problem, and IP. These formulations provide the basis for deriving their QUBO and HUBO encodings.

1.1.1 Gate Assignment Problem

The GAP is a well-studied challenge in operations research with significant practical relevance to the aviation industry [36]. It involves optimally assigning airport-gates to aircraft to minimize costs, such as passenger walking distances.

Various classical formulations of the GAP are established [36], and the problem has been studied in the context of quantum optimization [28, 37]. We consider the GAP with n flights and m airport-gates, where the objective is to minimize total passenger walking times (see Fig. 2).

We distinguish between two passenger types: those boarding or disembarking, and those transferring between flights. The first type only contributes to the linear coefficients, as the cost of




Figure 2: Illustration of the GAP with $m = 4$ airport-gates and $n = 3$ flights. Two of the flights have passengers who require a transfer connection, which is represented by the two-way symbol. A third flight has only arriving and departing passengers, illustrated by the luggage symbol. The goal is to assign airport-gates to flight such that the total passenger walking time is minimized.

assigning flight i to airport-gate v is given by

$$c_1(i, v) = p_i^{\text{arr}} t_v^{\text{arr}} + p_i^{\text{dep}} t_v^{\text{dep}}, \quad (5)$$

where p_i^{arr} and p_i^{dep} are the numbers of arriving and departing passengers on flight i , and t_v^{arr} and t_v^{dep} are the walking times from the check-in desk to the airport-gate v and from the airport-gate v to the luggage claim, respectively.

The transfer passengers contribute to the quadratic terms through

$$c_2(i, j, v, w) = p_{i,j}^{\text{trans}} t_{v,w}^{\text{trans}}, \quad (6)$$

where $p_{i,j}^{\text{trans}}$ is the number of passengers transferring from flight i to flight j , and $t_{v,w}^{\text{trans}}$ is the walking time between airport-gates v and w .

Finally, we enforce the constraint that temporally overlapping flights cannot share the same airport-gate through a penalty term of the same form as in Eq. (3), where P in the context of the GAP is the set of all conflicting flight pairs.

1.1.2 Maximum k-Colorable Subgraph Problem

The MkCS represents a fundamental COP that arises in diverse industrial applications, including network design, frequency assignment, and resource allocation [38–41]. It has attracted considerable attention in both classical and quantum optimization research [42–46].




Figure 3: Example instance for the MkCS problem: a properly 4-colored graph with $|V| = 7$. All adjacent vertices have distinct colors, so the objective value (number of monochromatic edges) is zero.

Given a graph $G = (V, E)$ with vertices V and edges E , the MkCS problem seeks to maximize the size of a properly vertex-colored subgraph with at most k colors (see Fig. 3). A properly vertex-colored subgraph is one where no two adjacent vertices share the same color, and the size of the subgraph is defined as the number of edges in the subgraph.

In the assignment formulation, each vertex is a variable that can assume k values, hence $n = |V|$ and $m = k$. The solution string $\mathbf{s}$ encodes the color of each vertex and s_i is the index for the color assigned to the vertex with index i . The objective function counts the number of edges in the graph that connect vertices of the same color. Minimizing the number of monochromatic edges is equivalent to maximizing the number of edges that connect vertices of different colors, which is the goal of the MkCS problem. Using the notation from Sec. (1.1), the objective function has the coefficients

$$c_2(i, j, v, v) = 1 \quad \forall (i, j) \in E, \quad (7)$$

and all other coefficients are zero.

1.1.3 Integer Programming

IP is fundamental to industrial applications such as production planning and resource allocation, as well as a key subject in classical optimization research [47–49]. Recently, quantum and hybrid

methods have been employed to address IP problems [20, 50, 51].

Here, we consider an IP problem with n variables and each variable can assume any value from the domain $\{y_1, \dots, y_m\}$ with $y_i \in \mathbb{Z}$. The task is to assign values to variables such that

$$\underset{\mathbf{u} \in \{y_1, \dots, y_m\}^n}{\text{minimize}} \quad \mathbf{q}^T \mathbf{u} + \mathbf{u}^T Q \mathbf{u} + \dots, \quad (8)$$

where $\mathbf{q} \in \mathbb{R}^n$ and $Q \in \mathbb{R}^{n \times n}$ are the coefficients of the objective function. In principle, the objective function Eq. (8) can be of any order. However, for the sake of clarity, here we write out only up to quadratic order. Typically, an IP problem also involves inequality constraints, which often are functions of many variables.

The objective function can be put into the assignment formulation of Sec. (1.1) by writing

$$\begin{aligned} c_1(i, v) &= q_i y_v, \\ c_2(i, j, v, w) &= Q_{ij} y_v y_w, \end{aligned} \quad (9)$$

for all $1 \leq v, w \leq m$. We find that computing the coefficients of the objective functions under consideration has computational complexity $\mathcal{O}(\binom{n}{2} m^2)$ (since in general there are $\binom{n}{2} m^2$ quadratic and $n \times m$ linear coefficients).

A range of methods exists to enforce the constraints of IP problems. Penalty-based techniques are most common, encoding constraints as additional terms in the objective function. Examples include slack variable methods that introduce additional variables to convert inequalities into penalizable equalities [34], as well as more recent variants such as unbalanced penalization [24], and Lagrangian approaches [52], both of which avoid increasing the variable space. Alternatively, oracle-based methods verify feasibility directly, without penalties [50]. Together, these strategies are broadly applicable to general constraint enforcement in IP.

In our benchmark examples, we restrict ourselves to constraints that each involve at most two variables. For such simple pairwise constraints, it is most straightforward to explicitly enumerate all violating value pairs, which can be done efficiently with computational complexity $\mathcal{O}(m^2)$ per constraint. Once identified, for each pair of values y_v, y_w violating a constraint with variables i and j the objective function coefficients are changed

$$c_2(i, j, v, w) \rightarrow c_2(i, j, v, w) + \lambda, \quad (10)$$

with a sufficiently large Lagrange multiplier λ . This direct penalty-based formulation is particularly efficient and effective for enforcing simple pairwise constraints in the framework considered here.

1.2 From Classical to Quantum Optimization

The task of mapping a COP to its quantum formulation is twofold. Firstly, we construct a unique mapping between classical solutions and quantum states $\mathbf{s} \leftrightarrow |\mathbf{x}\rangle$. Secondly, we construct a Hamiltonian $\hat{H}_C$ such that $\hat{H}_C |\mathbf{x}\rangle = C(\mathbf{x}) |\mathbf{x}\rangle$. We define $C(\mathbf{x}) \leftrightarrow C(\mathbf{s})$ in which $C(\mathbf{s})$ is given by Eq. (1). Finding the optimal solution to the COP corresponds to finding the ground state of the Hamiltonian. In the following, we discuss two different approaches to this mapping: the QUBO and the HUBO encoding.

1.2.1 QUBO Encoding

In the QUBO encoding, for each variable value pair (i, v) a binary variable $x_{i,v}$ is created. The i th variable taking the v th value, is encoded by setting the binary variable $x_{i,v} = 1$, while all other binary variables $x_{i,w}$ with $w \neq v$ are set to 0. The objective function Eq. (1) can be written in terms of binary variables $x_{i,v}$ as

$$C(\mathbf{x}) = \sum_{1 \leq i \leq n} \sum_{1 \leq v \leq m} c_1(i, v) x_{i,v} + \sum_{1 \leq i < j \leq n} \sum_{1 \leq v \leq w \leq m} c_2(i, j, v, w) x_{i,v} x_{j,w}. \quad (11)$$

To ensure that each variable i can only take one value, we must add a so-called one-hot penalty term

$$\lambda \left(1 - \sum_{v=1}^m x_{i,v} \right)^2, \quad (12)$$

to the objective function, which ensures that for each variable i exactly one of the binary variables $x_{i,v}$ is 1 and all others are 0. For constrained COPs, other penalty terms need to be added to the cost function to ensure that the constraints are satisfied. For example, the constraint introduced in Sec. (1.1), which requires that all variable pairs $(i, j) \in P$ for some set P cannot assume the same value, is enforced by the penalty term

$$\lambda \sum_{i,j \in P} \sum_{v=1}^m x_{i,v} x_{j,v}. \quad (13)$$

Similarly, other constraints can be added as penalty terms to the objective function. A discussion about the construction of the penalty terms in QUBO problems can be found in [35].

From the QUBO formulation, a quantum system is constructed where each QUBO variable is mapped to one qubit. Then, the value of the qubit $|j\rangle_{i,v}$ in the computational basis represents the value of the QUBO variable $x_{i,v}$. Hence, the total number of qubits scales with the number of QUBO variables $n \times m$. Now the Hamiltonian for this encoding can be found by substituting $x_{i,v} \rightarrow (\mathbb{I} - \hat{Z}_{i,v})/2$, where $\hat{Z}_{i,v}$ is the Pauli-Z operator acting on the qubit encoding the variable $x_{i,v}$. This gives the operator form of the Hamiltonian

$$\hat{H}_C = \sum_{1 \leq i \leq n} \sum_{1 \leq v \leq m} J_{i,v} \hat{Z}_{i,v} + \sum_{1 \leq i < j \leq n} \sum_{1 \leq v, w \leq m} J_{i,j,v,w} \hat{Z}_{i,v} \hat{Z}_{j,w}, \quad (14)$$

where the linear cost coefficients $c_1(i, v)$ contribute to the linear QUBO coefficients $J_{i,v}$ and the quadratic cost coefficients $c_2(i, j, v, w)$ contribute to $J_{i,v}$, $J_{j,w}$ and $J_{i,j,v,w}$. It is easy to check that with this Hamiltonian we have $\hat{H}_C |\mathbf{x}\rangle = C(\mathbf{x}) |\mathbf{x}\rangle$. In most numerical studies, the QUBO coefficients $J_{i,v}$ and $J_{i,j,v,w}$ are calculated from $c_1(i, v)$ and $c_2(i, j, v, w)$ through packages such as pyqubo [22, 53].

1.2.2 HUBO Encoding

In the HUBO encoding, each variable i is represented using $d = \lceil \log_2(m) \rceil$ bits. To retrieve the value assigned to variable i , we read its associated d bits, convert the binary string to the decimal value v , and conclude that variable i assumes the value v .

For a quantum system with $n \times d$ qubits encoding the COP solution, measuring the d -qubit state $|j\rangle_i$ in the computational basis yields a bit-string that, when converted to a decimal v , sets the variable with index i to the value with index v . While a similar encoding of integer values has been discussed in [23], here we use the encoded integers to denote value indices. This allows us to encode values beyond the set of integers, e.g., colors or gates. In this encoding, the objective

function can be written as

$$\begin{aligned}\hat{H}_C = & \sum_{1 \leq i \leq n} \sum_{1 \leq v \leq m} c_1(i, v) |v\rangle \langle v|_i \\ & + \sum_{1 \leq i < j \leq n} \sum_{1 \leq v, w \leq m} c_2(i, j, v, w) |v\rangle \langle v|_i \otimes |w\rangle \langle w|_j.\end{aligned}\quad (15)$$

It is straightforward to check that we have $\hat{H}_C |\mathbf{x}\rangle = C(\mathbf{x}) |\mathbf{x}\rangle$ with this encoding.

