

Out of the Loop: Structural Approximation of Optimisation Landscapes and non-Iterative Quantum Optimisation

Tom Krüger^{1,2} and Wolfgang Mauerer^{1,3}

¹Technical University of Applied Sciences Regensburg, Regensburg, Germany

²FI CODE, Universität der Bundeswehr München, Munich, Germany

³Siemens AG, Technology, Munich, Germany

The Quantum Approximate Optimisation Algorithm (qaoa) is a widely studied quantum-classical iterative heuristic for combinatorial optimisation. While qaoa targets problems in complexity class NP, the classical optimisation procedure required in every iteration is itself known to be NP-hard. Still, advantage over classical approaches is suspected for certain scenarios, but nature and origin of its computational power are not yet satisfactorily understood.

By introducing means of efficiently and accurately approximating the qaoa optimisation landscape from solution space structures, we derive a new algorithmic variant of unit-depth qaoa for two-level Hamiltonians (including all problems in NP): Instead of performing an iterative quantum-classical computation for each input instance, our non-iterative method is based on a quantum circuit that is instance-independent, but problem-specific. It matches or outperforms unit-depth qaoa for key combinatorial problems, despite reduced computational effort.

Our approach is based on proving a long-standing conjecture regarding instance-independent structures in qaoa. By ensuring generality, we link existing empirical observations on qaoa parameter clustering to established approaches in theoretical computer science, and provide a sound foundation for understanding the link between structural properties of solution spaces and quantum optimisation.

1 Introduction

With the advent of early commercially available quantum computers, interest in the field has spilled from academia to early industrial adopters, and the hope for possible advantage or even supremacy is manifest [11, 63, 72]. However, challenges arise not only from deficiencies of noisy, intermediate-scale quantum

(NISQ) hardware [17, 19, 40, 87], but quantum algorithmic theory in general lags behind the classical case. While fundamental complexity-theoretic boundaries have long been established [4, 16, 35], a more precise understanding of concrete algorithmic building blocks and how to construct them is required [2, 62, 78].

How to systematically construct quantum algorithms is a multi-faceted, highly non-trivial task that remains essentially unsolved. Instead, heuristics like variational quantum circuits (VQC) [23] have emerged as popular alternatives. Originally intended to extract useful computational power from NISQ devices despite their limitations, they may also prove relevant as resource-efficient primitives in the post-NISQ era. They are centred around an iterative quantum-classical process that learns a parameterised form of a quantum circuit such that sampling a produced quantum state obtains a valid solution with high probability. Depending on the variational ansatz, this approach defers considerable aspects (*i.e.*, finding an appropriate quantum circuit that eventually implements the optimisation routine) of the task to classical components [32, 69, 98].

Structured forms of VQCs, in particular the quantum approximate optimisation algorithm (QAOA), enjoy popularity in prototypical applications [14]. QAOA is a specific VQC for combinatorial optimisation. Similar to quantum annealing, to which QAOA is closely related [17, 70], it provides a strict framework on the quantum side. Given a specific problem, users only need to find a fitting problem representation, choose a suitable circuit depth, and pick a classical optimisation method. Devising such problem representations is not foreign to traditional computer science: In particular, techniques to reduce computational problems specified using an apt formalism into representations that are suitable for QAOA [24, 57, 76] are curricular knowledge [10].

The most essential ingredient of combinatorial optimisation is, of course, the optimisation landscape. This paper is centred around a new theorem, detailed in Section 4, that, for two level Hamiltonians (including all problems in NP), allows us to approximate the expected one layer QAOA optimisation landscape of a problem from existing solution space structures.

Tom Krüger: tom.krueger@oth-regensburg.de

Wolfgang Mauerer: wolfgang.mauerer@othr.de




Figure 1: Overview of our main foundational contributions. *Left*: Previously conjectured parameter clustering in QAOA optimisation landscapes $F_1^{(i)}(\beta, \gamma)$ (derived from a combinatorial optimisation objective function $c^{(i)}(\vec{z})$ of a decision problem instance) suggests that shared macroscopic similarities exist between instances i . Averaging over these reveals that the expected landscape $E(F_1(\beta, \gamma))$ is a common structure at the problem-global level. *Right*: Shared macroscopic features exist across all instances, based on structural properties manifest in the solution spaces. Aggregation (which may be possible analytically, but can always be performed using empirical sampling) leads to a macroscopic description of a problem-global solution space, from which the expected optimisation landscape can be efficiently approximated as $\tilde{E}(F_1(\beta, \gamma))$. Most importantly, we prove that the approximation has a bounded difference to the underlying exact quantity. The approximation approach (ingredients indicated by a blue background) and its consequences are subject of this paper.

Informally speaking, QAOA is concerned with multiple objects: A problem (*e.g.*, can a graph be partitioned into two halves such that only a certain amount of edges must be cut?), an instance (*e.g.*, a specific graph), a solution space (*e.g.*, lists of edges to cut), and parameters that define an appropriate quantum circuit specific for each *instance* to obtain solutions from sampling the circuit output.

Hamming distances between solutions offer an intuitive handle to discover and describe local structures in solution space. Our approach continues a line of research [21, 27, 63, 83] using such distance information, particularly the connection to state amplitudes and their interference, to obtain insights about the inner workings of QAOA. By using stochastic methods, we establish novel means of reasoning about QAOA landscapes on an *instance independent*, but *problem specific* level, and show that instance-specific quantum circuits can be replaced by one single problem-generic alternative, while obeying a strictly bounded maximal difference between approximated and exact ingredients of the overall combinatorial optimisation process. Our result contributes to both, foundational understanding of QAOA and practical implementations.

As for the fundamental aspects, we provide a rigorous mathematical proof of a long-standing empirically motivated conjecture [20, 37, 73, 84, 85] that instance independent structures of a problem form the global structure of the QAOA optimisation landscape. We combine methods from physics and computer science to ascertain this structural insight, and can use either efficient sampling techniques or an analytical approach to obtain an accurate approximation of the QAOA optimisation landscape. In particular, we can separate the instance sampling from specific QAOA parameters. These aspects of our contribution are visually summarised in Fig. 1.

In terms of practical impact, we provide a unit-depth non-iterative QAOA version for two-level problem Hamiltonians. Such two-level Hamiltonians are powerful enough to efficiently encode all problems in **NP**, allowing for the construction of QAOA circuits to solve said problems. Thus, for **NP**-Problems, we can replace the standard iterative QAOA heuristic comprising an interplay of quantum and classical components (that, notably, requires solving an **NP**-hard parameter optimisation problem [18] individually for *every* instance) with a two-phase, non-iterative algorithm that first approximates the instance-independent, but problem-specific expected landscape followed by sampling a fixed quantum circuit, as illustrated in Fig. 2. While omitting the outer loop has, in various ways, been suggested in prior work based on the aforementioned empirical observations [20, 84], our approach places the idea on a sound theoretical basis. It also introduces a new method of devising classical optimal parameters, and equips the method with a quantitative quality criterion. We demonstrate that our simplified

variant is identical or better to standard QAOA for several seminal subject problems in terms of successfully finding solutions to combinatorial optimisation problems.

By showing that the required structures exist for all problems in **NP**, we establish generality of our results. Our approach therefore opens up the discussion on parameter clustering and optimisation landscape similarities to the vast body of established results in theoretical computer science on structural properties of problem solutions.

The rest of this paper is organised as follows: Section 2 first presents related work, supporting and motivating our ideas, and then establishes precise terminology for QAOA fundamentals, particularly given the many subtly different variants in current use. Section 3 covers the QAOA optimisation landscape and prepares the necessary groundwork for Section 4, where the main approximation theorem is presented and proved. With the new theorem at hand, we discuss a list of five examples building on each other in Section 5. This illustrates how our theorem can be used to analyse synthetic and realistic computational problems. The practical application of our insights is subject of Section 6, where we introduce a novel non-iterative variant of QAOA based on our insights that matches or exceeds the performance of the standard heuristic for a set of subject problems, yet at substantially reduced computational cost. After setting our contributions in context with the current state of the art, and discussing potential further uses and future improvements in Section 7, we conclude in Section 8.

2 Foundations

2.1 Context and Related Work

QAOA and variants of the algorithm have been subject of intensive research [5, 6, 13, 15, 21, 22, 25, 28, 32–34, 37, 39, 41–43, 45, 49, 54–56, 59, 60, 64–67, 69, 73, 74, 76, 79–81, 85–87, 89–91, 93, 95, 96], as recently reviewed by Blekos *et al.* [19] or Zhou *et al.* [98]. Apart from improving the understanding of the heuristic construction, modifications to the structure of the quantum circuit itself (*e.g.*, Refs. [21, 22, 50, 90, 98]) aim at improving performance especially in NISQ scenarios. Likewise, changes to the classical optimisation procedure (*e.g.*, Refs. [13, 28, 34, 63, 84–86, 89]) have been proposed. Given that limitations of NISQ hardware restrict programs to shallow circuits, many analytical and practical studies focus on unit-depth QAOA [19, 21, 28, 29, 33, 37, 42, 48, 56, 73, 89]; our considerations are also based on this commonly employed scenario. Perhaps fuelled by the possibility of empirically investigating early-stage deployments, QAOA has been applied to a considerable variety of practical problems; see Bayerstadler *et al.* [14] for a review.




Figure 2: Overview of our main practical contributions (blue background ■ indicates *problem-global*, grey background ■ *instance-specific* components; yellow boxes ■ mark classical computation. Dashed lines - - - indicate transfer of classical data, and double lines — denote querying a resource). *Left*: Standard QAOA that iteratively (symbolised by ♻️) determines optimal parameters (β, γ) for every instance by repeatedly sampling a structured quantum circuit. *Right*: Two-phase QAOA approach introduced in this paper that first determines optimal parameters (β, γ) by sampling $\tilde{E}(F_1(\beta, \gamma))$ from the problem-global target space, and then uses the derived instance-independent optimal parameters (β, γ) to obtain instance-specific solutions by sampling from a quantum circuit that is constant for each instance $\hat{C}$. As we show, $\tilde{E}(F_1)$ has a strictly bounded difference to the expectation value of F_1 , the core quantity of interest in QAOA.

Our work is particularly motivated by efforts that observe concentrations of optimal QAOA parameters across instances, which has received substantial consideration in the literature. In 2020, Streif and Leib [84] observed a clustering of optimal parameters when training QAOA circuits for random Max-Cut instances, which inspired them to propose training QAOA without quantum hardware access. Wybo and Leib, recently used these methods to analyse fixed parameter QAOA landscapes for Ising formulations of the maximum cut and the maximum independent set problems [97]. Here, they also observed a promising performance of QAOA with fixed parameters. Sack and Serbyn [73] reported similar parameter concentrations in 2021 also for Max-Cut instances. They provided a theoretical definition of the effect based on the closeness of optimal sets in parameter space. Their highlighting the importance of picking good initialisation values for QAOA parameters evolved into *warm start* QAOA: Here, the focus lies on finding optimal initialisation parameters [28]; this is obviously tightly linked to optimal parameters concentrating in certain regions. Predicting the location of those clusters would benefit warm start efforts. Following that, Akshay et al. published a thorough analytic analysis, providing a new definition and concrete analytic insights [5]. They, for instance, showed that QAOA circuit parameters concentrate as an inverse polynomial in problem size. Galda et al. [37] demonstrated the transferability of QAOA parameters between different Max-Cut instances in 2023, and explained their findings with local graph properties. They also investigated optimisation landscapes of certain sub-graphs and observed stark similarities between different sub-graphs. This overlaps with earlier observations by Brandao et al. in 2018 [20] that fixed parameters of

the optimisation landscape concentrate around certain values for different instances. All these findings suggest the existence of higher level problem structures shared between instances that significantly influence the shape of the underlying optimisation landscapes.

Classical theoretical computer science has studied structural properties of problems for decades. Among the most fundamental structural observations are phase transitions in constraint satisfaction problems: At a certain point, problem instances transition from under-constrained to over-constrained. The parameter correlated with this effect depends on the problem: For the seminal problem of Boolean satisfiability (SAT), it is known to be the ratio of numbers of clauses to the number of variables in an instance [61]. For graph colouring, the connectivity of the underlying graph is the relevant determinant. Finding a solution is especially hard at the phase transition. While the phase transitions provide a high level description of problem hardness, more complex structural properties are known. Hogg argued in 1996 [47] that real-world applications model interactions between physical or social entities. Thus, interaction likelihood strongly depends on the distance between entities based on some appropriate distance measure, leading to recurring local structures. In 2004, Pari et al. [68] showed that even trivially sampled random SAT instances do possess structured solution spaces. The probability of a potential solution satisfying the SAT formula depends on the Hamming distance to other solutions. A few years later in 2011, Achlioptas et al. [3] discovered a clustering of the solution space of SAT instances. They further proved that under-constrained SAT formulas have exponentially many small clusters in their solution space. The presence of community structures

in the solution space of industrial SAT instances was demonstrated in 2012 by Ansótegui et al. [8].