Furthermore, if m is not a power of two, we also add a penalty term

$$\lambda \sum_{i=1}^n \sum_{v=m}^{2^d} |v\rangle \langle v|_i \quad (16)$$

to penalize all states that encode indices larger than m . Additionally, to incorporate the penalty term (3), we can add a term of the form

$$\lambda \sum_{i,j \in P} \sum_{v=1}^m |v\rangle \langle v|_i \otimes |v\rangle \langle v|_j, \quad (17)$$

to the Hamiltonian. The addition of a one-hot constraint (Eq. (12)) is not required in this encoding and, as we will show, this substantially reduces the resource requirements.

For the construction of the quantum circuit, however, we need to encode the Hamiltonian in terms of Pauli matrices. This can be done by noting that each projector-term can be written out as

$$\begin{aligned}|v\rangle \langle v|_i &= |v_1\rangle \langle v_1|_{i,1} \cdots |v_d\rangle \langle v_d|_{i,d} \\ &= \frac{\mathbb{I}_{i,1} + (-1)^{v_1} \hat{Z}_{i,1}}{2} \cdots \frac{\mathbb{I}_{i,d} + (-1)^{v_d} \hat{Z}_{i,d}}{2} \\ &= \frac{1}{2^d} \sum_{S \subseteq \{1, \dots, d\}} \prod_{a \in S} (-1)^{v_a} \hat{Z}_{i,a},\end{aligned}\quad (18)$$

where v_a is the a th bit of the binary representation of v and $\hat{Z}_{i,a}$ is the Pauli-Z operator acting on the qubit encoding the a th bit of the i th variable. Note that the sum over S runs over all subsets of the set $B = \{1, \dots, d\}$, which includes the empty set. The empty set contributes a global factor of $\mathbb{I}/2^d$ to the sum.

To simplify the notation, we define the product of Pauli-Z operators for a variable i and a set $S \subseteq B$ as

$$\hat{Z}_{i,S} = \prod_{a \in S} \hat{Z}_{i,a}. \quad (19)$$

Additionally, for a set S and a value v , we define

$$R_S(v) = \frac{1}{2^d} \prod_{a \in S} (-1)^{v_a}. \quad (20)$$

This allows us to express the projector in a more compact form as

$$|v\rangle \langle v|_i = \sum_{S \subseteq B} R_S(v) \hat{Z}_{i,S}. \quad (21)$$

Substituting (21) into the Hamiltonian (15) yields

$$\begin{aligned}\hat{H}_C = & \sum_{1 \leq i \leq n} \sum_{S \subseteq B} J_{i,S} \hat{Z}_{i,S} \\ & + \sum_{1 \leq i < j \leq n} \sum_{S_1, S_2 \subseteq B, B} J_{i,j,S_1,S_2} \hat{Z}_{i,S_1} \hat{Z}_{j,S_2},\end{aligned}\quad (22)$$

where

$$\begin{aligned}J_{i,S} &= \sum_{v=1}^m R_S(v) c_1(i, v), \\ J_{i,j,S_1,S_2} &= \sum_{v,w=1}^m R_{S_1}(v) R_{S_2}(w) c_2(i, j, v, w).\end{aligned}\quad (23)$$

Eq. (22) is a HUBO Hamiltonian with linear, and quadratic terms and terms of order up to $2 \times d$ in the Pauli-Z operators. The coefficients $J_{i,S}$ are determined by $c_1(i, v)$ and the coefficients J_{i,j,S_1,S_2} by $c_2(i, j, v, w)$. Hence, Eq. (23) provides a systematic way to find the coefficients of the HUBO Hamiltonian from the coefficients of the objective function.

In principle, the HUBO-coefficients $J_{i,S}$ and J_{i,j,S_1,S_2} can be computed from (23). However, this quickly becomes cumbersome, especially for large COPs. Instead, we use a technique based on the Walsh-Hadamard transform [54], which we present in Appendix A. We also prove that calculating the HUBO-coefficients has a numerical complexity that scales polynomially with the number of variables n and values m . For example, calculating the HUBO-coefficients for a COP with a quadratic objective function has numerical complexity $\mathcal{O}(n^2 m^6)$. Our open-source repository offers the Python package PyHUBO [31], along with detailed documentation, requirements, and installation instructions, that facilitates the computation of HUBO coefficients for COPs.




Figure 4: Schematic representation of the QAOA algorithm showing the initialization in the superposition state, the alternating application of cost and mixer unitaries across layers, followed by measurement and classical optimization of variational parameters γ and β .

1.3 Circuit Implementation

In the previous sections, we presented two methods that encode a COP into a Hamiltonian, ensuring that the ground state of the Hamiltonian corresponds to the optimal solution. Various techniques exist to find this ground state, such as adiabatic evolution [16, 55], quantum imaginary time evolution [56–58], and variational methods [59, 60]. In this work, we focus on the QAOA [18], a variational approach well-suited for near-term quantum devices and effective across many COPs [61–66]. QAOA is widely used in the community and has been implemented on multiple quantum devices, making it an ideal benchmark for our encoding schemes [44, 67, 68].

QAOA operates by alternating between two types of unitary operations: the cost unitary $U_C(\gamma) = \exp(-i\gamma\hat{H}_C)$ with the cost Hamiltonian $\hat{H}_C$ as developed in the Sec. (1.2.1) and Sec. (1.2.2), and the mixer unitary $U_M(\beta) = \exp(-i\beta\hat{H}_M)$ (see figure 4). Below, we outline the circuit constructions for these unitaries. For a comprehensive overview of QAOA, see [63].

The construction of the cost unitary $U_C(\gamma)$ for Hamiltonians with higher-order terms using only one- and two-qubit gates requires careful circuit implementation to maintain efficiency. A key observation that simplifies the circuit construction is that the cost Hamiltonian $\hat{H}_C$ for every encoding, when expressed in the computational basis, is a sum of diagonal terms. This diagonal structure implies that any two summands in the cost Hamiltonian commute with each other. This

commutativity property allows us to decompose the cost unitary for both QUBO and HUBO encodings into a product of unitaries that each implement single terms $J_T \hat{Z}_T$ of the Hamiltonians (14) and (22). Here, T is a set of qubit indices, J_T is the coefficient for these qubits and, similar to Sec. 1.2.2 we use the notation $\hat{Z}_T = \prod_{t \in T} \hat{Z}_t$. Obviously, in the QUBO encoding, the set T only involves at most two qubit indices, while it can involve up to $2 \times d$ qubit indices in the HUBO encoding.

To construct a circuit for the unitary $\exp(i\gamma J_T \hat{Z}_T)$ using only single- and two-qubit gates, we exploit the fact that the computational basis states $|\mathbf{b}\rangle = \bigotimes_{t \in T} |b_t\rangle_t$, with $b_t \in \{0, 1\}$, form a complete basis. The action of $\exp(i\gamma J_T \hat{Z}_T)$ on such a state results in $\exp((-1)^{\bigoplus_{t \in T} b_t} i\gamma J_T) |\mathbf{b}\rangle$. We therefore design a circuit composed of single- and two-qubit gates that reproduces this transformation on computational basis states and thus implements the unitary $\exp(i\gamma J_T \hat{Z}_T)$.

The circuit is constructed in three main steps, as shown in Fig. 5. First, a sequence of $|T| - 1$ CNOT gates is used to transform the initial state $|\mathbf{b}\rangle$ such that one qubit stores the parity $\bigoplus_{t \in T} b_t$ of the bitstring. Specifically, by arranging the $|T| - 1$ CNOT gates in a chain, where the t -th CNOT acts on the t -th (control) and $(t + 1)$ -th (target) qubits of T , the parity is encoded in the last qubit of the set T . Second, a single-qubit rotation $R_z(2\gamma J_T)$ is applied to the parity qubit, generating the phase factor $\exp(-iJ_T\gamma)$ or $\exp(iJ_T\gamma)$ for odd or even parity, respectively. Finally, the sequence of CNOT gates is applied in reverse order to restore the original product state, resulting in $\exp((-1)^{\bigoplus_{t \in T} b_t} i\gamma J_T) |\mathbf{b}\rangle$. Thus, the circuit implementing the cost unitary for a single $|T|$ -th order term in the cost Hamiltonian requires $2(|T| - 1)$ CNOT gates and one R_Z gate.

In principle, each term in (14) and (22) can be implemented in this way sequentially. However, as we show in the Appendix B, constructing the cost layer $U_C(\gamma)$ for all terms of a HUBO Hamiltonian (22) in this way is not resource efficient, and we present an optimized circuit which significantly reduces the number of CNOT gates. The key idea is that parity states generated for one term of the Hamiltonian can often be leveraged to efficiently construct the parity states needed for




Figure 5: Circuit implementation of the cost Hamiltonian for a three-qubit term, i.e. $\exp(i\gamma J_{1,2,3} \hat{Z}_1 \hat{Z}_2 \hat{Z}_3)$. The circuit consists of a sequence of CNOT gates to entangle the qubits, such that the initial state $|b_1\rangle_1 |b_2\rangle_2 |b_3\rangle_3$ in the computational basis becomes $|b_1\rangle_1 |b_1 \oplus b_2\rangle_2 |b_1 \oplus b_2 \oplus b_3\rangle_3$ through the first chain of CNOT gates. The single qubit rotation $R_z(2\gamma J_{1,2,3})$ acts on the last qubit. After the rotation, another sequence of CNOT gates disentangles the qubits, resulting in the final state $\exp((-1)^{b_1 \oplus b_2 \oplus b_3} i\gamma J_{1,2,3}) |b_1\rangle_1 |b_2\rangle_2 |b_3\rangle_3$.

other terms, thereby minimizing redundant operations. Consequently, as we show in Appendix C, the asymptotic CNOT and R_Z gate counts scale quadratically in the number of variables n and values m .

For the mixer Hamiltonian $\hat{H}_M$, we employ the standard Pauli-X mixer

$$\hat{H}_M = \sum_i \hat{X}_i, \quad (24)$$

which acts on all qubits of the circuit. This choice ensures that the mixer unitary can efficiently explore the computational basis states by creating a superposition that allows transitions between different solution candidates. Furthermore, the mixer unitary can be implemented using a simple layer of single-qubit R_x rotation gates.

To optimize the variational parameters γ and β , we apply gradient descent methods with the initialization scheme of Zhou et al. [61], which improves convergence and solution quality over random initialization. Gradients of the expectation value with respect to variational parameters are efficiently computed via implicit differentiation [69], avoiding the computational overhead and inaccuracy of finite-difference methods. Quantum circuits and optimization routines are implemented using PennyLane [70], enabling automatic differentiation and GPU acceleration for efficient optimization.

The Lagrange multipliers λ for the different

penalty terms in the objective function are chosen iteratively before running the QAOA benchmarking. Only after identifying λ values that ensure the QAOA algorithm consistently finds feasible solutions respecting all constraints do we proceed to run and evaluate the algorithm's performance using these selected penalty parameters. This iterative approach ensures that the penalty terms are effective in guiding the QAOA algorithm towards feasible solutions while still allowing for exploration of the solution space.

The quality of QAOA solutions is typically evaluated via the expectation value of the cost Hamiltonian. For COPs, however, feasibility must also be considered. Hence, we employ a metric similar to the one introduced in [71], here referred to as the approximation ratio, defined as

$$A = 1 - \sum_{\mathbf{b} \in B_{\text{feas}}} r(\mathbf{b}) |\langle \psi(\gamma, \beta) | \mathbf{b} \rangle|^2, \quad (25)$$

where $\{\mathbf{b} | \mathbf{b} \in B_{\text{feas}}\}$ is the set of computational basis states satisfying all constraints and

$$r(\mathbf{b}) = \frac{C_{\text{max}} - C(\mathbf{b})}{C_{\text{max}} - C_{\text{min}}}, \quad (26)$$

rescales $C(\mathbf{b})$ to the interval $[0, 1]$. Here, C_{min} and C_{max} are the minimal and maximal values of the objective function, determined using commercial solvers for the benchmark instances. The approximation ratio is particularly suitable for our benchmarks, as it immediately accounts for the treatment of invalid solutions as equivalent to maximum cost states. This metric combines solution utility, feasibility, and objective value, and can be applied to all three COPs discussed in the results section.

If the quantum state $|\psi(\gamma, \beta)\rangle$ corresponds to the optimal solution, then $A = 0$. If it yields only infeasible or maximal-cost solutions, $A = 1$. Thus, A can be interpreted as the expectation value of a normalized cost Hamiltonian, where feasible eigenstates have eigenvalues $[0, 1]$, while infeasible eigenstates can only have eigenvalue 1. Lower values of A indicate higher-quality solutions found by QAOA.

Since the set of feasible bitstrings S_{feas} grows exponentially with problem size, evaluation of the approximation ratio becomes numerically taxing even for modestly sized COPs. To address this, we approximate Eq. (25) using a sampling-based method. Further details on this approach and

	# Qubits	# CNOTs	# R_Z
QUBO	nm	$O(n^2m^2)$	$O(n^2m^2)$
HUBO	$n\lceil\log_2(m)\rceil$	$O(n^2m^2)$	$O(n^2m^2)$

Table 1: Exact qubit requirements and analytical scaling for CNOT and R_Z gates per QAOA layer in both encodings, where n is the number of variables and m the number of values per variable. In practice, gate counts are often lower (see Table 2).

	# Qubits	# CNOTS	# R_Z
GAP-QUBO	20	140	90
GAP-HUBO	10	68	27
MkCS-QUBO	20	132	86
MkCS-HUBO	10	90	27
IP-QUBO	16	120	76
IP-HUBO	8	38	31

Table 2: Quantum Resources required per QAOA layer to encode the GAP instance introduced in Fig. 6, an MkCS instance illustrated in Fig. 9 and an IP instance shown in Fig. 12.

other numerical techniques used for benchmarking are provided in Appendix D.

2 Results

To evaluate resource efficiency and solution quality, we analyze three complementary metrics across all benchmarks: (i) the number of qubits required to represent the problem, (ii) the gate count per QAOA layer after circuit compilation (i.e. # CNOTS, # R_Z), and (iii) the total number of gates needed to achieve target approximation ratios (A).

For qubit requirements, a COP with n variables each taking m values requires $n \times m$ qubits under QUBO, but only $n \times \lceil\log_2(m)\rceil$ under HUBO. Both encodings grow linearly in n , however, the growth in m is linear for QUBO but logarithmic for HUBO, yielding exponential savings when variables admit many values. This advantage is particularly critical in problems such as the GAP, where each flight-variable can assume a large number of airport-gate-values [36].

For gate counts per QAOA layer, we compile the cost unitary circuits to contain only rotation and CNOT gates, and report those as our primary resources. Both encodings scale quadratically in gate count (see Table 1), but HUBO

consistently compiles to fewer gates in practice (Table 2), often by large margins, for instance, a 69% reduction in CNOT gates for the considered IP instance. It should be noted that Hadamard gates for state initialization and R_X gates for the mixer are not reported separately since their counts equal the number of qubits.

Finally, while increasing the number of QAOA layers is known to improve solution quality [18], feasibility on real hardware depends on total gate counts rather than per-layer counts alone. In what follows, we will show that both the per-layer gate requirements and the convergence rate toward the optimum vary significantly with the chosen encoding. Since no analytical bounds exist for the minimum depth needed to reach a given solution threshold, we rely on empirical benchmarking. Here, the most notable advantage of HUBO over QUBO is evident, with HUBO achieving savings of up to 97.1% in CNOT counts for IP benchmarks. In the following subsections, we therefore compare QUBO and HUBO across three representative COPs, reporting solution quality at fixed depths, convergence behavior as a function of gate count, and total resources needed to achieve specified solution thresholds. All averages are over 100 independently optimized QAOA parameters, with error bars representing the standard deviation, as detailed in Appendix D.

2.1 Gate Assignment Problem

Using a GAP instance with five flights ($n = 5$) and four airport-gates ($m = 4$), we benchmark both QUBO-QAOA and HUBO-QAOA, with ten layers. Each layer contains gate counts as reported in Table 2. The GAP instance considered has 4 flights with transfer passengers, resulting in 24 nonzero quadratic terms (see Equation (6)) in the classical objective function. Additionally, 3 flights have gate conflicts with 2 other flights each, while 2 flights have gate conflicts with 1 other flight each. The motivation for this scenario, as well as further details such as passenger numbers and walking times, are provided in Appendix E. An illustration of the problem is shown in Figure 6.

The best feasible solution, determined by a commercial solver, yields an average passenger travel time of 9.6 minutes, while the worst feasible assignment totals 17.9 minutes. Within this range of solutions, HUBO-QAOA achieves




Figure 6: Upper panel: Airport-Gate layout for the benchmark instance with 4 airport-gates and 5 flights. Lower panel: The range of feasible objective values for total passenger time is shown, including the minimum (blue) and maximum (yellow) obtained by a commercial solver. The average objective values achieved by the QAOA algorithm with QUBO and HUBO encodings (10 layers) are shown for comparison.

an average objective value of 12.1 minutes, while QUBO-QAOA’s average objective value is 15.8 minutes, demonstrating a substantial solution quality improvement with HUBO for the same number of QAOA layers.

Furthermore, HUBO-QAOA achieves lower approximation ratios using fewer quantum resources. In Fig. 7, we compare the average approximation ratio as a function of total gate count, varying the number of QAOA layers from 1 to 10. Because the same number of gates is added per QAOA layer, the datapoints are evenly spaced along the x-axis. For any fixed number of QAOA layers, the HUBO encoding requires fewer gates in total. More importantly, for any given total gate count, HUBO-QAOA consistently attains much better approximation ratios. For example, with approximately 1040 total gates (1060 for HUBO, 1020 for QUBO), the approximation ratio achieved with QUBO is more than 3 times worse than that achieved with HUBO.

HUBO-QAOA consistently demonstrates resource savings over QUBO-QAOA as the problem size varies. Fig. 8 summarizes the number of gates required to achieve an approximation ratio of at least 0.50 for both encodings, as the number of flights increases from 1 to 5 with the number of airport-gates fixed at 4. The smaller




Figure 7: Average approximation ratio (Eq. (25)) versus total gate count for HUBO and QUBO encodings on the 4-airport-gate, 5-flight problem instance. The gate count includes all circuit gates used: CNOT, R_Z , R_X , and Hadamard, the latter applied only once for state initialization (10 for HUBO, 20 for QUBO). For each encoding, the number of gates per QAOA layer is fixed, so increasing the number of layers results in a uniform increment in total gates, producing evenly spaced datapoints along the x-axis. Each point corresponds to a fixed number of QAOA layers, ranging from 1 to 10.




Figure 8: Total number of CNOT and single-qubit gates required to attain an approximation ratio of 0.50, plotted against the problem size of the GAP (number of airport-gates fixed at 4). Problem size is defined as the product of the number of flights and the number of gates and directly corresponds to the qubit count in the QUBO encoding. CNOT and single-qubit gates are shown separately, reflecting their differing experimental costs. Across all problem sizes, the HUBO encoding achieves the target approximation ratio using substantially fewer CNOT and single-qubit gates than the QUBO encoding.




Figure 9: Upper panel: Graph with five vertices and nine edges used as a benchmark instance for the MkCS problem, colored with four colors. The displayed coloring corresponds to the optimal solution with zero monochromatic edges. Lower panel: Range of feasible objective values and average number of monochromatic edges identified by HUBO-QAOA and QUBO-QAOA for a fixed number of layers (3).

GAP instances are generated by repeatedly deleting flight assignments from the instance described in Fig. 6. In this scaling approach, the problem size corresponds directly to the number of qubits required by the QUBO encoding. The figure separately reports the number of CNOT and single-qubit gates, since CNOT gates are typically more experimentally demanding. Across all evaluated problem sizes, the HUBO encoding reduces the number of required CNOT gates between 90.0% and 100% (the single flight GAP instance in the HUBO encoding doesn't require CNOT gates) and the number of R_Z gates between 86.1% and 95.3% compared to the QUBO encoding. These results highlight that substantial quantum resource efficiency gains with the HUBO approach persist across varying problem sizes.

2.2 Maximum k-Colorable Subgraph Problem

We benchmark an MkCS problem with $|V| = n = 5$ vertices and $k = m = 4$ colors on the graph illustrated in Fig. 9, using a QAOA circuit with three layers. The objective is to minimize the number of monochromatic edges. For this instance, an optimal coloring results in zero monochromatic edges, while assigning the same




Figure 10: Average approximation ratio achieved by QAOA as a function of total gate count for HUBO and QUBO encodings. Results are shown for the same graph with five vertices and nine edges as in Fig. 9. The number of gates required for each layer can be found in Table 2. Data points are evenly spaced as gate counts scale proportionally with the number of QAOA layers.

color to all vertices yields a maximum objective value of nine. The range of feasible values and the average performance of each encoding are shown in the lower panel of Fig. 9. HUBO-QAOA finds, on average, 1.2 monochromatic edges, compared to an average of 6.9 for QUBO-QAOA.