2.2 Quantum Approximate Optimisation Algorithm

More often than not, the *Quantum Approximate Optimisation Algorithm* (QAOA) is associated with the *quadratic unconstrained binary optimisation* (QUBO) problem, which is NP-complete in its decision form and relatively straight forward to solve with QAOA. A QUBO problem is defined by Boolean quadratic formula of the form $\sum_{i \neq j} a_{i,j} x_i x_j + \sum_i a_i x_i$, where $a_{i,j}, a_i \in \mathbb{R}$ are real valued weights of the Boolean variables $x_i \in \mathbb{F}_2$, with $\mathbb{F}_2 := \{0, 1\}$. The goal is to find a variable assignment to maximise this QUBO formula. This can be easily formulated as a ground state problem of an Ising Hamiltonian $\sum_{i \neq j} -J_{i,j} \sigma_i^z \sigma_j^z - \sum_i h_j \sigma_i^z$, with $J_{i,j}, h_j \in \mathbb{R}$ and σ_i^z being the Pauli-Z operator on the i -th qubit. Therefore, it seems reasonable to view QAOA as a dedicated QUBO solver. This approach is analogous to how SAT solvers are employed in classical systems. There is a rich community of SAT experts working on newer and better solvers, while the users on the other side can rely on the interface of abstract SAT formulas to solve their concrete use cases without them needing to dive deep into the intricacies of Boolean satisfiability. Similarly, if we look at QAOA as just another QUBO solver, this takes the majority of *quantum* out of quantum computing as the QUBO formalism serves as a classical interface. This is arguably a major reason why QUBOs have been a welcoming entry point to quantum computing, particularly for researchers who do not feel the need to understand details of the computational process [52, 57]. Accompanying that, promising methods of solving QUBO problems with means other than quantum computing have enjoyed a certain amount of attention [7, 9, 44, 77, 100].

While our results apply to QUBO problems, our considerations are not restricted to this scenario, but consider QAOA as proposed by Farhi et al. to solve general combinatorial optimisation problems [29], of which QUBOs only form a restricted subset.

Definition 1. Let $z \in \mathbb{F}_2^n$ be a n -bit binary string. Further $\{c_\alpha\}_{\alpha=1}^m$ shall be a set of Boolean clauses with $c_\alpha(z) = 1$ iff z satisfies c_α and $c_\alpha(z) = 0$ otherwise. Maximising

$$c(z) = \sum_{\alpha=1}^m c_\alpha(z) \quad (1)$$

is known as the *Boolean constraint optimisation problem*.

We can bring Definition 1 to the quantum world by defining a corresponding basis state vector $|z\rangle$ for each bit string $z \in \mathbb{F}_2^n$. The obvious choice for $|z\rangle$ is the computational basis vector $|z\rangle \in \{|0\rangle, |1\rangle\}^{\otimes n}$ encoded

by z . Then we map the clause values $c_\alpha(z)$ to eigenvalues of quantum operators C_α representing the clauses. For a one to one mapping, we get a projector C_α per clause c_α that projects onto the subspace spanned by all states representing satisfying assignments of c_α . With Hermitian operators being closed under addition, we have that $C = \sum_{\alpha=1}^m C_\alpha$ is itself a Hermitian operator. This allows us to define the Hamiltonian time evolution operator $e^{-i\gamma C}$. Hamiltonian C describes the energy of a quantum system, and energy levels are eigenvalues of C . The solution z of Definition 1 maximises Eq. (1) and thus the corresponding eigenstate $|z\rangle$ of C has eigenvalue $\lambda_{\max} = \max \sigma(C)$. Here, $\sigma(C)$ is the set of eigenvalues of C . From the Hamiltonian point of view, solving the combinatorial optimisation problem is equivalent to finding a state with maximal energy of a system described by the Hamiltonian C . Farhi et al. came up with QAOA by trotterising an interpolated Hamiltonian time evolution from an easy to prepare maximal energy state of a system described by a simple Hamiltonian to the state of maximal energy of the system described by the constraint Hamiltonian C , in which the problem structure of Definition 1 is encoded.

Definition 2 (QAOA circuit as in Ref. [29]). We consider a constraint Hamiltonian C implementing the constraint cost function Eq. (1) for $z \in \mathbb{F}_2^n$ and a mixer Hamiltonian $X = \sum_{j=1}^n \sigma_j^x$, with $\sigma_i^x = \mathbb{1}^{\otimes i-1} \otimes \sigma^x \otimes \mathbb{1}^{n-i}$. Then, the QAOA circuit

$$U_p(\beta, \gamma) = e^{-i\beta_p X} e^{-i\gamma_p C} \dots e^{-i\beta_1 X} e^{-i\gamma_1 C} \quad (2)$$

produces the state

$$|\beta, \gamma\rangle_p = U_p(\beta, \gamma)|+\rangle^{\otimes n} \quad (3)$$

with real angles $\beta, \gamma \in \mathbb{R}^p$, $p \in \mathbb{N}$.

Definition 2 actually defines a parameterised family of circuits $\{U_p(\beta, \gamma)\}_{\beta, \gamma, p}$. For the Hamiltonian implementation of Eq. (1), the evaluation $c(z)$ can be performed by calculating the expectation value $\langle z|C|z\rangle$. Farhi et al. showed that $\lim_{p \rightarrow \infty} \langle \beta, \gamma|_p C |\beta, \gamma\rangle_p = \max_z c(z)$. This lets one define an algorithm to approximate the combinatorial optimisation problem with the circuit defined in Eq. (2).

Definition 3 (QAOA). Consider a combinatorial optimisation problem with constraint cost function $c(z)$ and $z \in \mathbb{F}_2^n$. For a fixed $p \in \mathbb{N}$, choose a set of angles $\beta, \gamma \in \mathbb{R}^p$ that maximises the QAOA cost function

$$F_p(\beta, \gamma) = \langle \beta, \gamma|_p C |\beta, \gamma\rangle_p \quad (4)$$

Then construct the circuit $U_p(\beta, \gamma)$, with which the state $|\beta, \gamma\rangle_p$ will be prepared and measured in the computational basis to produce a binary string $z \in \mathbb{F}_2^n$. Repeat this sampling step m times with the same circuit to get a binary string that is close to $\max_z c(z)$ with high probability.

Strictly speaking, Definition 3 defines a heuristic and not an algorithm. The process of finding optimal angles $\beta, \gamma \in \mathbb{R}^p$ is not further specified and open to interpretation. As shown by Farhi et al. [29], F_p can be simplified for specific problems—they consider Max-Cut on 3-regular graphs—which allows for finding an efficient classical evaluation of F_p . It would also be feasible to evaluate F_p on quantum hardware. The possible parameter optimisation methods are plentiful [19, 32, 69], and their impact is subject to ongoing research.

Even for $p = 1$, the parameter optimisation problem of $F_1(\beta, \gamma)$ is **NP**-hard [18]. We restrict our considerations to a single layer, as is common practice [19, 21, 28, 29, 33, 37, 42, 48, 56, 73, 89]. For the sake of simplicity, we also focus on constraint Hamiltonians with two-level eigenspectra. This avoids some effort in notation for the following theorems, but still allows us to solve problems in **NP**. Note that the structural properties we prove below also exist for constraint Hamiltonians C with more than two distinct eigenvalues. Also, two-level constraint Hamiltonians include all decision problems with classical proofs $\mathbf{z} \in \mathbb{F}_2^n$, most notably the complete class of **NP**. Despite the restriction to $|\sigma(C)| = 2$ and $p = 1$, our setting therefore covers a large body of non-trivial, interesting computational problems.

Before we proceed with our analysis, let us fix terminology regarding QAOA. The unitary gates $e^{-i\gamma_i C}$ defined by the Hamiltonian time evolution of C are usually called phase separation gates or just phase separators. As becomes clear in Eq. 7 below, it separates the solution space from the search space by a complex phase, and hence earns its name. The term $e^{-i\beta_i X}$ is usually called mixer, and X is the mixer Hamiltonian. Together, a phase separator and mixer pair forms a layer $e^{-i\beta_i X} e^{-i\gamma_i C}$. In this case we speak of the i -th layer of a p -layer QAOA circuit $U_p(\beta, \gamma)$.

3 The Optimisation Landscape

Now we want to take a look at the optimisation landscape induced by the decision problem derived from Definition 1, which asks whether there exists an assignment $\mathbf{z} \in \mathbb{F}_2^n$ satisfying all clauses—or, equivalently, if $c(\mathbf{z}) = \prod_{\alpha} c_{\alpha}(\mathbf{z}) = 1$. This directly translates to the Hamiltonian implementation $C = \prod_{\alpha} C_{\alpha}$ which, as a product of projectors, remains a projector itself. Thus, the eigenspectrum of C is $\sigma(C) = \{0, 1\}$. Since C is Hermitian, there exists a unitary diagonalisation $C = U \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} U^{\dagger}$. Note that we can collect all eigenvalues $\lambda = 1$ in the upper left block of the diagonal matrix by simply applying a permutation operator, which we can subsume into the unitaries U and $U^{\dagger}$. Using the power series expansion of the exponential function, we see that the exponentiation can be passed through the diagonalisation. Let H be a diagonalisable operator with $H = UDU^{\dagger}$ with

$D = \text{diag}(d_1, \dots, d_n)$ and unitary U . Then using the power series expansion of the exponential function, we get $e^H = \sum_{k=0}^{\infty} \frac{1}{k!} (UDU^{\dagger})^k$. Since U is a unitary operator, the inner contributions $U^{\dagger}U = \mathbf{1}$ vanish in the product, and $(UDU^{\dagger})^k = UD^kU^{\dagger}$. It follows that

$$e^H = \sum_{k=0}^{\infty} \frac{1}{k!} UD^kU = U \left(\sum_{k=0}^{\infty} \frac{1}{k!} D^k \right) U^{\dagger}. \quad (5)$$

Since D is diagonal, we find that

$$\begin{aligned} \sum_{k=0}^{\infty} D^k &= \text{diag} \left(\sum_{k=0}^{\infty} d_1^k, \dots, \sum_{k=0}^{\infty} d_n^k \right) \\ &= \text{diag}(e^{d_1}, \dots, e^{d_n}). \end{aligned} \quad (6)$$

Consequently, we can express the phase separator by $e^{-i\gamma C} = U \begin{pmatrix} e^{-i\gamma} \mathbf{1} & 0 \\ 0 & \mathbf{1} \end{pmatrix} U^{\dagger}$. Let $|z\rangle$ be a computational basis state. Then $(e^{-i\gamma} \mathbf{1} \oplus \mathbf{1})U^{\dagger}|z\rangle = e^{-i\gamma}U^{\dagger}|z\rangle = U^{\dagger}e^{-i\gamma}|z\rangle$ iff $|z\rangle$ has eigenvalue 1 under C , and $(e^{-i\gamma} \mathbf{1} \oplus \mathbf{1})U^{\dagger}|z\rangle = U^{\dagger}|z\rangle$ otherwise. If we apply the full phase separator gate on the state, we obtain $e^{-i\gamma C}|z\rangle = U \begin{pmatrix} e^{-i\gamma} \mathbf{1} & 0 \\ 0 & \mathbf{1} \end{pmatrix} U^{\dagger}|z\rangle = UU^{\dagger}e^{i\gamma}|z\rangle = e^{i\gamma}|z\rangle$ iff $|z\rangle$ has eigenvalue 1 under C , and $UU^{\dagger}|z\rangle = |z\rangle$ otherwise. We conclude that $e^{-i\gamma C}$ adds a global phase to a computational basis state $|z\rangle$ if and only if $|z\rangle$ has the eigenvalue 1 under C , and leaves the state invariant otherwise. Applying $e^{-i\gamma C}$ to an arbitrary superposition of computational basis states $|\psi\rangle = \sum_{z=0}^{2^n-1} \omega_z |z\rangle$ can be characterised by

$$e^{-i\gamma C}|\psi\rangle = e^{-i\gamma} \sum_{|z\rangle \in T} \omega_z |z\rangle + \sum_{|z\rangle \notin T} \omega_z |z\rangle. \quad (7)$$

Here, $T = \{|z\rangle \mid c(z) = 1\}$ is the target space of C , which directly maps to the solution space of Eq. (1). At this point, the amplitude symmetry is broken in the QAOA circuit, which allows for establishing interference effects whose pattern can be controlled by γ and β . This, eventually, allows us to benefit from quantum effects in the computational process.

The angle parameters β also have an interesting effect on state amplitudes. As we will shortly show in detail, solely the Hamming distances between states and the target/non-target partition of the state space suffice to describe the inner workings of QAOA circuits for decision problems. However, we need to analyse the effect of mixer layers before we can commence to proving this statement.

Lemma 1 (Projector version of [27]). *Let $X = \sum_{j=1}^n \sigma_j^x$ be the n -qubit mixer Hamiltonian. The effect of $e^{-i\beta X}$ on an arbitrary basis state $|z\rangle \in \mathbb{F}_2^n$ can be characterised by*

$$e^{-i\beta X}|z\rangle = \sum_{k \in \mathbb{F}_2^n} f(\beta, z, k) |k\rangle \quad (8)$$

where f is defined as

$$f(\beta, z, k) := (\cos \beta)^{n-d_H(z,k)} (-i \sin \beta)^{d_H(z,k)}, \quad (9)$$

and $d_H(z, k)$ is the Hamming distance of z and k .

Proof. Given that X is a 1-local Hamiltonian, we can analyse its time evolution by only considering its effect on single qubits. Let us look at $n = 1$ first: In this case, $X = \sigma^x$, and $e^{-i\beta X}|z\rangle = \cos \beta|z\rangle - i \sin \beta|z \oplus 1\rangle$ for $z \in \mathbb{F}_2$, which performs a bit flip with probability $(\sin \beta)^2$. For arbitrary $n \in \mathbb{N}$, this generalises to

$$e^{-i\beta X}|z\rangle = \bigotimes_{l=1}^n (\cos \beta|z_l\rangle - i \sin \beta|z_l \oplus 1\rangle). \quad (10)$$

We would now like to express this state as a superposition of computational basis states, which can be achieved by reconstructing each of its components $|k \in \mathbb{F}_2^n\rangle$. From Eq. (10), we see that the amplitude of $|k\rangle$ acquires a multiplicative pre-factor of either $-i \sin \beta$ for each flipped, or $\cos \beta$ for each preserved qubit, respectively. In other words, the eventual amplitude of state $|k\rangle$ is given by $f(\beta, z, k) = (\cos \beta)^{n-d_H(z, k)} \cdot (-i \sin \beta)^{d_H(z, k)}$. $\square$

Now that we have characterised the mixer and phase separation layers individually, we have the tools to further our analysis. Next in line is the optimisation landscape, which is a central point of interest when analysing QAOA circuits.

Now that we have characterised the mixer and phase separation layers individually, we have the tools to further our analysis. Next in line is the optimisation landscape, which is a central point of interest when analysing QAOA circuits. In Lemma 2 we use that the landscape F_1 is a scaled sum of terms c_k that for fixed angles β and ϕ only depend on the Hamming distances between k and all other states $z \in \mathbb{F}_2^n$, which allows us to simplify the equation by collecting all terms with equal hamming distance. For a visualization see Fig. 3.

Lemma 2. *A QAOA circuit with $p = 1$ constructed for general decision problems induces the optimisation landscape*

$$F_1(\beta, \gamma) = \frac{1}{2^n} \sum_{k \in T} |c_k(\beta, \gamma)|^2 \quad (11)$$

with

$$c_k(\beta, \gamma) := \sum_{d=0}^n \left(\#_d(k) (e^{-i\gamma} - 1) + \binom{n}{d} \right) f_n(\beta, d) \quad (12)$$

where $\#_d(k) = |\{z \in T \mid d_H(z, k) = d\}|$ and $f_n(\beta, d) = (\cos \beta)^{n-d} (-i \sin \beta)^d$. Here, $d_H(z, k)$ is the Hamming distance between the bit strings of the binary representations of z and k .

Proof. By fixing $p = 1$ in Eq. (3), we obtain the state $|\beta, \gamma\rangle_1 = e^{-i\beta X} e^{-i\gamma C} |+\rangle^{\otimes n}$. After applying the phase separating gates to $|+\rangle^{\otimes n}$ and linearly pulling the mixer operators into the sums, we arrive at

$\frac{1}{\sqrt{2^n}} (e^{-i\gamma} \sum_{z \in T} e^{-i\beta X} |z\rangle + \sum_{z \notin T} e^{-i\beta X} |z\rangle)$. From Lemma 1, it follows that

$$|\beta, \gamma\rangle_1 = \frac{1}{\sqrt{2^n}} \left(e^{-i\gamma} \sum_{z \in T} \sum_{k \in \mathbb{F}_2^n} f(\beta, z, k) |k\rangle + \sum_{z \notin T} \sum_{k \in \mathbb{F}_2^n} f(\beta, z, k) |k\rangle \right) \quad (13)$$

By reordering the sums and factoring out $|k\rangle$ in Eq. (13), we can express the state $|\beta, \gamma\rangle_1$ as $\frac{1}{\sqrt{2^n}} \sum_{k \in \mathbb{F}_2^n} c_k(\beta, \gamma) |k\rangle$ with $c_k(\beta, \gamma) = e^{-i\gamma} \sum_{z \in T} f(\beta, z, k) + \sum_{z \notin T} f(\beta, z, k)$ (note that we omit parameters (β, γ) on c_k and other quantities below when the dependency is clear from the context). If we apply C to $|\beta, \gamma\rangle$ we basically filter out all $|k\rangle \notin T$ by linearly pulling C into the sum, therefore obtaining $C|\beta, \gamma\rangle = \frac{1}{\sqrt{2^n}} \sum_{k \in T} c_k(\beta, \gamma) |k\rangle$. We conclude that

$$\begin{aligned} F_1(\beta, \gamma) &= \langle \beta, \gamma |_1 C | \beta, \gamma \rangle_1 \\ &= \frac{1}{2^n} \left(\sum_{k \in \mathbb{F}_2^n} \overline{c_k(\beta, \gamma)} \langle k | \right) \left(\sum_{k \in T} c_k(\beta, \gamma) |k\rangle \right) \end{aligned} \quad (14)$$


$$= \frac{1}{2^n} \sum_{k \in T} \overline{c_k(\beta, \gamma)} c_k(\beta, \gamma) \quad (15)$$

$$= \frac{1}{2^n} \sum_{k \in T} |c_k(\beta, \gamma)|^2. \quad (16)$$

The step from Eq. (14) to Eq. (16) follows since all pairs in $\{|k\rangle \mid k \in \mathbb{F}_2^n\} \times \{|k\rangle \mid k \in T\}$ comprise orthogonal states. Note that $c_k(\beta, \gamma)$ essentially contains a sum over all $z \in \mathbb{F}_2^n$ with an optional phase factor of $e^{-i\gamma}$ if $z \in T$. Thus, we find that

$$\begin{aligned} c_k(\beta, \gamma) &= \sum_{z \in \mathbb{F}_2^n} \chi_T(z) e^{-i\gamma} f(\beta, z, k) + \\ &\quad (1 - \chi_T(z)) f(\beta, z, k). \end{aligned} \quad (17)$$

Here $\chi_T(z) = 1$ iff $z \in T$ and otherwise $\chi_T(z) = 0$. By recalling the definition of $f(\beta, z, k)$ in Eq. (9), we recognise $(\cos \beta)^{n-d} (-i \sin \beta)^d$ as the basic shape of f , where d is the Hamming distance between two concrete states $k, z \in \mathbb{F}_2^n$. This means that for a fixed angle β_c , $f(\beta_c, k, z)$ is actually a function of the Hamming distance between two states. Therefore, while the sum in Eq. (17) iterates over all 2^n states z , there are effectively only $n+1$ different basic terms in the sum—one for each possible distance d . Thus, if we count the number of occurrences $\#_d(k) = |\{z \in T \mid d_H(z, k) = d\}|$ of each distance d between k and other states in the target set, we can reorder the sum while still capturing



(a) Lemma 2 relates QAOA ($p = 1$) landscapes with the Hamming distance structures of instance solution spaces by counting equally distanced states in $\#_d(k) = |\{z \in T \mid d_H(z, k) = d\}|$. We can visualize this by drawing the graph of all solutions k in the solution space T , where two solutions $k, z \in T$ are connected by an edge iff they have hamming distance $d_H(z, k) = 1$, i.e. you can reach one from the other by flipping just one bit. Now, to visualize $\#_d(k)$ we place k in the center and order all $z \in T$ in orbits around k , where z is in the i -th orbit iff $d_H(z, k) = i$. Then, $\#_d(k)$ is the number of vertices on the d -th orbit.

 Derivation of $\#_d(k)$

$$\begin{aligned}
 \overline{\#_{d=2}} &= \overline{\left\{ \#_2(001010111100), \#_2(010011000110), \#_2(010001000110), \#_2(010101000110), \right. \\
 &\quad \left. \#_2(010111000110), \#_2(010001000111), \#_2(010101000111), \#_2(010111000010) \right\}} \\
 &= \overline{\left\{ \# \bullet \left(\text{diagram 1} \right), \# \bullet \left(\text{diagram 2} \right), \# \bullet \left(\text{diagram 3} \right), \# \bullet \left(\text{diagram 4} \right), \right. \\
 &\quad \left. \# \bullet \left(\text{diagram 5} \right), \# \bullet \left(\text{diagram 6} \right), \# \bullet \left(\text{diagram 7} \right), \# \bullet \left(\text{diagram 8} \right) \right\}} \\
 &= \overline{\{1, 5, 8, 6, 6, 6, 5, 4\}} = 5.125
 \end{aligned}$$

(b) By looking at the upper visualisation of solutions of a randomly generated SAT formula (see Section 5.3), the global similarities between equal orbits across different instances become apparent. This is the key motivation for Theorem 1, where global problem structures are captured by averaging $\#_d(k)$ over all $k \in T$. Making use of this visualisation, we can get a graphical intuition on how to derive $\overline{(\#_d)}(k)$, as depicted for $d = 2$ and the 8 solutions plotted above.

Figure 3: A central part of our approximation theorem is the idea of grouping equal Hamming distances. How this captures structural information of a problem solution space is visualized in (a). Mathematically this concept is expressed by a counting function $\#_d$, its derivation is depicted in (b).

all possible bit-flip mappings from z to k :

$$c_k(\beta, \gamma) = \sum_{d=0}^n \#_d(k) e^{-i\gamma} f_n(\beta, d) + \left(\binom{n}{d} - \#_d(k) \right) f_n(\beta, d)$$

with

$$f_n(\beta, d) := (\cos \beta)^{n-d} (-i \sin \beta)^d. \quad (18)$$

Now it is simply a matter of factoring out f_n and $\#_d$ to arrive at Eq. (12). $\square$

4 The QAOA Approximation Theorem

We now state and prove our main result.

4.1 Intuition

A straightforward description of an optimisation landscape might be directly derived by considering all point to point interactions between all states in superposition. Obviously, there are exponentially many such state interactions to consider. In Section 3, we captured those effects in $f(\beta, k, z)$ (see Eq. (9)) and showed that there are only $n + 1$ different outcomes that can occur based on the Hamming distance between two states. Now we use this reduction in complexity of the effect domain to derive an approximation theorem for the *expected optimisation landscape* $E(F_1)$ of a specific problem.