HUBO-QAOA consistently achieves better approximation ratios than QUBO-QAOA while using fewer quantum resources. Fig. 10 shows the average approximation ratio obtained by QAOA as a function of total gate count for both encodings. Using the gate numbers reported in Table 2, HUBO-QAOA always reaches approximation ratios below 0.18. In contrast, even with the largest considered number of QAOA layers (10), the QUBO-QAOA algorithm only attains an average approximation ratio of 0.37. Each QUBO-QAOA layer requires 238 gates, compared to only 127 gates for HUBO-QAOA. This substantial difference in resource demands can be attributed to the penalty overhead required by the QUBO encoding.

We find that the quantum resource advantage of HUBO over QUBO persists across all tested MkCS problem sizes. As summarized in Fig. 11, the number of gates required to achieve an approximation ratio of 0.20 increases much more rapidly for QUBO than for HUBO as problem size grows. Problem size is defined as the product of the number of vertices and the number of colors in the graph, which matches the qubit requirement for the QUBO encoding. For these bench-




Figure 11: Scaling of quantum resource requirements for MkCS instances. The number of CNOT and R_Z gates required to achieve an approximation ratio of 0.20 is shown as a function of problem size (number of vertices times number of colors, with $m = 4$ fixed). The HUBO encoding consistently requires substantially fewer gates than QUBO across all tested sizes, with the difference increasing for larger instances.

marks, the number of colors is fixed at $k = 4$ and the number of vertices varies from 1 to 5. The smaller graphs are generated by deleting vertices from the instance in Fig. 9. CNOT and single-qubit gate requirements are reported separately (note that the one-vertex instance is solved trivially by HUBO and doesn't require any CNOT or R_Z gates). Across all evaluated problem sizes, the HUBO encoding reduces the number of required CNOT gates by 89.6 – 100% and the number of R_Z gates by 90.8 – 100% compared to the QUBO encoding.

2.3 Integer Programming

To provide context for the subsequent analyses, the upper panel of Fig. 12 illustrates the IP instance considered in this work. This instance consists of $n = 5$ variables, each with $m = 4$ possible values, two linear constraints (each depending on two variables), and an objective function containing both linear and quadratic terms. The minimum feasible objective value is -27 , which is achieved by the assignment $x_1 = 0$, $x_2 = 3$, $x_3 = 3$, $x_4 = 0$, and $x_5 = 3$. This instance is used to benchmark HUBO-QAOA; however, QUBO-QAOA was not computationally feasible at this problem size due to excessive computational resource requirements.

To enable a direct comparison between the two encodings, we benchmark QAOA on a reduced version of the problem with $x_5 = 3$ fixed. The




Figure 12: Upper panel: Structure of the IP instance considered, with $n = 5$ variables ($m = 4$ values each), two linear constraints, and both linear and quadratic objective function terms. This instance is referenced here for later HUBO-QAOA benchmarking, as QUBO-QAOA is not computationally feasible for this size. Lower panel: Range of feasible objective values for the reduced IP instance with $x_5 = 3$ fixed. The optimal (-27) and maximal (18) values are shown. Average objective values attained by HUBO-QAOA and QUBO-QAOA after nine layers are indicated for comparison.

lower panel of Fig. 12 displays the range of feasible objective values for this simplified instance, which retains the optimum (-27) and maximum (18) values of the original. For 9 QAOA layers, HUBO-QAOA yields an average objective value of -18.55 , while QUBO-QAOA produces an average of 0.21 .

Additionally, HUBO-QAOA demonstrates significantly greater resource efficiency and superior approximation ratio compared to QUBO-QAOA across all layer settings, as shown in Fig. 13. The per-layer CNOT and R_Z gate counts for each encoding are provided in Table 2. Evidently, QUBO-QAOA requires over three times more CNOT gates and more than twice as many R_Z gates per layer compared to HUBO-QAOA.

For a targeted approximation threshold, HUBO-QAOA requires substantially fewer gates compared to QUBO-QAOA. For instance, to reach an approximation ratio of 0.60, QUBO-QAOA requires 10 layers, while HUBO-QAOA achieves the same result with only 1 layer. When accounting for both the higher resources per layer and the greater number of layers required, QUBO-QAOA, compared to HUBO-QAOA, consumes 31.6 times more CNOT gates (1200 vs. 38) and 24.5 times more R_Z gates (760 vs. 31).




Figure 13: Average approximation ratio versus total gate count for QAOA applied to the reduced IP instance, comparing HUBO and QUBO encodings. HUBO-QAOA consistently achieves higher approximation ratios at lower quantum resource cost than QUBO-QAOA. For example, to reach a target approximation ratio of 0.60, HUBO-QAOA and QUBO-QAOA require 85 and 2136 gates, respectively.




Figure 14: Number of gates required to achieve an approximation ratio of 0.30 as a function of IP problem size, comparing HUBO and QUBO encodings. Both CNOT and single-qubit gate requirements are shown. Problem size is modified by changing the number of variables in the IP problem. HUBO-QAOA consistently achieves the target approximation ratio with dramatically fewer gates than QUBO-QAOA, and remains tractable even for the largest instances where QUBO-QAOA cannot be executed due to computational resource constraints.

The resource advantage of HUBO-QAOA over QUBO-QAOA remains robust across all benchmarked IP problem sizes. Fig. 14 reports the number of gates required to reach an approximation ratio of 0.30 for different problem sizes, with CNOT and single-qubit gate requirements shown separately. Problem size is varied by starting from the largest instance (as shown in Fig. 12, with $n = 5$ variables, and $m = 4$ values) and generating progressively smaller instances by fixing variables to their optimal assignments. This approach preserves the minimal objective value while producing a sequence of simpler problems, resulting in increments of 4 qubits per data point.

Across all benchmarked problem sizes, the HUBO encoding reduces CNOT requirements by 94.4 – 100% and the number of R_Z gates by 91.6 – 95.3% compared to QUBO (again, the one variable instance doesn’t require any CNOT gates in the HUBO formulation). These results, together with those from earlier figures, demonstrate the persistent resource efficiency advantage of HUBO-QAOA over QUBO-QAOA throughout the tested range of IP problem sizes.

3 Conclusion

We showed that using HUBO formulations circumvents resource bottlenecks that often constrain the practical quantum optimization of COPs. Although one might assume that higher-order interactions would increase quantum hardware demands, our results on three representative problems reveal the opposite. HUBO encoding requires exponentially fewer qubits compared to QUBO, and more importantly, leads to substantial reductions in the number of CNOT and single-qubit gates. For each problem class, we observe a reduction of CNOT gates by 89.6 – 100%, which makes experimental realization more feasible for current and near-term quantum devices.

In addition to these resource savings, we have developed a general framework for constructing HUBO models, together with the open-source Python package [31]. This equips the community with a framework to efficiently formulate COPs on quantum hardware, helping to mitigate current resource limitations, and, for quadratic problems, ensures that the Hamiltonian construction scales polynomially with problem size.

By helping to lower the resource bottleneck,

our approach brings quantum optimization of challenging, real-world problems closer to practical reality. As quantum technology continues to advance, resource-efficient methods such as ours will play a key role in enabling broader applications in industry and research.

4 Acknowledgements

The authors gratefully acknowledge the funding provided by the Hamburgische Investitions- und Förderbank for the project "Efficient Quantum Algorithms for Aviation". The authors also thank Felix Herbot for his technical support and Robert Glöckner and Shaham Jafarpisheh for helpful discussions. SP and DJ acknowledge support from the Hamburg Quantum Computing Initiative (HQIC) project EFRE. The EFRE project is co-financed by ERDF of the European Union and by "Fonds of the Hamburg Ministry of Science,

Research, Equalities and Districts (BWFGB)". DJ acknowledges support from the Cluster of Excellence 'Advanced Imaging of Matter' of the Deutsche Forschungsgemeinschaft (DFG) - EXC 2056 - project ID 390715994, the European Union's Horizon Programme (HORIZON-CL42021- DIGITALEMERGING-02-10) Grant Agreement 101080085 QCFD, and DFG project 'Quantencomputing mit neutralen Atomen' (JA 1793/1-1, Japan-JST-DFG-ASPIRE 2024). For this work the HPC-cluster Hummel-2 at University of Hamburg was used. The cluster was funded by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - 498394658. Für diese Arbeit wurde das HPC-Cluster Hummel-2 an der Universität Hamburg genutzt, das durch die Deutsche Forschungsgemeinschaft (DFG) - 498394658 - gefördert wurde. We acknowledge funding from the Federal Ministry of Research, Technology and Space (BMFTR) under the grant BeRyQC.

A Mathematical Foundation of HUBO Encoding

While in principle it is possible to find the HUBO coefficients using Eq. (23), it quickly becomes cumbersome and numerically expensive to iterate over all required subsets. Instead, we adopt a more efficient, less error-prone, and clearer method to determine the HUBO coefficients. The following derivation is inspired by [54], where the authors show that finding the coefficients of a HUBO Hamiltonian from a Boolean function amounts to Fourier transforming it.

We first note that a single qubit projector can be written out as

$$\begin{pmatrix} c_0 \\ c_1 \end{pmatrix} \begin{pmatrix} |0\rangle \langle 0| \\ |1\rangle \langle 1| \end{pmatrix} = \frac{1}{2} \begin{pmatrix} c_0 + c_1 \\ c_0 - c_1 \end{pmatrix} \begin{pmatrix} \mathbb{I} \\ \hat{Z} \end{pmatrix}. \quad (27)$$

In the following, we will call $\begin{pmatrix} |0\rangle \langle 0| \\ |1\rangle \langle 1| \end{pmatrix}$, p -basis and $\begin{pmatrix} \mathbb{I} \\ \hat{Z} \end{pmatrix}$, j -basis and rewrite Eq. (27) as

$$\begin{pmatrix} c_0 \\ c_1 \end{pmatrix}_p = \frac{1}{2} \begin{pmatrix} c_0 + c_1 \\ c_0 - c_1 \end{pmatrix}_j = \hat{H} \begin{pmatrix} c_0 \\ c_1 \end{pmatrix}_j := \begin{pmatrix} J_0 \\ J_1 \end{pmatrix}_j, \quad (28)$$

where $\hat{H}$ is the Hadamard matrix providing a transformation from the p -basis to the j -basis.

Generalizing this idea to an d -qubit projector, the coefficient $c_{i_1 i_2 \dots i_d}$ of the projector $|i_1 i_2 \dots i_d\rangle \langle i_1 i_2 \dots i_d|$ in the p -basis is transformed to the j -basis using the Hadamard matrix $\hat{H}^{\otimes d}$, which is the tensor product of d Hadamard matrices. This can be seen by noting that the projector $|i_1 i_2 \dots i_d\rangle \langle i_1 i_2 \dots i_d|$ can be written as the tensor product of the single-qubit projectors $|i_1\rangle \langle i_1| \otimes |i_2\rangle \langle i_2| \otimes \dots \otimes |i_d\rangle \langle i_d|$, each being transformed into the j -basis by a Hadamard matrix, hence the transformation for the entire d -qubit projector is done by the tensor product of Hadamard matrices for all single qubit coefficients.