4.2 Formalisation

In Section 3, we presented a closed form of the optimisation landscape based on the Hamming distances between states. The components $|c_k(\beta, \gamma)|^2$ in Eq. (12) play a crucial role in $F_1(\beta, \gamma)$ (see Eq. (11)). As they are functions of the optimisation angles, we find it prudent to first obtain an intuitive visual understanding of their interrelationship. Each contribution $|c_k(\beta, \gamma)|^2$ can be seen as a landscape for a specific value of k itself. Consider Fig. 4 that depicts, as an illustration, vertical slices of three $|c_k(\beta, \gamma)|^2$ for $k \in T = \{10, 13, 14\}$ and dimension $n = 5$. The macroscopic similarities between the three are visually obvious, which motivate the idea to take the mean over all $|c_k(\beta, \gamma)|^2$. To further quantify this idea, we look at the overlaid cross sections for all three components compared with their mean in Fig. 5. Here, we evaluate $|c_k(\beta, \gamma)|^2$ for all $k \in T = \{10, 13, 14\}$ across $0 \leq \beta \leq \pi$ at a fixed $\gamma_c = 1.2$ and calculate the mean of all $|c_k(\beta, \gamma_c)|^2$ at each point. The chosen constant 1.2 does not have any particular computational or physical significance, but results in a “typical” two-dimensional sub-space of the optimisation landscape. We again observe how the mean value

nicely captures the higher level structures shared between all individual components. In fact, this intuition can be expressed as a precise mathematical relationship, as F_1 is nothing more than this mean value of c_k terms scaled by the relative cardinality $\frac{|T|}{2^n}$ of the target space T , since

$$\begin{aligned} \frac{|T|}{2^n} \overline{|c_k(\beta, \gamma)|^2} &= \frac{|T|}{2^n} \frac{1}{|T|} \sum_{k \in T} |c_k|^2 \\ &= \frac{1}{2^n} \sum_{k \in T} |c_k|^2 = F_1(\beta, \gamma). \end{aligned} \quad (19)$$

As remarked above, the global structure of F_1 is very similar across different instances of a given problem. This renders the *expected value* of an optimisation landscape for a problem an interesting object of study, as it captures the salient properties *across* instances. In particular, we argue that the ability to efficiently approximate is well suited to improve the understanding of QAOA, as we show in Section 5, and also leads to practical implications that can help optimise the use of QAOA, as we detail in Section 6.

Before that, and based on these observations, let us however first state and prove our main theorem for efficiently approximating the expected value of $F_1(\beta, \gamma)$ across instances of a computational problem. Recall that the specific optimisation Landscape of a single instance can be expressed as the mean of its $|c_k(\beta, \gamma)|^2$ weighted by the relative cardinality of its solution space compared to the full state space ($|T|2^{-n}$). Now we show, that we can approximate the expected landscape of a random problem instance by calculating the expected values of these two parameters: $E(\overline{|c_k(\beta, \gamma)|^2})$ and $E(|T|)$.

Theorem 1. *Let $\#_d(k)$ be defined as in Lemma 2, then we can approximate the expected value of $F_1(\beta, \gamma)$ for a random instance of a problem by*

$$\tilde{E}(F_1(\beta, \gamma)) = \frac{E(|T|)}{2^n} E(\overline{|c_k(\beta, \gamma)|^2})$$

with

$$E(\overline{|c_k(\beta, \gamma)|^2}) = \sum_{d_1, d_2=0}^n w_{d_1, d_2}(\gamma) f_n(\beta, d_1) f_n(\beta, d_2)^*, \quad (20)$$

where $w_{d_1, d_2}(\gamma)$ is

$$\begin{aligned} w_{d_1, d_2}(\gamma) &= E(\overline{\#_{d_1}(k) \#_{d_2}(k)}) (e^{-i\gamma} - 1) (e^{i\gamma} - 1) + \\ &\quad E(\overline{\#_{d_1}(k)}) (e^{-i\gamma} - 1) \binom{n}{d_2} + \\ &\quad E(\overline{\#_{d_2}(k)}) (e^{i\gamma} - 1) \binom{n}{d_1} + \binom{n}{d_1} \binom{n}{d_2}. \end{aligned}$$

We calculate mean values over all k in the instance specific target set T . The absolute approximation error is bounded by $|E(F_1(\beta, \gamma)) - \tilde{E}(F_1(\beta, \gamma))| \leq \sqrt{\text{Var}(|T|) \text{Var}(\overline{|c_k|^2})}$.




Figure 4: Three landscape components $|c_k(\beta, \gamma)|^2$ for $k \in \{10, 13, 14\} = T$ and a state space of dimension $n = 5$ are depicted in this order from left to right. The landscapes are represented by an array of vertical cross sections along β .




Figure 5: Cross section of the optimisation landscape components $|c_k(\beta, \gamma_c)|^2$ depicted in Fig. 4 at $\gamma_c = 1.2$ for $k \in \{10, 13, 14\} \subset T$ (note matching colours). The dashed line shows $E(|c_k(\beta, \gamma_c)|^2)$ for these three components, and illustrates how the mean value captures the globally relevant features of instance-specific information.

Proof. We start by reformulating $|c_k(\beta, \gamma)|^2$. Recall that $|c_k(\beta, \gamma)|^2 = c_k(\beta, \gamma)c_k(\beta, \gamma)^*$ and $c_k(\beta, \gamma) = \sum_{d=0}^n (\#_d(k)(e^{-i\gamma} - 1) + \binom{n}{d}) f_n(\beta, d)$. Then we use the distributivity of the complex conjugate over addition and multiplication to pull it into the sum:

$$\begin{aligned} c_k(\beta, \gamma)^* &= \sum_{d=0}^n \left(\#_d(k) \left((e^{-i\gamma})^* - 1 \right) + \binom{n}{d} \right) f_n(\beta, d)^* \\ &= \sum_{d=0}^n \left(\#_d(k) (e^{i\gamma} - 1) + \binom{n}{d} \right) f_n(\beta, d)^* \end{aligned}$$

where the second equality follows from $(e^z)^* = e^{z^*}$.

From this we see, that

$$\begin{aligned} |c_k(\beta, \gamma)|^2 &= \left(\sum_{d_1=0}^n \left(\#_{d_1}(k)(e^{-i\gamma} - 1) + \binom{n}{d_1} \right) f_n(\beta, d_1) \right) \\ &\quad \left(\sum_{d_2=0}^n \left(\#_{d_2}(k)(e^{i\gamma} - 1) + \binom{n}{d_2} \right) f_n(\beta, d_2)^* \right) \end{aligned}$$

and further

$$\begin{aligned} |c_k(\beta, \gamma)|^2 &= \sum_{d_1, d_2=0}^n \left(\binom{n}{d_1} \binom{n}{d_2} + \right. \\ &\quad \#_{d_1}(k)(e^{-i\gamma} - 1) \binom{n}{d_2} + \#_{d_2}(k)(e^{i\gamma} - 1) \binom{n}{d_1} + \\ &\quad \left. \#_{d_1}(k) \#_{d_2}(k)(e^{-i\gamma} - 1)(e^{i\gamma} - 1) \right) f_n(\beta, d_1) f_n(\beta, d_2)^* \end{aligned}$$

Now from $\overline{|c_k(\beta, \gamma)|^2} = \frac{1}{|T|} \sum_{k \in T} |c_k(\beta, \gamma)|^2$ it follows by reordering the sum that

$$\begin{aligned} \overline{|c_k(\beta, \gamma)|^2} &= \sum_{d_1, d_2=0}^n \left(\binom{n}{d_1} \binom{n}{d_2} + \right. \\ &\quad \overline{\#_{d_1}(k)}(e^{-i\gamma} - 1) \binom{n}{d_2} + \overline{\#_{d_2}(k)}(e^{i\gamma} - 1) \binom{n}{d_1} + \\ &\quad \left. \overline{\#_{d_1}(k) \#_{d_2}(k)}(e^{-i\gamma} - 1)(e^{i\gamma} - 1) \right) f_n(\beta, d_1) f_n(\beta, d_2)^* \end{aligned}$$

where $\overline{\#_d(k)} = \frac{1}{|T|} \sum_{k \in T} \#_d(k)$. Finally, Eq. (20) follows from the linearity of the expected value. From Eq. (19) we deduce that $E(F_1) = E\left(\frac{|T|}{2^n} |c_k|^2\right) = \frac{E(|T|)}{2^n} E(|c_k|^2) + \text{Cov}(|T|, |c_k|^2)$. Therefore we conclude that $|\tilde{E}(F_1) - E(F_1)| = |\text{Cov}(|T|, |c_k|^2)| \leq \sqrt{\text{Var}(|T|) \text{Var}(|c_k|^2)}$. $\square$

4.3 Generality

The approximation theorem only operates on the Hamming distance structure of problem target spaces. Recall, however, that we arrive at the target space T by defining it through an isomorphism from the set of all solutions $\{z \in \mathbb{F}_2^n \mid c(z) = 1\}$ to a subset of quantum basis states $T = \{|z\rangle \mid c(z) = 1\}$, let that be $f : z \mapsto |z\rangle$. A more complex constraint Hamiltonian might also need to operate on ancilla qubits that are not part of the target space, possibly embedding f in a structure like $g : xz \mapsto |h(x)\rangle \otimes |z\rangle$. This additional degree of freedom allows for inconsistent side effects on ancilla registers while still maintaining functional correctness. As a result, this could thus alter the hamming distance between pairs of states encoding the same solutions. As a consequence the approximation theorem would require a mechanism to take computations on arbitrary ancilla registers into account. We would thus lose the elegance of being independent of concrete Hamiltonian constructions. We now show that for every problem in **NP**, it is possible to efficiently construct an ancilla register independent constraint Hamiltonian. This guarantees a wide applicability of Theorem 1 only focusing on the concrete solution structure of the problem at hand. For the following theorem and its proof we use verifier definition of **NP**, that is: A problem is in **NP** if and only if there exists an algorithm $v : \mathbb{F}_2^n \rightarrow \mathbb{F}_2$ that efficiently verifies solutions $z \in \mathbb{F}_2^n$, i.e. $v(z) = 1$ iff z is a valid solution to the problem and otherwise $v(z) = 0$.

Theorem 2. *For every problem in **NP** there exists a family of constraint two-level Hamiltonians $\{C_n\}_{n \in \mathbb{N}}$ that can be efficiently constructed such that*

$$C_n|z\rangle|0\rangle = \begin{cases} |z\rangle|0\rangle & z \text{ proves the decision property} \\ 0 & \text{otherwise} \end{cases} \quad (21)$$

for proofs z of length n .

Proof. By definition a problem is in **NP** if and only if there exist an efficient verifier v that returns $v(z) = 1$ on all valid proofs z of the decision property and $v(z) = 0$ on all other inputs. Thus, there also exists an efficiently constructible family of quantum circuits $\{V_n\}_{n \in \mathbb{N}} : V_n|z\rangle|0\rangle = |z\rangle|v\rangle$ with $v = (v(z), v_2, \dots, v_{p(n)})$, where $|0\rangle = |0\rangle^{\otimes p(n)}$ is an ancilla register of size $p(n)$ for some polynomial p . Now, the construction of C_n looks as follows: $C_n := V_n^\dagger(\mathbb{1}^{\otimes n} \otimes |1\rangle\langle 1| \otimes \mathbb{1}^{p(n)-1})V_n$. $\square$

Remark 3. *Theorem 2 ensures us that if a problem is in **NP**, we only have to consider its target space T of problems when applying Theorem 1. Remember that T is isomorphic to the solution or proof space of a problem. In essence, Theorem 1 addresses structural properties of the expected solution space of a problem. So, the combination of Theorem 1 and Theorem 2 shows the existence of macroscopic structures on a*

*problem level that significantly influence the optimisation landscapes induced by QAOA circuits for at least all **NP** problems.*

On a more practical note, this result means that we don't have to take computational side effects of concrete circuit implementations into account, when using Theorem 1 to analyse the effect of problem structures on QAOA. Let's say for example, one wants to solve a k -CLIQUE problem, which is the corresponding decision problem to the max-CLIQUE optimisation problem. Then we can construct a QAOA circuit using Theorem 2 to construct the constraint Hamiltonian needed for the problem layer. The resulting QAOA circuit will then solve the k -CLIQUE problem by preparing a superposition $\sum_i \alpha_i |z_i\rangle|0\rangle$ of states where z_i representing cliques of size k . For a detailed description of exactly this example see Section 5.4.

4.3.1 A Note about **NP**

We now want to argue why covering **NP** is practically sufficient to prove generality for our framework. We acknowledge that in the quantum computing community usually optimization problems like for example the ever present Max-Cut problem are the subject of investigation. Thus, it might seem unintuitive to base a framework on a decision problem class like **NP**. In fact, most problems relevant in practice are in some variant part of **NP**, see the seminal list of **NP** problems [51] by Karp. Indeed, foundational work by Lucas that arguably significantly contributed to the commercialisation of quantum computing [57] provided Ising formulations of many important **NP** problems, bridging the gap between classical industrial problems solving usually powered by highly sophisticated SAT-solvers [46], which are basically **NP** solvers, and quantum computing at that time.