We use the notation $c_{i_1 i_2 \dots i_d}$ for the coefficient of the projector $|i_1 i_2 \dots i_d\rangle \langle i_1 i_2 \dots i_d|$ in the p -basis and the coefficient $J_{i_1 i_2 \dots i_d}$ in the j -basis encodes the interaction strength of the HUBO term




Figure 15: Circuit implementation of the cost unitary for two-qubit and three-qubit terms, i.e. $\exp(i\gamma J_{1,2} \hat{Z}_1 \hat{Z}_2)$ and $\exp(i\gamma J_{1,2,3} \hat{Z}_1 \hat{Z}_2 \hat{Z}_3)$, respectively. The circuit consists of a sequence of CNOT gates to entangle the qubits, such that the initial states $|b_1\rangle_1 |b_2\rangle_2$ and $|b_1\rangle_1 |b_2\rangle_2 |b_3\rangle_3$ in the computational basis become $|b_1\rangle_1 |b_1 \oplus b_2\rangle_2$ and $|b_1\rangle_1 |b_1 \oplus b_2\rangle_2 |b_1 \oplus b_2 \oplus b_3\rangle_3$ through the first chain of CNOT gates. The single qubit rotation $R_Z(2\gamma J_{0,1})$ and $R_Z(2\gamma J_{0,1,2})$ acts on the last qubit. After the rotation, another sequence of CNOT gates disentangles the qubits, resulting in the final state $\exp((-1)^{b_1 \oplus b_2} i\gamma J_{1,2}) |b_1\rangle_1 |b_2\rangle_2$ and $\exp((-1)^{b_1 \oplus b_2 \oplus b_3} i\gamma J_{1,2,3}) |b_1\rangle_1 |b_2\rangle_2 |b_3\rangle_3$.

$\prod_{j=1}^d \delta_{i,j,1} \hat{Z}_j$. Using this notation, we can write the transformation from the p -basis to the j -basis as

$$\begin{pmatrix} c_{0\dots 00} \\ c_{0\dots 01} \\ \vdots \\ c_{1\dots 10} \\ c_{1\dots 11} \end{pmatrix}_p = \hat{H}^{\otimes d} \begin{pmatrix} c_{0\dots 00} \\ c_{0\dots 01} \\ \vdots \\ c_{1\dots 10} \\ c_{1\dots 11} \end{pmatrix}_j = \begin{pmatrix} J_{0\dots 00} \\ J_{0\dots 01} \\ \vdots \\ J_{1\dots 10} \\ J_{1\dots 11} \end{pmatrix}_j. \quad (29)$$

For example, a quadratic COP has the HUBO model

$$\hat{H} = \sum_{1 \leq i \leq n} \sum_{1 \leq k \leq m} c_1(i, k) |k\rangle \langle k|_i + \sum_{1 \leq i < j \leq n} \sum_{1 \leq k, l \leq m} c_2(i, j, k, l) |k\rangle \langle k|_i \otimes |l\rangle \langle l|_j. \quad (30)$$

We write the linear coefficients for each variable i in a vector $\mathbf{c}(i) = (c_1(i, 1), c_1(i, 2), \dots, c_1(i, m))^T$ and apply $\hat{H}^{\otimes d}$ to each vector $\mathbf{c}_1(i)$ to obtain the coefficients in the j -basis $\mathbf{J}_i = \hat{H}^{\otimes d} \mathbf{c}_1(i)$. Note that because $2^d = m$, the Hadamard transformation is an $m \times m$ matrix, hence finding the linear coefficients of all n variables in the j -basis has a computational complexity of n matrix multiplications on an m -dimensional vector, hence $O(nm^3)$. Similarly, we can find the quadratic coefficients by applying the Hadamard transform to the vector $\mathbf{c}_{ij} = (c_2(i, j, 1, 1), c_2(i, j, 1, 2), \dots, c_2(i, j, m, m))^T$ for each pair of variables (i, j) and obtain the linear coefficients in the j -basis $\mathbf{J}_{ij} = \hat{H}^{\otimes d} \mathbf{c}_{ij}$. Calculating the quadratic coefficients for all pairs of variables amounts to a computational complexity of n^2 (since there are $\binom{n}{2}$ combinations of variables) and with the vector $\mathbf{c}_{i,j}$ being of size m^2 the computational complexity for finding the quadratic coefficients is $O(n^2 m^6)$. In general, for a COP of k th order with n variables and m values per variable, the computational complexity for finding the coefficients in the j -basis is $O(\binom{n}{k} m^{3k})$.

This formalism forms the basis of the PyHUBO Python package [31], which is made available as part of this publication.

B Cost-Efficient Circuit Implementation of HUBO Hamiltonians

An optimized circuit implementation of a HUBO Hamiltonian can be achieved using a Gray code approach [72], which has been successfully employed to problems like the traveling salesperson problem [26] and the circuit representation of boolean functions [73].

To understand this approach, we examine the effect of the circuit introduced in Sec. (1.3) on a computational basis state $|\mathbf{b}\rangle = \otimes_{t \in T} |b_t\rangle_t$ for a set of qubit indices T . To construct a circuit for

the term $\exp(iJ_T\hat{Z}_T)$, we create entangled states where one qubit encodes the parity $\oplus_{t \in T} b_t$. Acting on an R_Z gate on this qubit results in a phase factor $\exp(i\gamma J_T)$ if $\oplus_{t \in T} b_t = 1$ and $\exp(-i\gamma J_T)$ if $\oplus_{t \in T} b_t = 0$. This is exactly how $\exp(iJ_T\hat{Z}_T)$ acts on a computational basis state. We illustrate the circuit in Fig. 15, where we show the circuit diagrams to implement the terms $\exp(i\gamma J_{1,2}\hat{Z}_1\hat{Z}_2)$ and $\exp(i\gamma J_{1,2,3}\hat{Z}_1\hat{Z}_2\hat{Z}_3)$.

To implement all terms of the HUBO-Hamiltonian $\sum_{T \subseteq \{1, \dots, d\}} J_T \hat{Z}_T$, we construct a circuit that creates parity states $|\oplus_{t \in T} b_t\rangle$ and acts R_Z gates on these states for all subsets $T \subseteq \{1, \dots, d\}$ at least once. This can, in principle, be done by using circuits like in Fig. 15 sequentially. However, this is suboptimal in the number of CNOT gates. The construction of parity states $|\oplus_{t \in T} b_t\rangle$ for all subsets $T \subseteq \{1, \dots, d\}$ with minimal CNOT gate costs is a routing problem, and as shown in [72] and [74] the Gray Code implementation is the optimal approach to it.

The key observation for optimizing the circuit is that once the parity state corresponding to a subset $T_1 \subseteq \{1, \dots, d\}$ is prepared, the parity state for any subset $T_2 \subseteq \{1, \dots, d\}$ differing from T_1 by exactly one element can be obtained by a single additional operation. In particular, the phase factor $\exp(iJ_{T_2}\hat{Z}_{T_2})$ is implemented by applying a single CNOT gate, with the differing element as control and the state encoding $\bigoplus_{a \in T_1} b_a$ as target.

To introduce the Gray Code-based circuit, we denote each subset $T_1 \subseteq \{1, \dots, d\}$ through a binary string $\mathbf{g}$ of length d , where the j -th element is 1 if $j \in T_1$ and 0 otherwise. If two subsets T_1, T_2 differ in exactly one element, the corresponding binary strings $\mathbf{g}_1, \mathbf{g}_2$ differ in exactly one bit.

Binary strings of length d map naturally to d -dimensional hypercubes, and implementing all subsets of $\{1, \dots, d\}$ is equivalent to traversing all vertices of this hypercube. Each vertex of the hypercube corresponds to one binary string and therefore to one subset $T \subseteq \{1, \dots, d\}$. Two neighboring vertices differ by one element. The construction of parity states $|\oplus_{i \in T} b_i\rangle$ for all subsets $T \subseteq \{1, \dots, d\}$ with minimal CNOT gate costs, now amounts to traversing the hypercube with as few steps as possible.

The Gray code provides a systematic way to traverse the vertices of this hypercube. Crucially, the Gray code visits each vertex of the hypercube exactly once, thereby avoiding any redundant operations or inefficiencies that would arise from revisiting previously implemented terms. Starting from the vertex corresponding to the global phase, i.e., the bitstring $\mathbf{0} = (0, 0, \dots, 0)$ and following the Gray code sequence, we can implement all terms in the HUBO Hamiltonian using only one CNOT gate and one R_Z gate per term. This is illustrated in Fig. 16 for a four-qubit HUBO Hamiltonian. Welch et al. [74] have proven that this Gray code implementation is optimal in the sense that it requires the minimum number of CNOT gates to implement all terms in the HUBO Hamiltonian.

In practice, this means that a HUBO Hamiltonian of the form $\sum_{T \subseteq \{1, \dots, d\}} J_T \hat{Z}_T$ can be implemented with $2^{|d|} - 2$ CNOT gates and $2^{|d|} - 1$ R_Z gates. The number of CNOT gates is significantly reduced compared to the naive approach, which, as a calculation shows, requires $2^{|d|}(|d| - 2) + 2$ CNOT gates. An illustration of the optimized circuit for the three-qubit HUBO Hamiltonian, in comparison to the non-optimized implementation shown in Fig. 17, is presented in Fig. 18.

C Derivation: Scaling of Quantum Resources for QUBO and HUBO Formulations

In this section, we derive scaling relations for the quantum resources required to encode a quadratic COP in the QUBO and HUBO encodings. The COP we consider has a quadratic objective function of the form

$$C(\mathbf{s}) = \sum_{1 \leq i \leq n} c_1(i, s_i) + \sum_{1 \leq i < j \leq n} c_2(i, j, s_i, s_j). \quad (31)$$

We will consider the most general instance of this problem, where the linear coefficients $c_1(i, v)$ are nonzero for all variables i and all values v , and the quadratic coefficients $c_2(i, j, v, w)$ are nonzero for all variable combinations $i \neq j$ and all value pairs v, w . This instance then has $n \cdot m$ linear coefficients




Figure 16: Gray code traversal for optimal implementation of a 4-qubit HUBO Hamiltonian. Each vertex of the 4-dimensional hypercube represents a subset $T \subseteq \{1, 2, 3, 4\}$ of qubits, with subscripts indicating which qubits are included. The red path shows the Gray code sequence starting from the coefficient J_0 (bitstring 1000, the bitstring 0000 corresponds to a global phase and is ignored). Each edge corresponds to flipping exactly one bit, enabling implementation of consecutive terms with only one additional CNOT gate per term. This systematic traversal visits all $2^4 - 1 = 15$ vertices exactly once, requiring only $2^4 - 2 = 14$ CNOT gates total for the complete HUBO Hamiltonian $\sum_{T \subseteq \{1, 2, 3, 4\}} J_T \hat{Z}_T$.