Coming back now to the other perspective, starting with an optimisation problem, how do we apply our framework. Let's consider the Max-Cut problem for instance. Instead of asking for an optimal cut, we reformulate the problem to ask whether there exists a cut above a certain threshold $k \in \mathbb{N}$. We can as proven above for all optimization problems which have a threshold variant in **NP**. In practice this applies many interesting problems like Max-Cut, Partition, Vertex Cover, etc. [51]. Then we formulate a Hamiltonian that *verifies* states encoding a cut satisfying the threshold and apply our framework to it. Compared to the classical SAT-solver case, the Hamiltonian plays the role of the stat formula and our QAOA approach takes the place of the SAT-solver.

5 Application

After having set the foundations for a methodology to understand structural properties of optimisation problems across instances, let us now commence with

applying the framework to several concrete examples. We consider five different subject problems respectively scenarios (uniform random sampling, clustered sampling, Boolean satisfiability, k-clique, and one-way functions in the form of qr factoring) that build on each other to best introduce the application of the QAOA approximation theorem on real problems. Overall, the selection of examples is carefully curated to highlight different aspects of using our approximation theorem: We first compare the two possible approaches of either analytically or empirically analysing the target space structure of a problem at hand. Then we demonstrate with a purposefully constructed sampling method that the existence of a stochastic dependence of two states with a certain Hamming distance in a random instance significantly influences the optimisation landscape. Following that, we use SAT as a first straight forward real world example with an easy to construct constraint Hamiltonian. After that, we show that for all problems in **NP**, even with more complicated constraints, there is a construction for an appropriate constraint Hamiltonian that satisfies the preconditions of our approximation theorem. As a final example, we highlight the case of integer factoring to argue that problems based on one-way functions are interesting subjects for our methods as their target space can be relatively easily characterised.

5.1 Uniform Random Sampling

Central in our theory is the target space T of the constraint projector C . Given a concrete interpretation, every state in T corresponds to a problem solution.¹ Thus, our notion of target spaces T is an abstraction of concrete problem solution spaces. Sampling a random target space T equals sampling a random problem minus the abstraction of a target state interpretation. Furthermore, the description of a sampling procedure for T defines an abstract random problem, where the problem itself becomes a random variable in the probabilistic point of view. To provide a smooth onramp to more complex examples further below, we start by simply sampling T by randomly drawing states from the state space with uniform probability.

From Theorem 1, we see that $E(F_1)$ can be efficiently approximated if the quantities (a) $|T|$, (b) $E(\overline{\#_d(k)})$ for all $0 \leq d \leq n$, and (c) $\text{Cov}(\overline{\#_{d_1}(k)}, \overline{\#_{d_2}(k)})$ for all $0 \leq d_1, d_2 \leq n$ are provided. Recall that the expected value is calculated over all instances of a specific problem, thus $\overline{\#_d(k)}$ are random variables describing the mean value $\frac{1}{|T|} \sum_{k \in T} \#_d(k)$ over the target set T of a random problem instance. There are, in general, two approaches: We can, if possible, determine the distribution

of the random variables $\overline{\#_d(k)}$ by analytical means, or by an empirical numerical approach. For uniform sampling as considered in this motivating example, it is relatively straightforward to execute the necessary analytic calculations.

5.1.1 Analytic Approach

Definition 4. We consider an urn model with balls of $m \in \mathbb{N}$ different colours. The population of all balls in the urn is described by $\mathbf{X} = (X_1, X_2, \dots, X_m)$, where X_i is the number of balls with colour i . Then, if n balls are drawn without replacement, the probability of having x_i balls of colour i in the sample is given by the multivariate hypergeometric distribution $P(\mathbf{x}) = \text{HypG}(\mathbf{X}, \mathbf{x})$ with $\mathbf{x} = (x_1, x_2, \dots, x_m)$. Recall that, by textbook knowledge, the hypergeometric distribution is defined by $P(\mathbf{x}) = \prod_{i=1}^m \binom{X_i}{x_i} / \binom{N}{n}$, with $N := \sum_{i=1}^m X_i$.

The size of T is trivially given, as we always sample a fixed number of states. In the following discussion we will encounter expected values over different sample spaces, we will differentiate this by $E_T(\cdot)$ and $E_{\mathcal{I}}(\cdot)$ being defined to be the mean over all states in the target space T and over all instances in the set of problem instances $\mathcal{I}$. No subscript is used if the sample space is clear from context. Further note that $\overline{\#_d(k)}$ is the sample mean of T where T is sampled from the sample space of all problem instances $\mathcal{I}$. Therefore, $E_{\mathcal{I}}(\overline{\#_d(k)}) = E_T(\overline{\#_d(k)})$ as the sample mean $\overline{\#_d(k)}$ is an unbiased estimator of the expected value $E_T(\overline{\#_d(k)})$. The same also holds for $\overline{\#_{d_1}(k)}\overline{\#_{d_2}(k)}$. To calculate $E_T(\overline{\#_d(k)})$, we need to know the distributions of $\{\#_d(k)\}_{d=0}^n$. To compute $E_{\mathcal{I}}(\overline{\#_{d_1}(k)}\overline{\#_{d_2}(k)}) = E_T(\overline{\#_{d_1}(k)}\overline{\#_{d_2}(k)})$ we use that $E_T(\overline{\#_{d_1}(k)}\overline{\#_{d_2}(k)}) = E_T(\overline{\#_{d_1}(k)})E_T(\overline{\#_{d_2}(k)}) + \text{Cov}(\overline{\#_{d_1}(k)}, \overline{\#_{d_2}(k)})$. For $\text{Cov}(\overline{\#_{d_1}(k)}, \overline{\#_{d_2}(k)})$, the joint probability distribution of $\{(\#_{d_1}(k), \#_{d_2}(k))\}_{d_1, d_2=0}^n$ is required.

As T is uniformly sampled from the complete state space, this basically leads to an urn model without replacement. Therefore, except for some edge cases, the random variables $\#_d(k)$ can be described by a Hypergeometric distribution, see Lemma 3.

¹Note that given different interpretations, a state in T can encode different solutions of different problems, too.

Lemma 3. In case of uniform target sampling the probabilities $P(\#_d(k) = x)$ and $P(\#_{d_1}(k) = x_1, \#_{d_2}(k) = x_2)$ are defined as follows:

$$P(\#_d(k) = x) = \begin{cases} 1 & d = 0 \wedge x = 1 \\ 0 & d = 0 \wedge x \neq 1 \\ \text{HypG}(\mathbf{X}, \mathbf{x}) & \text{otherwise} \end{cases} \quad (22)$$

with $\mathbf{X} = \left(\binom{n}{d}, 2^n - 1 - \binom{n}{d}\right)$ and $\mathbf{x} = (x, |T| - 1 - x)$. Furthermore,

$$P(\#_{d_1}(k) = x_1, \#_{d_2}(k) = x_2) = \begin{cases} P(\#_{d_1}(k) = x_1)P(\#_{d_2}(k) = x_2) & d_1 = 0 \vee d_2 = 0 \\ P(\#_{d_1}(k) = x_1) & d_1 = d_2 \wedge x_1 = x_2 \\ 0 & d_1 = d_2 \wedge x_1 \neq x_2 \\ \text{HypG}(\mathbf{X}, \mathbf{x}) & \text{otherwise} \end{cases} \quad (23)$$

with $\mathbf{X} = \left(\binom{n}{d_1}, \binom{n}{d_2}, 2^n - 1 - \binom{n}{d_1} - \binom{n}{d_2}\right)$ and $\mathbf{x} = (x_1, x_2, |T| - 1 - x_1 - x_2)$.

Proof. This proof is structured in two parts. We start by showing the correctness of Eq. (22), which in turn necessitates distinguishing between two cases. As $\#_d(k)$ describes the distance relationship of a random state in T to all other states in T (including itself), $\#_0(k)$ is always 1, from which the first two cases of Eq. (22) follow. For $d \geq 0$, we need to consider the sampling process of T . Assume we start with just one random state in T , and then sample the rest. This is described by an urn model with $X_1 = \binom{n}{d}$ balls of colour d , and $X_2 = 2^n - 1 - \binom{n}{d}$ other balls. Note that by picking a random reference state in T , there are only $2^n - 1$ balls left in the urn to draw the $|T| - 1$ states left to fill the target space. Therefore, the probability of sampling x states of Hamming distance d from the random initial state is given by

$$P(\#_d(k) = x) = \text{HypG}((X_1, X_2), (k, |T| - 1 - x)),$$

which concludes the proof of Eq. (22).

We need to consider three cases for Eq. (23). If $d_1 = 0$, the only possible samples (x_1, x_2) must have $x_1 = 1$, with the same reasoning as above. Thus, $P(\#_{d_1}(k) = x_1, \#_{d_2}(k) = x_2) = \delta_{1,x_1}P(\#_{d_2}(k) = x_2)$, where $\delta_{i,j}$ is the Kronecker delta function. Further note that $\delta_{1,x_1} = P(\#_{d_1}(k) = x_1)$. The case of a vanishing second argument, $d_2 = 0$, is symmetric to $d_1 = 0$. Therefore, it follows that $P(\#_{d_1}(k) = x_1, \#_{d_2}(k) = x_2) = P(\#_{d_1}(k) = x_1)P(\#_{d_2}(k) = x_2)$ iff $d_1 = 0$ or $d_2 = 0$. If $d_1 = d_2$, then obviously $x_1 = x_2$ needs to be satisfied, in which case $P(\#_{d_1}(k) = x_1, \#_{d_2}(k) = x_2)$ reduces to $P(\#_{d_1}(k) = x_1) = P(\#_{d_2}(k) = x_2)$. For all other d_1 and d_2 , re-consider the sampling process from an urn without replacement: This time, we are interested in two colours d_1 and d_2 . Consequently, the urn contains $\binom{n}{d_1}$ balls of colour d_1 , $\binom{n}{d_2}$ balls of colour d_2 , and $2^n - 1 - \binom{n}{d_1} - \binom{n}{d_2}$ other balls. The probability of obtaining k_1 states of Hamming distance d_1 and k_2 states of Hamming distance d_2 in a sample of size $|T|$

is therefore given by

$$P(\#_{d_1}(k) = x_1, \#_{d_2}(k) = x_2) = \text{HypG}(\mathbf{X}, \mathbf{x})$$

with $\mathbf{X} = \left(\binom{n}{d_1}, \binom{n}{d_2}, 2^n - 1 - \binom{n}{d_1} - \binom{n}{d_2}\right)$ and $\mathbf{x} = (x_1, x_2, |T| - 1 - x_1 - x_2)$. $\square$

Now that we know the distribution of our random variables $\#_d(k)$, we can deduce properties like expected values and covariances. In Lemma 4, we use Lemma 3 to determine $E_T(\#_d(k))$ and $\text{Cov}(\#_{d_1}(k), \#_{d_2}(k))$.

Lemma 4.

$$E_T(\#_d(k)) = \begin{cases} 1 & d = 0 \\ |T| \frac{\binom{n}{d}}{2^n} & \text{otherwise} \end{cases} \quad (24)$$

$$\text{Cov}(\#_{d_1}(k), \#_{d_2}(k)) =$$

$$\begin{cases} 0 & d_1 = 0 \vee d_2 = 0 \\ |T| \frac{2^n - |T|}{2^{n-1}} \frac{\binom{n}{d}}{2^n} \left(1 - \frac{\binom{n}{d}}{2^n}\right) & d_1 = d_2 = d > 0 \\ -|T| \frac{2^n - |T|}{2^{n-1}} \frac{\binom{n}{d_1} \binom{n}{d_2}}{2^{2n}} & \text{otherwise} \end{cases} \quad (25)$$

Proof. In large, this proof follows the structure of the proof of Lemma 3. We start by show the correctness of Eq. (24). If $d = 0$, we trivially have $E(\#_0) = 0P(\#_0 = 0) + 1P(\#_0 = 1) = 1$. For all $d > 0$, $P(\#_d(k) = x)$ is described by the Hypergeometric distribution as in the second case of Eq. (22), with an expected value of $|T| \frac{\binom{n}{d}}{2^n} 2^{-n}$ as required.

We now prove Eq. (25). Assume that $d_1 = 0$, then with $\#_{d_1 d_2}(k) := \#_{d_1}(k) \#_{d_2}(k)$ we have

$$\begin{aligned} E(\#_{d_1 d_2}(k)) &= \sum_k 0 \cdot \#_{d_2}(k) P(\#_{d_1}(k) = 0, \#_{d_2}(k) = x) + \\ &\quad \sum_k 1 \cdot \#_{d_2}(k) P(\#_{d_1}(k) = 1, \#_{d_2}(k) = x) \\ &= \sum_k \#_{d_2}(k) P(\#_{d_1}(k) = 1, \#_{d_2}(k) = x) \end{aligned}$$

and of course $P(\#_{d_1}(k) = 1) = 1$ for $d_1 = 0$ and thus $E(\#_{d_1}(k)\#_{d_2}(k)) = \sum_k \#_{d_2}(k)P(\#_{d_2}(k) = k) = E(\#_{d_2}(k))$. Together with Eq. (24), it follows that

$$\begin{aligned} \text{Cov}(\#_{d_1}(k), \#_{d_2}(k)) &= E(\#_{d_1}(k)\#_{d_2}(k)) - \\ &\quad E(\#_{d_1}(k))E(\#_{d_2}(k)) \\ &= E(\#_{d_2}(k)) - E(\#_{d_2}(k)) = 0. \end{aligned}$$

The case for $d_2 = 0$ is symmetric. We conclude that $\text{Cov}(\#_{d_1}(k), \#_{d_2}(k)) = 0$ iff $d_1 = 0$ or $d_2 = 0$. Note that $\text{Cov}(\#_d(k), \#_d(k)) = \text{Var}(\#_d(k))$, which is given by

$$\text{Var}(\#_d(k)) = |T| \frac{2^n - |T|}{2^n - 1} \frac{\binom{n}{d}}{2^n} \left(1 - \frac{\binom{n}{d}}{2^n} \right)$$

since the sampling process is described by a multivariate hypergeometric distribution as shown in Lemma 3. The third and second case also directly follow from the covariance of the multivariate hypergeometric distribution from Lemma 3. $\square$

With the two lemmas above, we can efficiently calculate $\tilde{E}(F_1)$. This demonstrates a use case that benefits from our approximation theorem: Given a problem for which the target space structure with respect to Hamming distances between targets can be modelled analytically, the expected optimisation landscape can be approximately derived solely from this model. While instance specific analytic formulations of F_1 are known for the Ising model problem [65] from a physics-centric point of view, our approach accesses the problem from a previously unexplored angle that benefits from insights from theoretical computer science. We also conclude that an efficient description of such models allows for an efficient approximation of a QAOA optimisation landscape. Further investigations into the complexity theoretical implications of this could provide insights into the link between the structure target spaces, the complexity of problems and the computational power of QAOA, albeit we need to leave these questions to future research, as our primary goal in this paper is to establish the framework and derive direct practical utility.

5.1.2 Sampling Approach

Approximating $E(F_1)$ with an efficient theoretical model of the target space is obviously the preferred way to address a given problem, if feasible analytically. However, it is also possible to determine $\left\{ E(\overline{\#_d(k)}) \right\}_{d=0}^n$ and $\left\{ E(\overline{\#_{d_1}(k)\#_{d_2}(k)}) \right\}_{d_1, d_2=0}^n$ empirically by sampling from a set of problem instances. Above, we described the problem instances of the uniform sampling example by giving a sampling routine for a random target space. Recall that this defines problem instances as random variables. If we now want to determine the expected values and covariance matrix empirically, we need a concrete sample of

instances. Which in this example will be a set of realisations of problem random variables.

In our experiments, we sampled 500 random instances for each state space dimension from $8 \leq n \leq 11$. The upper bound of $n = 11$ keeps computational cost of explicitly evaluating F_1 at a reasonable level, whereas the lower bound $n = 8$ ensures a sufficiently large state space. For the scope of this paper, this range is sufficient to show the effects of increasing the state space dimension. Every instance has a target set of 2^{n-1} states, making it scale with the size of the whole state space. As defined above, $\overline{\#_d(k)} = \frac{1}{|T|} \sum_{k \in T} \#_d(k)$. Let $\mathcal{T}$ be the set of target sets of all sampled instances. Then, empirically determining $E(\overline{\#_d(k)})$ means calculating the mean of $\overline{\#_d(k)}$ over the set $\mathcal{T}$ of all sampled target sets. To clarify over which set a mean is calculated, we introduce the notation $\overline{f(x)}^A = \frac{1}{|A|} \sum_{x \in A} f(x)$. Then,

$$\overline{\overline{\#_d(k)}}^{\mathcal{T}} = \frac{1}{|\mathcal{T}|} \sum_{T \in \mathcal{T}} \frac{1}{|T|} \sum_{k \in T} \#_d(k).$$

The same holds for $E(\overline{\#_{d_1}(k)\#_{d_2}(k)})$, with

$$\overline{\overline{\#_{d_1}(k)\#_{d_2}(k)}}^{\mathcal{T}} = \frac{1}{|\mathcal{T}|} \sum_{T \in \mathcal{T}} \frac{1}{|T|} \sum_{k \in T} \#_{d_1}(k)\#_{d_2}(k).$$

5.1.3 Comparison

As we can see in Fig. 6, both approaches lead to approximations that fit the expected value of F_1 exceptionally well. In all our experiments we used the sample mean (over all instances) as an estimator for the expected landscape $E(F_1)$. We calculated the mean of F_1 by evaluating $F_1(\gamma, \beta)$ for every instance at 100 sample points for $0 \leq \beta \leq \pi$ at the cross section for $\gamma = 1.2$. This shows the versatility of our approximation: Theoretical models can be used to get excellent results. This is not always feasible for more complex problems where the inherent structural properties of the target space are not as straightforward. However, we also showed that an empirical approach leads to equally excellent results in such cases. We therefore argue that our approximation theorem can be a useful tool that is equally well suited for theoretical considerations and empirical analyses.

5.2 Clustered Sampling

The uniform sampling example was chosen deliberately trivial to showcase the analytical approach. As a consequence, one important aspect of more realistic problems remains unaddressed: Owing to the uniform sampling process, the states in T are independent of each other. This is, in general, not the




Figure 6: Comparison of analytic and sampling-based landscape approximation for the uniform sampling experiment. *Left:* Landscapes for sampling methods. The difference between the analytically modelled and the empirically sampled landscape approximations is negligibly small. *Top Right:* A two-dimensional heat-map $\blacksquare$ (grey) that illustrates the vertical distribution of $F_1(\beta, \gamma_0 = 1.2)$ values of 500 random instances sampled as described in Section 5.1.2. The heat-map is calculated over a grid of 100×100 bins, where each column collects all F_1 values in the corresponding β interval and sums up to 1. *Bottom Right:* Cross-section through the above landscape at $\gamma_c = 1.2$. The expected landscape estimated by the mean value of the data points displayed in the heat map above $\blacksquare$ (black line) and two approximations of F_1 , where $\{\overline{\#_d(k)}\}_{d=0}^n$ and $\{(\overline{\#_{d_1}(k)}, \overline{\#_{d_2}(k)})\}_{d_1, d_2=0}^n$ are determined empirically --- (orange line) and modelled theoretically --- (green line) are shown. Again, exact and approximate results are in excellent agreement, with little variation.




Figure 7: Landscape cross-sections and approximation quality for clustered sampling for different state space dimensions in columns. *Top*: Two-dimensional heat-map of sampled values for $F_1(\beta, \gamma_c)$ at $\gamma_c = 1.2$ and across $0 \leq \beta \leq \pi$, overlaid with the expected landscape estimated by the mean value of the sampled data points $\bar{F}_1$ (black line) and the approximated expectation $\hat{E}(F_1)$ (orange line). *Bottom*: Absolute approximation error (solid line) and standard deviation of F_1 (dashed line). The data demonstrate a small approximation error that lies strictly and considerably below the standard deviation of F_1 .

case for real-world problems. SAT instances, for example, often have clustered solution spaces. Industrial SAT instances exhibit community-like structures in their solution spaces [8]. But contrary to intuition, even uniformly sampled SAT instances possess solution structures [68]. Inspired by this observation, we conduct further experiments with a random clustered sampling process: Three cluster seeds are sampled uniform at random from the complete state space. Then, per seed, 30 new states are added to the target set. Each state is reached by a random walk starting from its corresponding seed by flipping a random bit each step, with a step probability of $\frac{1}{2}$. Note that reaching a basis state $|k\rangle$ from another basis state $|z\rangle$, with $k, z \in \mathbb{F}_2^n$ means that qubit l was flipped if and only if $k_l \neq z_l$. The process of empirically determining $\left\{E\left(\overline{\#_d(k)}\right)\right\}_{d=0}^n$ and $\left\{E\left(\overline{\#_{d_1}(k)\#_{d_2}(k)}\right)\right\}_{d_1, d_2=0}^n$ stays the same as in the uniform example above.

Figure 7 shows the empirical estimate of $E(F_1)$ compared to the actual mean value $\bar{F}_1$ and the introduced absolute error. As can be seen, we clearly introduce some error by approximating. But although the variance of F_1 drops significantly with higher dimensions, the absolute error never surpasses the standard deviation of F_1 . Also note how the whole co-domain of F_1 is compressed. This is a direct consequence of

the scaling coefficient $|T|/2^n$ in Eq. (11) as we chose to generate fixed dimension independent sized target sets. Recall that $E(F_1)$ describes the result probability of sampling a solution state from $|\gamma, \beta\rangle_1$, which unveils an interesting relation between the size of the solution space and the solution sample probability.

5.3 Boolean Satisfiability

After two artificially constructed sampling examples, it is time to consider our first real problem. While SAT is one of the cornerstones of **NP**-complete problems in theoretical computer science, it also has considerable practical impact [38], and has also been intensively studied in the quantum domain [36, 53, 75, 94, 99].

Most importantly, it is known to often exhibit structured solution spaces [3], and is a natural combinatorial decision problem, which makes it perfectly suited for our framework.

Note that in both above sampling examples, we only discussed the target set and completely disregarded any actual realisation of the QAOA circuit in our analysis. However, if the complete target set were known, it would be trivial to construct a constraint Hamiltonian $C = \prod_{k \in T} |k\rangle\langle k|$ by simply projecting onto every target state. This is obviously not an adequate approach for hard computational problems, as full knowledge about the solution space cannot be ex-




Figure 8: Approximate expected value $\tilde{E}(F_1)$ of $F_1(\beta, \gamma)$ for a random SAT instance perfectly fits the actual expected landscape $E(F_1)$ over all sampled instances, even though SAT induces a higher variance in the values of F_1 . This is illustrated by a cross section at $\gamma = \gamma_0 = 1.2$. The two-dimensional heatmap $\blacksquare$ (gray) illustrates the vertical distribution of $F_1(\beta, \gamma_0)$ over a grid of 100×100 bins, where each column collects all F_1 values in the corresponding β interval and sums up to 1. This histogram is overlaid by the expected Landscape $E(F_1(\beta, \gamma_0))$ $\blacksquare$ (black, solid) and the approximated mean value $\tilde{E}(F_1(\beta, \gamma_0))$ --- (orange, dashed) obtained from the approximation theorem. All three quantities are in very good agreement.

pected. Unfortunately, explicitly projecting onto every state in T violates this requirement for difficult problems. We therefore show now that we still can apply our approximation theorem even if this requirement is added to the picture.

Let $\mathcal{C} = \{c_i\}_{i=1}^m$ a Boolean formula in conjunctive normal form, with c_i being the i -th clause. The SAT problem asks whether there exists a Boolean variable assignment such that all clauses are satisfied. Each clause is a disjunction of literals. A literal is a possibly negated Boolean variable. Let's consider the clause $c = x_1 \vee x_3 \vee \overline{x_8}$. Then c is satisfied by all possible assignments except for the characteristic unsatisfying assignment $[x_1 \mapsto 0, x_3 \mapsto 0, x_8 \mapsto 1]$. Let $\overline{C}_i$ be a projector onto this characteristic unsatisfying assignment of the clause c_i . Given $\{\overline{C}_i\}_{i=0}^m$ we can construct a constraint Hamiltonian $C = \prod_{i=0}^m (\mathbb{1} - \overline{C}_i)$ that projects onto all satisfying assignments of $\mathcal{C}$. This allows us to just focus on the solution space, and set the stage for using the approximation theorem.

For our experiments, we generated 500 random satisfiable SAT instances for each $n \in \{8, 9, 10, 11\}$, where the number of variables directly translates to the dimension ($|V| = n$) and the number of clauses is always $|C| = 4n$. Each clause has three literals that are negated with probability $\frac{1}{2}$. Glucose 4.2 [12] was used to enumerate all satisfying assignments for each

instance. In contrast to the previous sampling examples, where we always sampled a target set of fixed size $|T|$, the size of T depends on the instance and its number of satisfying assignments in this scenario. Therefore, we also have to determine $|T|$ empirically.