Figure 17: Non-optimized implementation of the cost unitary for a three-qubit HUBO Hamiltonian $\sum_{T \subseteq \{1, 2, 3\}} J_T \hat{Z}_T$, where T runs over all subsets of the three qubits. Each term is implemented sequentially using parity circuits as shown in Fig. 15, resulting in $2^4(4 - 2) + 2 = 10$ CNOT gates and $2^4 - 1 = 7$ R_Z gates. This method incurs a higher CNOT count than the Gray code-based optimal circuit presented in Fig. 18, highlighting the benefit of circuit optimization for efficient quantum computation.




Figure 18: Optimal quantum circuit for implementing the cost unitary of a three-qubit HUBO Hamiltonian using the Gray code approach. Here, each subset $T \subseteq \{1, 2, 3\}$ defines a term $J_T \hat{Z}_T$ in the Hamiltonian. The upper panel illustrates the Gray code traversal of the three-dimensional hypercube, enabling efficient implementation with one additional CNOT and one R_Z gate per term. For three qubits, this results in $2^3 - 2 = 6$ CNOT gates and $2^3 - 1 = 7$ R_Z gates for all nontrivial subsets T , demonstrating the reduction in CNOT gates achieved by this optimized construction. Compared to the non-optimized implementation shown in Fig. 17, this approach significantly reduces the CNOT gate count.

and $\binom{n}{2} \cdot m^2$ quadratic coefficients. Of course, in practice, the objective function often has a simpler form, but here we consider the worst-case scaling of the quantum resources required for the problem.

C.1 QUBO encoding

The QUBO encoding uses $n \cdot m$ qubits, where each qubit represents a unique variable-value pair. The linear coefficients in the objective function give rise to one-body terms in the cost Hamiltonian. Since each one-body term is implemented using a single R_Z gate, the number of R_Z gates required to implement all linear terms is equal to the number of qubits.

The quadratic coefficients translate to two-body terms in the cost Hamiltonian, each requiring a circuit with two CNOT gates and one R_Z gate (see Fig. 15, left circuit). Since there are $\binom{n}{2} \cdot m^2$ quadratic coefficients, the number of CNOT gates required to implement all quadratic terms is $2 \cdot \binom{n}{2} \cdot m^2$, and the number of R_Z gates is $\binom{n}{2} \cdot m^2$. The asymptotic scaling of the number of CNOT and R_Z gates is therefore $\mathcal{O}(n^2 m^2)$.

C.2 HUBO encoding

In the HUBO encoding, each variable is represented by $d = \lceil \log_2(m) \rceil$ qubits. Therefore, the total number of qubits required is $n \cdot \lceil \log_2(m) \rceil$. The linear coefficients in the objective function give rise to a Hamiltonian term $\sum_{T \subseteq T_i} J_T \hat{Z}_T$, where T_i denotes the set of d qubits encoding variable i . This generates many-body interactions ranging from single-qubit terms (1st order) up to d -qubit terms (d th order) for each of the n variables.

The quadratic coefficients for variable combinations $i \neq j$ produce a Hamiltonian term $\sum_{T_1 \subseteq T_i, T_2 \subseteq T_j} J_{T_1, T_2} \hat{Z}_{T_1 \cup T_2}$. Here, T_1 and T_2 are subsets of the qubit indices, T_i and T_j , encoding variables i and j , respectively. Importantly, since $T_i \subseteq T_i \cup T_j$, the linear terms can be encoded simultaneously with the quadratic terms, eliminating the need for separate linear term implementation. As discussed in Appendix B, implementing the coefficients for all subsets of a set of qubit indices T requires $2^{|T|} - 1$ R_Z gates and $2^{|T|} - 2$ CNOT gates. Since the union $T_i \cup T_j$ involves $2d$ qubits, this

translates to $2^{2d} - 1$ R_Z gates and $2^{2d} - 2$ CNOT gates per variable pair. With $\binom{n}{2}$ variable pairs, the total number of R_Z gates required is $\binom{n}{2} \cdot (2^{2d} - 1)$ and the total number of CNOT gates is $\binom{n}{2} \cdot (2^{2d} - 2)$. With $d = \lceil \log_2(m) \rceil$, this gives us the asymptotic scaling of the CNOT and R_Z gates as $\mathcal{O}(n^2 m^2)$.

D Technical Details on QAOA Experiments

The exact computation of the approximation ratio in Eq. (25) requires summing over all feasible solutions, a process that scales exponentially with problem size in the considered benchmarks. Therefore, we estimate the approximation ratio using sampling-based methods. Specifically, the sum over all feasible states is replaced by a sum over sampled bitstrings,

$$\sum_{\mathbf{b} \in B_{\text{feas}}} \rightarrow \sum_{\mathbf{b} \in B_{\text{sample}}} \delta(\mathbf{b}), \quad (32)$$

where $\delta(\mathbf{b})$ equals one if the bitstring $\mathbf{b}$ satisfies all constraints, and zero otherwise. The required sample complexity is determined using Hoeffding’s inequality [75]. Employing $|B_{\text{sample}}| = 10000$ samples in all experiments ensures a 99% probability that the estimated value is within 2% of the true approximation ratio.

Furthermore, for all benchmarked COP instances and QAOA depths, we performed 100 independent QAOA runs. This ensures that all reported averages and standard deviations are directly comparable across experiments. For the target approximation benchmarks, we ran the QAOA algorithm 100 times and picked the QAOA experiment with the lowest layer requirements to reach the target threshold.

E Details on the GAP Instance

	Check-in/Luggage claim	Gate 1	Gate 2	Gate 3	Gate 4
Check-in/Luggage claim	0	10	10	20	20
Gate 1	10	0	20	20	20
Gate 2	10	20	0	20	20
Gate 3	20	20	20	0	1
Gate 4	20	20	20	1	0

Table 3: Walking times (in minutes) between check-in/luggage claim and airport gates 1-4 for the benchmarked GAP instance.

Aircraft i	$p_i^{\text{arr}} + p_i^{\text{dep}}$
0	75
1	74
2	62
3	88
4	61

Table 4: Number of arriving and departing passengers for each of the five aircraft in the benchmarked GAP instance.

For benchmarking, we use a GAP instance with $n = 5$ flights and $m = 4$ airport-gates, numbered 1 to 4, as shown in Fig. 6. The numbers of arriving and departing passengers per flight are given in Table 4, with values ranging from 61 to 88. Out of the total, there are 13 transfer passengers between flights 1 and 4, and 29 transfer passengers between flights 0 and 2. Transfer passengers contribute to the quadratic cost term $c_2(i, j, v, w)$ in Eq. (6), while the remaining passengers are included in the linear term $c_1(i, v)$ in Eq. (5).

The walking times between check-in/luggage claim and the airport-gates, as well as in between the airport-gates, are given in Table 3. In this instance, the walking times from each gate to check-in and from each gate to luggage claim are identical. For this reason, there is no difference in the cost contribution from arriving or departing passengers. Both groups affect the objective function in the same way, hence we don't distinguish between them in Table 4.

Gate assignments are further constrained by overlapping flight schedules. Any flights that overlap cannot be assigned to the same gate. The overlapping flight pairs in this instance are $\{(0, 1), (1, 2), (2, 3), (3, 4)\}$.

References

- [1] M. Marzec. Portfolio optimization: Applications in quantum computing, 2016. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118593486.ch4>.
- [2] A. Perdomo-Ortiz, N. Dickson, M. Drew-Brook, G. Rose, and A. Aspuru-Guzik. Finding low-energy conformations of lattice protein models by quantum annealing. *Scientific Reports*, 2(1): 571, August 2012. ISSN 2045-2322. DOI: 10.1038/srep00571. URL <https://www.nature.com/articles/srep00571>.
- [3] R. J. Boucherie, A. Braaksma, and H. Tijms. *Operations Research*. WORLD SCIENTIFIC, 2021. DOI: 10.1142/12343.
- [4] P. Crescenzi and V. Kann. A compendium of NP optimization problems, 1995. URL <https://cs.pwr.edu.pl/zielinski/lectures/om/compendium.pdf>.
- [5] Y. Fu and P. W. Anderson. Application of statistical mechanics to NP-complete problems in combinatorial optimisation. *Journal of Physics A: Mathematical and General*, 19(9):1605, June 1986. ISSN 0305-4470. DOI: 10.1088/0305-4470/19/9/033.
- [6] S. Kirkpatrick, Jr. Gelatt, C. D., and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983. DOI: 10.1126/science.220.4598.671.
- [7] F. Glover, E. TaiUard, and D. de Werra. A user's guide to tabu search. *Annals of Operations Research*, 41(1):1–28, March 1993. ISSN 1572-9338. DOI: 10.1007/BF02078647.
- [8] IBM CPLEX. IBM ILOG CPLEX Optimization Studio, November 2024. URL <https://www.ibm.com/products/ilog-cplex-optimization-studio>.
- [9] Gurobi Optimization. Gurobi optimization, July 2025. URL <https://www.gurobi.com/>.
- [10] L. Xu and L. Liberti. Relaxations for binary polynomial optimization via signed certificates, 2024. URL <https://arxiv.org/abs/2405.13447>.
- [11] J. Puchinger, G. R. Raidl, and U. Pferschy. The Multidimensional Knapsack Problem: Structure and Algorithms. *INFORMS Journal on Computing*, 22(2):250–265, May 2010. ISSN 1091-9856, 1526-5528. DOI: 10.1287/ijoc.1090.0344.
- [12] T. Packebusch and S. Mertens. Low Autocorrelation Binary Sequences. *Journal of Physics A: Mathematical and Theoretical*, 49(16):165001, April 2016. ISSN 1751-8113, 1751-8121. DOI: 10.1088/1751-8113/49/16/165001.
- [13] M. Danilova, P. Dvurechensky, A. Gasnikov, E. Gorbunov, S. Guminov, D. Kamzolov, and I. Shibaev. Recent Theoretical Advances in Non-Convex Optimization. In Ashkan Nikeghbali, Panos M. Pardalos, Andrei M. Raigorodskii, and Michael Th. Rassias, editors, *High-Dimensional Optimization and Probability: With a View Towards Data Science*, pages 79–163. Springer International Publishing, Cham, 2022. DOI: 10.1007/978-3-031-00832-0_3.
- [14] S. Burer and A. N. Letchford. Non-convex mixed-integer nonlinear programming: A survey. *Surveys in Operations Research and Management Science*, 17(2):97–106, July 2012. ISSN 1876-7354. DOI: 10.1016/j.sorms.2012.08.001. URL <https://www.sciencedirect.com/science/article/pii/S1876735412000037>.
- [15] C. A. Floudas, I. G. Akrotirianakis, S. Caratzoulas, C. A. Meyer, and J. Kallrath. Global optimization in the 21st century: Advances and challenges. *Computers & Chemical Engineering*,

- 29(6):1185–1202, May 2004. ISSN 0098-1354. DOI: [10.1016/j.compchemeng.2005.02.006](https://doi.org/10.1016/j.compchemeng.2005.02.006). URL <https://www.sciencedirect.com/science/article/pii/S0098135405000104>.
- [16] T. Albash and D. A. Lidar. Adiabatic quantum computation. *Reviews of Modern Physics*, 90(1):015002, January 2018. ISSN 0034-6861, 1539-0756. DOI: [10.1103/RevModPhys.90.015002](https://doi.org/10.1103/RevModPhys.90.015002).
- [17] S. Ebadi, A. Keesling, M. Cain, T. T. Wang, H. Levine, D. Bluvstein, G. Semeghini, A. Omran, J.-G. Liu, R. Samajdar, X.-Z. Luo, B. Nash, X. Gao, B. Barak, E. Farhi, S. Sachdev, N. Gemelke, L. Zhou, S. Choi, H. Pichler, S.-T. Wang, M. Greiner, V. Vuletić, and M. D. Lukin. Quantum Optimization of Maximum Independent Set using Rydberg Atom Arrays. *Science*, 376(6598):1209–1215, June 2022. ISSN 0036-8075, 1095-9203. DOI: [10.1126/science.abo6587](https://doi.org/10.1126/science.abo6587).
- [18] E. Farhi, J. Goldstone, and S. Gutmann. A quantum approximate optimization algorithm, 2014. URL <https://arxiv.org/abs/1411.4028>.
- [19] A. Lucas. Ising formulations of many NP problems. *Frontiers in Physics*, 2, February 2014. ISSN 2296-424X. DOI: [10.3389/fphy.2014.00005](https://doi.org/10.3389/fphy.2014.00005).
- [20] K. Goswami, R. Mukherjee, H. Ott, and P. Schmelcher. Solving optimization problems with local light shift encoding on Rydberg quantum annealers. *Physical Review Research*, 6(2):023031, April 2024. ISSN 2643-1564. DOI: [10.1103/PhysRevResearch.6.023031](https://doi.org/10.1103/PhysRevResearch.6.023031).
- [21] C. Lai, C. Blank, P. Schmelcher, and R. Mukherjee. Towards arbitrary qubo optimization: Analysis of classical and quantum-activated feedforward neural networks, 2024.
- [22] M. Zaman, K. Tanahashi, and S. Tanaka. Pyqubo: Python library for mapping combinatorial optimization problems to qubo form. *IEEE Transactions on Computers*, 71(4):838–850, 2022. DOI: [10.1109/TC.2021.3063618](https://doi.org/10.1109/TC.2021.3063618).
- [23] F. Dominguez, J. Unger, M. Traube, B. Mant, C. Ertler, and W. Lechner. Encoding-independent optimization problem formulation for quantum computing. *Frontiers in Quantum Science and Technology*, 2, September 2023. ISSN 2813-2181. DOI: [10.3389/frqst.2023.1229471](https://doi.org/10.3389/frqst.2023.1229471). URL <http://dx.doi.org/10.3389/frqst.2023.1229471>.
- [24] J. A. Montañez-Barrera, D. Willsch, A. Maldonado-Romo, and K. Michielsen. Unbalanced penalization: A new approach to encode inequality constraints of combinatorial problems for quantum optimization algorithms. *Quantum Science and Technology*, 9(2):025022, April 2024. ISSN 2058-9565. DOI: [10.1088/2058-9565/ad35e4](https://doi.org/10.1088/2058-9565/ad35e4).
- [25] S. V. Romero, A.-M. Visuri, A. Gomez Cadavid, A. Simen, E. Solano, and N. N. Hegade. Bias-field digitized counterdiabatic quantum algorithm for higher-order binary optimization. *Communications Physics*, 8(1), August 2025. ISSN 2399-3650. DOI: [10.1038/s42005-025-02270-3](https://doi.org/10.1038/s42005-025-02270-3). URL <http://dx.doi.org/10.1038/s42005-025-02270-3>.
- [26] A. Glos, A. Krawiec, and Z. Zimborás. Space-efficient binary optimization for variational quantum computing. *npj Quantum Information*, 8(1):39, April 2022. ISSN 2056-6387. DOI: [10.1038/s41534-022-00546-y](https://doi.org/10.1038/s41534-022-00546-y). URL <https://www.nature.com/articles/s41534-022-00546-y>.
- [27] S. V. Romero, A. Gomez Cadavid, P. Nikačević, E. Solano, N. N. Hegade, M. A. Lopez-Ruiz, C. Giroto, M. Yamada, P. Kl. Barkoutsos, A. Kaushik, and M. Roetteler. Protein folding with an all-to-all trapped-ion quantum computer, 2025. URL <https://arxiv.org/abs/2506.07866>.
- [28] C. Yuhui, E. Epifanovsky, K. Jansen, A. Kaushik, and S. Kühn. Simulating the flight gate assignment problem on a trapped ion quantum computer, 2023. URL <https://arxiv.org/abs/2309.09686>.
- [29] K. Wintersperger, F. Dommert, T. Ehmer, A. Hoursanov, J. Klepsch, W. Maurer, G. Reuber, T. Strohm, M. Yin, and S. Luber. Neutral atom quantum computing hardware: performance and end-user perspective. *EPJ Quantum Technology*, 10(1), August 2023. ISSN 2196-0763. DOI: [10.1140/epjqt/s40507-023-00190-1](https://doi.org/10.1140/epjqt/s40507-023-00190-1).
- [30] B. Fauseweh. Quantum many-body simulations on digital quantum computers: State-of-the-art and future challenges. *Nature Communications*, 15(1):2123, March 2024. ISSN 2041-1723. DOI: [10.1038/s41467-024-46402-9](https://doi.org/10.1038/s41467-024-46402-9).
- [31] F. Koch. PyHUBO, October 2025. URL <https://github.com/frederikKoch/PyHUBO>.

- [32] A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*, volume B. Journal of Computer and System Sciences - JCSS, 2003. URL <https://link.springer.com/book/9783540443896>.
- [33] T. G. Crainic, M. Gendreau, and A. Frangioni, editors. *Combinatorial Optimization and Applications: A Tribute to Bernard Gendron*, volume 358 of *International Series in Operations Research & Management Science*. Springer Nature Switzerland, Cham, 2024. ISBN 978-3-031-57602-7 978-3-031-57603-4. DOI: [10.1007/978-3-031-57603-4](https://doi.org/10.1007/978-3-031-57603-4).
- [34] J. Chen, H. Westerheim, Z. Holmes, I. Luo, T. Nuradha, D. Patel, S. Rethinasamy, K. Wang, and M. M. Wilde. Slack-variable approach for variational quantum semidefinite programming. *Physical Review A*, 112(2):022607, August 2025. ISSN 2469-9926, 2469-9934. DOI: [10.1103/lwxq-4myj](https://doi.org/10.1103/lwxq-4myj).
- [35] F. Glover, G. Kochenberger, and Y. Du. Quantum bridge analytics I: A tutorial on formulating and using QUBO models. *Annals of Operations Research*, 314(1):141–183, July 2022. ISSN 1572-9338. DOI: [10.1007/s10479-022-04634-2](https://doi.org/10.1007/s10479-022-04634-2).
- [36] A. Bouras, M. A. Ghaleb, U. S. Suryahatmaja, and A. M. Salem. The Airport Gate Assignment Problem: A Survey. *The Scientific World Journal*, 2014:1–27, 2014. ISSN 2356-6140, 1537-744X. DOI: [10.1155/2014/923859](https://doi.org/10.1155/2014/923859). URL <http://www.hindawi.com/journals/tswj/2014/923859/>.
- [37] Y. Chai, L. Funcke, T. Hartung, K. Jansen, S. Kühn, P. Stornati, and T. Stollenwerk. Towards Finding an Optimal Flight Gate Assignment on a Digital Quantum Computer. *Physical Review Applied*, 20(6):064025, December 2023. ISSN 2331-7019. DOI: [10.1103/PhysRevApplied.20.064025](https://doi.org/10.1103/PhysRevApplied.20.064025).
- [38] M. Bentert, R. van Bevern, and R. Niedermeier. Inductive k -independent graphs and k -colorable subgraphs in scheduling: A review. *Journal of Scheduling*, 22(1):3–20, February 2019. ISSN 1094-6136, 1099-1425. DOI: [10.1007/s10951-018-0595-8](https://doi.org/10.1007/s10951-018-0595-8).
- [39] M. M. Halldórsson, J. Y. Halpern, L. (Erran) Li, and V. S. Mirrokni. On spectrum sharing games. In *Proceedings of the Twenty-Third Annual ACM Symposium on Principles of Distributed Computing*, pages 107–114, St. John’s Newfoundland Canada, July 2004. ACM. ISBN 978-1-58113-802-3. DOI: [10.1145/1011767.1011783](https://doi.org/10.1145/1011767.1011783).
- [40] A. Hertz, R. Montagné, and F. Gagnon. Constructive algorithms for the partial directed weighted improper coloring problem. *Journal of Graph Algorithms and Applications*, 20(2):159–188, February 2016. ISSN 1526-1719. DOI: [10.7155/jgaa.00389](https://doi.org/10.7155/jgaa.00389).
- [41] A.M.C.A. Koster and M. Scheffel. A Routing and Network Dimensioning Strategy to reduce Wavelength Continuity Conflicts in All-Optical Networks, November 2006. URL <https://optimization-online.org/?p=10032>.
- [42] D. Liu, J. Li, X. Cheng, S. Zhang, Y. Chang, and L. Yan. Efficient hybrid variational quantum algorithm for solving graph coloring problem, 2025. URL <https://arxiv.org/abs/2504.21335>.
- [43] R. Quintero, D. Bernal, T. Terlaky, and L. F. Zuluaga. Characterization of qubo reformulations for the maximum k -colorable subgraph problem, 2021. URL <https://arxiv.org/abs/2101.09462>.
- [44] Z. Wang, N. C. Rubin, J. M. Dominy, and E. G. Rieffel. XY -mixers: Analytical and numerical results for QAOA. *Physical Review A*, 101(1):012320, January 2020. ISSN 2469-9926, 2469-9934. DOI: [10.1103/PhysRevA.101.012320](https://doi.org/10.1103/PhysRevA.101.012320).
- [45] M. Streif, M. Leib, F. Wudarski, E. Rieffel, and Z. Wang. Quantum algorithms with local particle number conservation: Noise effects and error correction. *Physical Review A*, 103(4):042412, April 2021. ISSN 2469-9926, 2469-9934. DOI: [10.1103/PhysRevA.103.042412](https://doi.org/10.1103/PhysRevA.103.042412).
- [46] R. Sotirov, O. Kuryatnikova, and J. Vera. The maximum k -colorable subgraph problem and related problems, 2021. URL <https://arxiv.org/abs/2001.09644>.
- [47] L. Wolsey. Integer programming. In *Integer Programming*, chapter 1, pages 1–23. John Wiley & Sons, Ltd, 2020. ISBN 978-1-119-60647-5. DOI: [10.1002/9781119606475.ch1](https://doi.org/10.1002/9781119606475.ch1).
- [48] P. Yves and A. W. Laurence. *Production Planning by Mixed Integer Programming*. Springer Series in Operations Research and Financial Engineering. Springer New York, 2006. ISBN 978-0-387-29959-4. DOI: [10.1007/0-387-33477-7](https://doi.org/10.1007/0-387-33477-7).