In Fig. 8 we see once again how well our approximation fits the real mean values for F_1 . Note that compared to the examples above, there seems to be a unusual amount of variance on the y-axis. This is a result of different target set sizes $|T|$ across instances, which causes global vertical shifts of function F_1 . If one is just interested in the overall structure of F_1 one could consider ignoring the scaling factor induced by $|T|$ altogether.

5.4 k -CLIQUE

An efficient in-place projection like the SAT constraint Hamiltonian can not always be expected. So what if ancilla bits are needed? In this example, we demonstrate an approach following Theorem 2 by showing how to construct an ancilla register independent of the constraint Hamiltonian for k -CLIQUE.

Definition 5. Given a graph G and $k \in \mathbb{N}$, then the decision problem whether G has a clique (i.e., a fully connected sub-graph) of size k is called the k -CLIQUE Problem.

To solve k -CLIQUE with QAOA, we first need to define the state space. Let $G = (V, E)$ be a graph with n vertices, then we have a n -dimensional state space as we map each vertex to a specific qubit. A basis state $|z\rangle$ with $z \in \mathbb{F}_2^n$ marks vertices such that a vertex v_i is marked iff $z_i = 1$. Then, $|z\rangle$ represents a valid solution to k -CLIQUE iff $K = \{v_i \in V \mid z_i = 1\}$ is a clique in G and $|K| = k$.

Let $\overline{G}$ be the complement of G and (i, j) an edge in $\overline{G}$. Then,

$$C_{\text{cliques}} = \prod_{(i,j) \in \overline{G}} (\mathbb{1} - P_{ij})$$

with $P_{ij} = \mathbb{1}_{0,i-1} \otimes |1\rangle\langle 1| \otimes \mathbb{1}_{i+1,j-1} \otimes |1\rangle\langle 1| \otimes \mathbb{1}_{j+1,n}$ (26)

projects onto all cliques in G . However, this also includes cliques of sizes different from k , such as trivial cliques as single vertices and edges. A second step is needed to filter out the k -cliques. For this we have to define an unitary operator D_H that calculates the Hamming weight of $|z\rangle$. This is done by writing $d_H(z)$ on an ancilla register but note that $|z\rangle|y\rangle \mapsto |z\rangle|d_H(z)\rangle$ is only invertible for a fixed y (e.g., $y = 0$), which conflicts with D_H being unitary. Thus, we define $D_H : \mathcal{B}^{\otimes n} \otimes \mathcal{B}^{\otimes [\text{ld}(n)]} \rightarrow \mathcal{B}^{\otimes n} \otimes \mathcal{B}^{\otimes [\text{ld}(n)]}$ as $D_H : |z\rangle|y\rangle \mapsto |z\rangle|y + d_H(z) \bmod 2^m\rangle$ with $m := [\text{ld}(n)] + 1$. Since $\mathbb{N}/m\mathbb{N}$ forms a group under addition, there exists a unique inverse element for each $d_H(z)$ which ensures D_H to be invertible. Our construction is illustrated in Fig. 9.

With D_H we can construct a second constraint Hamiltonian $C_{d_H=k} = D_H^\dagger(\mathbb{1}^{\otimes n} \otimes |k\rangle\langle k|)D_H$ that projects onto the space spanned by $\{|z\rangle|y\rangle \mid y + d_H(z) \bmod 2^m = k\}$. Therefore, if the ancilla register is initialised to $|0\rangle$ the application of $C_{d_H=k}$ projects onto states with Hamming distance k . Applying both projectors results in

$$C := C_{d_H=k} \left(C_{\text{cliques}} \otimes \mathbb{1}^{\lceil \text{ld}(n) \rceil} \right) \quad \text{with} \\ C|z\rangle|0\rangle = \begin{cases} |z\rangle|0\rangle & z \text{ is } k\text{-clique in } G \\ 0 & \text{otherwise} \end{cases} \quad (27)$$

This results in a QAOA state $|\beta, \gamma\rangle_1 = e^{-i\beta X^{\otimes n}} e^{-i\gamma C} \frac{1}{\sqrt{2^n}} |+\rangle^{\otimes n} |0\rangle$ after the phase separation introduced by $e^{-i\gamma C}$ the state evolves to $\frac{1}{\sqrt{2^n}} e^{-i\beta X^{\otimes n}} (e^{-i\gamma} \sum_{z \in K} |z\rangle|0\rangle + \sum_{z \notin K} |z\rangle|0\rangle)$ finally, after factoring out the ancilla register and applying the mixer we end up with

$$\frac{1}{\sqrt{2^n}} \left(e^{-i\gamma} \sum_{z \in K} \sum_{x \in \mathbb{F}_2^n} f(\beta, z, x) |x\rangle + \sum_{z \notin K} \sum_{x \in \mathbb{F}_2^n} f(\beta, z, x) |x\rangle \right) \otimes |0\rangle$$

Note that this QAOA circuit is invariant on the ancilla register, which allows us to trace the ancilla qubits out. Then, $|\beta, \gamma\rangle_1$ is identical to Eq. (13). The ancilla qubits have no influence on the distribution of Hamming weights. Therefore, they can be ignored and once again we only have to reason about the problem specific solution space.

5.5 One-Way Functions

Theorem 1 is especially useful if we either have a theoretical understanding of the solution space of a problem or if the solution space can be efficiently sampled. This is the case for one-way functions: Instances can be easily generated by first picking a state from solution space and then generating the original input instance by applying the inverse of the one-way function, as it is easy to compute by definition. We will showcase this scenario with qr -FACTORING.

Definition 6. Given a value $x = qr$ with $q, r \in \mathbb{P}$, qr -FACTORING is the problem of computing q and r given x .

$$f_{qr} : \{qr \mid q, r \in \mathbb{P}\} \rightarrow \mathbb{P} \times \mathbb{P}, \quad f_{qr} : qr \mapsto (q, r) \quad (28)$$

Obviously the inverse of Eq. (28) is simply the integer multiplication $f_{qr}^{-1}(q, r) = q \cdot r$. So, given a pre-computed set of primes $\mathcal{P} \subset \mathbb{P}$, we can efficiently sample a pool of qr -FACTORING instances by $f_{qr}^{-1}(\{(q, r)_i\}_{i=1}^m)$ with $(q, r)_i$ being sampled from $\mathcal{P} \times \mathcal{P}$ at uniform random for all $0 \leq i \leq m$. In fact,

as we argued in Section 5.4, the target space T without ancillary qubits fully suffices for our approximate analysis. For qr -FACTORING, a solution (q, r) can be represented in T as follows: Let $x_{(2)}$ be the binary representation of $x \in \mathbb{N}$, and let $z \in \mathbb{F}_2^k$ with $m \leq n$. Then $p_n(z)$ is a padding of z with $n - k$ leading zeros. Now, $\{(q, r)_i\}_{i=1}^m$ is mapped to T by $(q, r) \mapsto p_n(q_{(2)}) \circ p_n(r_{(2)})$ with $n = \max\{\lceil \text{ld}(q) \rceil, \lceil \text{ld}(r) \rceil\}$, where $\circ$ denotes concatenation of two bit strings.

With this mapping we performed a series of experiments for different target space dimensions $n \in \{12, 14, 16, 18\}$, where we sampled $\{(q, r)_i\}_{i=1}^{100}$ from $\mathcal{P}_{\frac{n}{2}} \times \mathcal{P}_{\frac{n}{2}}$, with $\mathcal{P}_{\frac{n}{2}} := \{q \in \mathbb{P} \mid q < 2^{\frac{n}{2}}\}$. Again, $\left\{E(\overline{\#_d(k)})\right\}_{d=0}^n$ and $\left\{E(\overline{\#_{d_1}(k)\#_{d_2}(k)})\right\}_{d_1, d_2=0}^n$ were empirically determined as described above. The cardinality of the target set is $|T| = 2$, as $f_{qr}^{-1}(q, r) = f_{qr}^{-1}(r, q)$ and thus for each $x = qr$, solution and target space each contain exactly two elements. Figure 10 shows the approximated and actual mean of F_1 as well as the absolute approximation error for $n = 18$. Although the margin for error apparently is extremely tight for qr -FACTORING, our approximation nicely lies on top of the actual mean. Again, the absolute error remains below the standard deviation. As is likewise visible in Fig. 10, qr -FACTORING exhibits only small deviations between instances in their optimisation landscapes. But even with this small margin of error, the approximation tightly fits the real distribution, and the error is bounded by the standard deviation of F_1 .

6 Practical Utility

It is well known that the classical task of optimising parameters in QAOA requires substantial amounts of computational effort: This aspect of the heuristic is **NP-hard** [18] even for classically tractable systems; likewise, every polynomial time algorithm is susceptible to instances for which the relative resulting error can be arbitrarily large, thus rendering approximate approaches likewise troublesome. The QAOA quantum circuit, in addition, usually needs to be evaluated for many different sets of angles to gain the necessary information on the optimisation landscape that is required as input for the classical parameter optimisation routine, albeit efforts differ depending on the concrete choice [71, 80, 87]

Our approximation approach allows us to separate instance sampling from optimisation landscape sampling. Even at a fixed point (β, γ) , directly sampling $F_1(\beta, \gamma)$ is intractable for exact computation, as it contains a sum over exponentially many terms. Likewise, determining an expectation value $E(F_1)$ requires to consider a substantial amount of points (β, γ) and instances. However, given structural information




Figure 9: Circuits for Hamming weight computation. D_H is used to construct the constraint Hamiltonian for the k -CLIQUE problem. It is used to check whether a potential clique is of size k or not. For this D_H adds the Hamming weight of register $|z\rangle$ to another register $|y\rangle$, while $D_H^\dagger$ subtracts it from $|y\rangle$. State $|y\rangle$ is initialised with $y = 0$ and holds the Hamming weight of z encoded as a bit string, after the application of D_H . The D_H gate adds the Hamming weight to $|y\rangle$ by conditionally applying U_{+1} to it for each qubit in $|z\rangle$, with $U_{+1}|y\rangle = |y + 1\rangle$. The inverse operation $D_H^\dagger$ subtracts the Hamming weight of z from $|y\rangle$, by conditionally applying U_{-1} for each qubit, with $U_{-1}|y\rangle = |y - 1\rangle$.




Figure 10: *Top*: 2D histogram of the F_1 values for all instances for $n = 18$ from the qr -FACTURING example, overlaid by the approximated landscape $\tilde{E}(F_1)$ (orange) and the expected landscape $E(F_1)$ (black). *Bottom*: Absolute approximation error (solid line) compared to the standard deviation of F_1 (dashed line). The variance in F_1 is negligible, which makes the margin of error for approximations extremely tight. Our approximation provides excellent results even under such challenging conditions.

about the target space in form of $\left\{E\left(\overline{\#_d(k)}\right)\right\}_{d=0}^n$ and $\left\{E\left(\overline{\#_{d_1}(k)\#_{d_2}(k)}\right)\right\}_{d_1,d_2=0}^n$, only a sum over linear many terms has to be evaluated to approximate $E(F_1)$ efficiently at a point (β, γ) . The inputs $\left\{E\left(\overline{\#_d(k)}\right)\right\}_{d=0}^n$ and $\left\{E\left(\overline{\#_{d_1}(k)\#_{d_2}(k)}\right)\right\}_{d_1,d_2=0}^n$ of our approximation method can reasonably be obtained empirically by statistical methods, or even theoretically modelled as demonstrated in Section 5.1. Additionally, Theorem 1 takes $E(|T|)$ as an input. This quantity is in fact the expected value of the solution of the corresponding counting problem. Counting problems and approximate counting have been studied extensively by the theoretical computer science community since the early days [26, 82, 92]. In general even approximate counting is a hard problem in most cases. Luckily, $E(|T|)$ only acts as a scaling factor. In most cases where one is only interested in the overall structure of the optimisation landscape, *i.e.* local optima, plateaus, etc. the global scaling factor can be discarded. Thus, an accurate approximation of $E(|T|)$ is not as important for the application of our main theorem.