- [49] L. Magatão, L.V.R. Arruda, and F. Neves Jr. A Mixed Integer Programming Approach for Scheduling Commodities in a Pipeline. In Johan Grievink and Jan van Schijndel, editors, *Computer Aided Chemical Engineering*, volume 10 of *European Symposium on Computer Aided Process Engineering-12*, pages 715–720. Elsevier, January 2002. DOI: [10.1016/S1570-7946\(02\)80147-X](https://doi.org/10.1016/S1570-7946(02)80147-X).
- [50] K. Goswami, P. Schmelcher, and R. Mukherjee. Qudit-based scalable quantum algorithm for solving the integer programming problem, 2025. URL <https://arxiv.org/abs/2508.13906>.
- [51] M. Svensson, M. Andersson, M. Grönkvist, P. Vikstål, D. Dubhashi, G. Ferrini, and G. Johansson. Hybrid Quantum-Classical Heuristic to Solve Large-Scale Integer Linear Programs. *Physical Review Applied*, 20(3):034062, September 2023. ISSN 2331-7019. DOI: [10.1103/PhysRevApplied.20.034062](https://doi.org/10.1103/PhysRevApplied.20.034062).
- [52] M. Sharma and H.C. Lau. Cutting slack: Quantum optimization with slack-free methods for combinatorial benchmarks, 2025. URL <https://arxiv.org/abs/2507.12159>.
- [53] K. Tanahashi, S. Takayanagi, T. Motohashi, and S. Tanaka. Application of ising machines and a software development for ising machines. *Journal of the Physical Society of Japan*, 88(6):061010, 2019. DOI: [10.7566/JPSJ.88.061010](https://doi.org/10.7566/JPSJ.88.061010).
- [54] S. Hadfield. On the representation of Boolean and real functions as Hamiltonians for quantum computing. *ACM Transactions on Quantum Computing*, 2(4):1–21, December 2021. ISSN 2643-6809, 2643-6817. DOI: [10.1145/3478519](https://doi.org/10.1145/3478519).
- [55] 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(5516):472–475, April 2001. ISSN 0036-8075, 1095-9203. DOI: [10.1126/science.1057726](https://doi.org/10.1126/science.1057726).
- [56] S. McArdle, T. Jones, S. Endo, Y. Li, S. C. Benjamin, and X. Yuan. Variational ansatz-based quantum simulation of imaginary time evolution. *npj Quantum Information*, 5(1):75, September 2019. ISSN 2056-6387. DOI: [10.1038/s41534-019-0187-2](https://doi.org/10.1038/s41534-019-0187-2).
- [57] M. J. S. Beach, R. G. Melko, T. Grover, and T. H. Hsieh. Making trotters sprint: A variational imaginary time ansatz for quantum many-body systems. *Physical Review B*, 100(9):094434, September 2019. ISSN 2469-9950, 2469-9969. DOI: [10.1103/PhysRevB.100.094434](https://doi.org/10.1103/PhysRevB.100.094434).
- [58] P. J. Love. Cooling with imaginary time. *Nature Physics*, 16(2):130–131, February 2020. ISSN 1745-2481. DOI: [10.1038/s41567-019-0709-z](https://doi.org/10.1038/s41567-019-0709-z).
- [59] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien. A variational eigenvalue solver on a photonic quantum processor. *Nature Communications*, 5(1):4213, July 2014. ISSN 2041-1723. DOI: [10.1038/ncomms5213](https://doi.org/10.1038/ncomms5213).
- [60] J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth, and J. Tennyson. The Variational Quantum Eigensolver: A review of methods and best practices. *Physics Reports*, 986:1–128, November 2022. ISSN 0370-1573. DOI: [10.1016/j.physrep.2022.08.003](https://doi.org/10.1016/j.physrep.2022.08.003).
- [61] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin. Quantum Approximate Optimization Algorithm: Performance, Mechanism, and Implementation on Near-Term Devices. *Physical Review X*, 10(2):021067, June 2020. ISSN 2160-3308. DOI: [10.1103/PhysRevX.10.021067](https://doi.org/10.1103/PhysRevX.10.021067).
- [62] J. Basso, D. Gamarnik, S. Mei, and L. Zhou. Performance and limitations of the QAOA at constant levels on large sparse hypergraphs and spin glass models. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 335–343, October 2022. DOI: [10.1109/FOCS54457.2022.00039](https://doi.org/10.1109/FOCS54457.2022.00039).
- [63] K. Blekos, D. Brand, A. Ceschini, C.-H. Chou, R.-H. Li, K. Pandya, and A. Summer. A review on Quantum Approximate Optimization Algorithm and its variants. *Physics Reports*, 1068:1–66, June 2024. ISSN 0370-1573. DOI: [10.1016/j.physrep.2024.03.002](https://doi.org/10.1016/j.physrep.2024.03.002). URL <https://linkinghub.elsevier.com/retrieve/pii/S0370157324001078>.
- [64] J. Golden, A. Bärttschi, D. O’Malley, and S. Eidenbenz. Numerical Evidence for Exponential Speed-up of QAOA over Unstructured Search for Approximate Constrained Optimization. In *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, pages 496–505, September 2023. DOI: [10.1109/QCE57702.2023.00063](https://doi.org/10.1109/QCE57702.2023.00063).

- [65] J. Weidenfeller, L. C. Valor, J. Gacon, C. Tornow, L. Bello, S. Woerner, and D. J. Egger. Scaling of the quantum approximate optimization algorithm on superconducting qubit based hardware. *Quantum*, 6:870, December 2022. ISSN 2521-327X. DOI: [10.22331/q-2022-12-07-870](https://doi.org/10.22331/q-2022-12-07-870).
- [66] K. Kurowski, T. Pecyna, M. Słysz, R. Różycki, G. Waligóra, and J. Węglarz. Application of quantum approximate optimization algorithm to job shop scheduling problem. *European Journal of Operational Research*, 310(2):518–528, October 2023. ISSN 0377-2217. DOI: [10.1016/j.ejor.2023.03.013](https://doi.org/10.1016/j.ejor.2023.03.013).
- [67] S.-S. Wang, H.-L. Liu, Y.-Q. Song, F. Gao, S.-J. Qin, and Q.-Y. Wen. Quantum alternating operator ansatz for solving the minimum exact cover problem. *Physica A: Statistical Mechanics and its Applications*, 626:129089, September 2023. ISSN 0378-4371. DOI: [10.1016/j.physa.2023.129089](https://doi.org/10.1016/j.physa.2023.129089).
- [68] N. Sachdeva, G. S. Hartnett, S. Maity, S. Marsh, Y. Wang, A. Winick, R. Dougherty, D. Canuto, Y. Q. Chong, M. Hush, P. S. Mundada, C. D. B. Bentley, M. J. Biercuk, and Y. Baum. Quantum optimization using a 127-qubit gate-model ibm quantum computer can outperform quantum annealers for nontrivial binary optimization problems, 2024. URL <https://arxiv.org/abs/2406.01743>.
- [69] M. Blondel, Q. Berthet, M. Cuturi, R. Frostig, S. Hoyer, F. Llinares-López, F. Pedregosa, and J.-P. Vert. Efficient and modular implicit differentiation, 2022. URL <https://arxiv.org/abs/2105.15183>.
- [70] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, S. Ahmed, V. Ajith, M. S. Alam, G. Alonso-Linaje, B. AkashNarayanan, A. Asadi, J. M. Arrazola, U. Azad, S. Banning, C. Blank, T. R. Bromley, B. A. Cordier, J. Ceroni, A. Delgado, O. Di Matteo, A. Dusko, T. Garg, D. Guala, A. Hayes, R. Hill, A. Ijaz, T. Isaacsson, D. Ittah, S. Jahangiri, P. Jain, E. Jiang, A. Khandelwal, K. Kottmann, R. A. Lang, C. Lee, T. Loke, A. Lowe, K. McKiernan, J. J. Meyer, J. A. Montañez-Barrera, R. Moyard, Z. Niu, L. J. O’Riordan, S. Oud, A. Panigrahi, C.-Y. Park, D. Polatajko, N. Quesada, C. Roberts, N. Sá, I. Schoch, B. Shi, S. Shu, S. Sim, A. Singh, I. Strandberg, J. Soni, A. Száva, S. Thabet, R. A. Vargas-Hernández, T. Vincent, N. Vitucci, M. Weber, D. Wierichs, R. Wiersema, M. Willmann, V. Wong, S. Zhang, and N. Killoran. PennyLane: Automatic differentiation of hybrid quantum-classical computations, 2022. URL <https://arxiv.org/abs/1811.04968>.
- [71] S. Schulz, D. Willsch, and K. Michielsen. Guided quantum walk. *Physical Review Research*, 6(1):013312, March 2024. ISSN 2643-1564. DOI: [10.1103/PhysRevResearch.6.013312](https://doi.org/10.1103/PhysRevResearch.6.013312).
- [72] Z. Verchère, S. Elloumi, and A. Simonetto. Optimizing variational circuits for higher-order binary optimization, 2023. URL <https://arxiv.org/abs/2307.16756>.
- [73] M. Amy, P. Azimzadeh, and M. Mosca. On the CNOT-complexity of CNOT-PHASE circuits. *Quantum Science and Technology*, 4(1):015002, September 2018. ISSN 2058-9565. DOI: [10.1088/2058-9565/aad8ca](https://doi.org/10.1088/2058-9565/aad8ca).
- [74] J. Welch, D. Greenbaum, S. Mostame, and A. Aspuru-Guzik. Efficient quantum circuits for diagonal unitaries without ancillas. *New Journal of Physics*, 16(3):033040, March 2014. ISSN 1367-2630. DOI: [10.1088/1367-2630/16/3/033040](https://doi.org/10.1088/1367-2630/16/3/033040).
- [75] W. Hoeffding. Probability Inequalities for Sums of Bounded Random Variables. *Journal of the American Statistical Association*, 58(301):13–30, March 1963. ISSN 0162-1459. DOI: [10.1080/01621459.1963.10500830](https://doi.org/10.1080/01621459.1963.10500830).