The above mentioned separation of instance and object landscape sampling enables a different approach to QAOA that does not require an unbounded iteration involving the quantum circuit. It starts by sampling (partial) target spaces of random instances, from which the structural metrics needed as input for Theorem 1 are gathered. Then, a classical optimiser is utilised to determine the optimal QAOA angles with




Figure 11: Each red cross marks a combination (β, γ) obtained from the standard QAOA algorithm for one single instance of the subject problem. The parameter pairs are plotted on top the expected optimisation landscapes. As is visually apparent, optimal parameters for different instances follow the *problem-specific*, but *instance-independent* patterns expected from our considerations. Solutions for local optima of the instance-averaged landscape, as they arise for some QAOA executions, are typical artefacts expected from numerical optimisation.

regard to the approximate optimisation landscape, resulting from the previously sampled target spaces. Because we used the expected landscape of a random instance of the problem at hand, the found parameters apply to the complete problem. Thus, parameter optimisation only needs to be performed once per problem to find one single set of parameters for all instances. Finally, the resulting angles are used to initialise a QAOA circuit from which potential problem solutions will be sampled on quantum hardware. Following this approach, only the final sampling step needs to be performed on real quantum hardware. The standard QAOA procedure is in fact more of a heuristic than a quantum algorithm, where we have to optimise a quantum circuit for each instance. In our approach to QAOA we optimise the parameters on a error bounded approximation of the expected landscape. Thus, we end up with one circuit for all problem instances. This mathematically sound splitting of computation into a problem-global and instance-specific phase significantly moves QAOA towards the realm of true quantum algorithms, in stark contrast to empirically motivated heuristics. Recall Fig. 2, where we illustrated the difference between both approaches.

We explored this approach empirically (using numerical simulations; see Section A for details about our reproduction package that allows researchers to directly employ our approach or inspect our implementation) for the examples introduced above in Section 5. For every example, we randomly generated 50 instances to be solved with both QAOA methods: standard QAOA and non-iterative QAOA. The approximate optimisation landscape for the non-iterative QAOA was calculated based on the dataset generated earlier in Section 5. After optimisation, we visually




Figure 12: Probability of states sampled from the QAOA circuits to encode valid solution for the various subject problems considered in this paper with a standard QAOA approach (individual parameter optimisation for each instance), and with our non-iterative schema using a-priori parameters obtained from the approximated landscape for all instances.

verified each set of parameters for non-iterative QAOA resides on one of the two main extrema of the landscape. After that, the parameters were used to create a QAOA circuit for each problem instance. This circuit then was compared with a QAOA circuit trained on the actual instance. From both circuits, we sampled 50 potential solutions per instance (100 states were sampled per instance for *qr*-FACTORING to more accurately capture the extremely low maximal possible success probability). Fig. 14 shows the resulting approximate landscapes.

As for practical performance, we can see from Fig. 11 that parameters obtained by standard QAOA are mostly located around extrema of the approximate optimisation landscapes. However, the optimisation apparently also often ends in local maxima, especially for *qr*-FACTORING. This suggests matching success rates for standard and non-iterative QAOA. Indeed, we can see in Fig. 12 that states sampled from non-iterative QAOA are valid solutions with identical (or slightly better) probability than states sampled from standard QAOA. Hard combinatorial constraint satisfaction problems are known to have phase transitions between trivially under- and over-constrained instances, and the actually hard instances reside in the parameter region of this very phase transition. For SAT, the level of under- or over-constrainedness is measured by the coefficient of variables to clauses $\alpha = |C|/|V|$, where $|C|$ and $|V|$ denote the number of clauses and variables, respectively. The phase transition of SAT happens at $\alpha \approx 4$. For our comparison, we generated 50 instances for each value of $\alpha = 2$, $\alpha = 4$ and $\alpha = 6$. As before, the non-iterative QAOA approach is on par with standard QAOA for under- ($\alpha = 2$) and over-constrained ($\alpha = 6$) SAT instances, as well as for hard SAT instances right in the phase transition at $\alpha = 4$. This underlines the utility of our approach not only for “average”, but also hard instances. Fig. 13 visualises the outcome distribution in detail.

Judging from the above evaluation, our non-iterative QAOA algorithm at least matches standard QAOA success probabilities, and occasionally shows slight advantage. However, QAOA involves classical parameter optimisation for every instance, while our approach is more resource efficient: Only a constant number of quantum circuit evaluations instead of sampling the circuit in every iteration of the classical optimisation loop is required, and we avoid instance-specific classical optimisation altogether: One single up-front classical optimisation process is required to infer optimal parameters from the approximated landscape per problem. Costly circuit evaluations in the quantum-classical iteration of QAOA are replaced by classical instance sampling in non-iterative QAOA: The quantum resource is only used for instance optimisation, not for preparatory work.




Figure 13: Probability of states sampled from the QAOA circuits to encode a valid solution for SAT instances of different hardness α . The non-iterative QAOA variant is on par with standard QAOA for trivially under- ($\alpha = 2$) and over-constrained ($\alpha = 6$) instances, as well as for usually hard instances in-between at $\alpha = 4$ (a comparison between standard and non-iterative QAOA over all SAT instances is included in Fig. 12).

7 Discussion

We have considered QAOA from a mixed perspective comprising computer science and physics, and could not only establish a new approximation theorem for the optimisation landscape, but have also suggested an algorithm that translates these insights directly into useful advantages. The evaluation of both aspects in the preceding sections has shown that either improves upon the state of the art in various ways.

However, our results also raise new questions. Obvious issues include how to extend the approach to QAOA depths larger than one; possible gains in post-NISQ systems; a more comprehensive evaluation on specific industrial problems and larger instance sizes; and the robustness to noise. Likewise, the question of how computational power that eventually arises from our algorithm is distributed between classical and quantum resources, and what impact this distribution has on the properties like runtime or solution quality, remains open. We need to leave addressing these questions to future research.

For all subject problems analysed above in Section 5 our approximation matches the expected landscape extremely well, with negligibly small errors. Given that we could show the absolute approximation error to be bounded by $\sqrt{\text{Var}(|T|) \text{Var}(\overline{|c_k|^2})}$, it follows that for all problems with constant solution space sizes (*i.e.*, $\text{Var}(|T|) = 0$), our approximation actually delivers exact results. It should be noted that this requires knowledge of exact values for $\left\{ E\left(\overline{\#_d(k)}\right) \right\}_{d=0}^n$, $\left\{ E\left(\overline{\#_{d_1}(k)\#_{d_2}(k)}\right) \right\}_{d_1, d_2=0}^n$ and $E(|T|)$, which is usually not practically achievable, especially for empirical sampling. This explains the approximation errors observed for all subject problems with fixed target




Figure 14: Overview of all approximated landscapes for the considered subject problems (100 random instances per problem and dimension). Each panel shows the approximated optimisation landscape $\tilde{E}(F_1(\beta, \gamma))$ obtained by Theorem 1 as surrogate for the expected value of $F_1(\beta, \gamma)$. Macroscopic similarities do not only arise between instances of different dimensions (left to right), but also to a certain extent across problems (top to bottom).

sizes above—except for SAT. This observation naturally raises questions about the relationship between the estimation accuracy of the target space structures and the resulting approximation error. Here, a trade-off between sample size and approximation quality is to be expected and should be analysed in future work. An answer to this question also could give interesting insight into the variance of optimisation landscapes between individual problem instances.

Regarding the scalability of our approach, we note that given the structural information about the target space the landscape approximation according to Theorem 1 results in a evaluation of $\mathcal{O}(n^2)$ terms, where n is the number of qubits. Therefore, the landscape approximation itself scales efficiently with the number of qubits, obtaining the needed structural solution space information on the other hand depends on the chosen method. We demonstrated that for this both analytical and empirical approaches can be taken. In the first case the efficiency depends on the theoretical model and the closed form or approximation arising from it. For the latter, it would be interesting to investigate the approximation robustness of 1 under limited sampling. Another empirical approach where our framework could prove beneficial would be to deploy it in an online learning fashion, where the needed solution samples are gathered on the way. Although, out of scope and unaddressed in this paper, this could provide an interesting path investigate in future work. One could also be curious about the feasibility of our proposed landscape approximation for deeper circuits with an increasing number of layers. With more layers QAOA sees more of the graph, while it generally does not see the complete graph with $p < \log n$ [30, 31]. With Theorem 1 operating on global structures, we argue that these global structures should even become more important for deeper QAOA circuits scanning the complete qubit graph.

Our optimisation landscape approximation is solely based on structural information about a problem. This allowed us to prove the existence of conjectured instance invariants that stem from observed effects like parameter clustering [20, 84]. The approximation allows us to obtain a smooth representation of the complete optimisation landscape. In Fig. 14, we collect such landscapes for the subject problems discussed in Section 5 for different dimensions of the target space T . This provides us with some interesting observations: On the one hand, it can be seen that macroscopic similarities exist between all problems. They share very similar high level features (with some problem specific differentiation). These macroscopic features are also present for structure-less target sets like in the case of the uniform sampling example. We thus conjecture that these features are truly problem independent and could be explained by either just the approximation theorem presented in this work, or in combination with straight forward and minimalis-

tic examples like uniform sampled target sets. Additional structure causes some displacement and skewing. This structure can be plausibly explained by local effects like membership of a state to the target set depending on the membership of other neighbouring states. As we can additionally observe from Fig. 14, simply increasing the state space dimension (*i.e.*, going from left to right on the panels for a specific problem) seems to compress the landscape along $\beta = 0$ along the β -axis. This effect is independent of the subject problem (*i.e.*, does not change when traversing the panels from top to bottom in the plot). As it can also be observed in the uniform sampling example, it must be independent of structural properties. While further examining these effects could provide valuable insight into the behaviour of QAOA optimisation landscapes, we need to leave the efforts required for such an investigation to future work, as they go beyond the scope of this paper. The approximation theorem devised in this paper will likely serve as an analytical basis for such considerations.

Focusing on two-level Hamiltonians and QAOA circuits keeps the approximation framework simple and straightforward to use. We thus provide a well defined base framework to analyse the essential connection of problem structure and QAOA landscapes. At the same time, it is extendable to accommodate, for instance, Hamiltonians with more complex eigenspectra or QAOA circuits with more layers, albeit we also need to leave such efforts to future work. However, it seems pertinent to note that the landscapes for the subject problems considered in this work match empirical results with higher level eigenspectra and deeper QAOA circuits in their macroscopic structure [70, 84]. The same holds true for Ising model QAOA landscapes based on an instance-specific analytical derivation [65].

We believe our work lays a foundation to map significant insights from classical theoretical computer science to quantum approaches along the lines of QAOA. In the classical case, the analysis of solution space structures is well established [3, 8, 47, 68]. With the approximation theorem (and the techniques introduced to derive it), we present a handle to apply this knowledge to quantum optimisation. This provides an opportunity for future progress in understanding and utilising the class of algorithms initiated by QAOA.

8 Conclusion

Our new perspective on QAOA from a mixed physics and theoretical computer science point of view allowed us to accurately approximate the QAOA optimisation landscape based on the inherent structural properties of a problem, and prove a long-standing set of hypotheses about relationships between optimal QAOA parameters for different instances of a problem. By directly linking the solution space structure of prob-

lems to their QAOA optimisation landscapes, we could devise an approximation theorem that does not only provide structural insights, but also impacts the practical use of QAOA.

Based on our results, we have constructed a non-iterative QAOA algorithm that is more resource-efficient than standard QAOA for various measures. Our evaluation on five different scenarios and subject problems has shown excellent agreement from the theoretical and practical point of view, and provides concrete advantages over standard QAOA. Our new perspective on QAOA opens various possibilities for future research to understand unsolved issues about quantum optimisation.

Acknowledgements: We thank Jonathan Reyersbach for many active discussions that helped shape the mathematical exposition of our ideas, and acknowledge important comments from Lukas Schmidbauer, Simon Thelen, and Maja Franz to ascertain the correctness of our main theorem. We are grateful for funding from German Federal Ministry of Education and Research (BMBF), funding program “Quantum Technologies—from Basic Research to Market”, grant #13NI6092. WM also acknowledges support by the High-Tech Agenda Bavaria.

A Reproduction Package

To make our experiments reproducible by other researchers [1], we provide the complete source code for the calculations performed in the paper, in form of a long-term stable [58] [reproduction package](#) (link in PDF; a [DOI-safe](#) version [88] is also available). We ascertain that the package is fully self-contained, and does not rely on resources that may eventually vanish from public access. In particular, we provide code that is easily extensible to more subject problems than considered in the paper, and include routines to (a) perform target space sampling; (b) implement non-iterative QAOA; (c) perform numerical simulations to explore performance and compare with standard QAOA. Any raw data obtained from our simulations, together with a complete pipeline to perform the evaluation and visualisation presented in the paper, are included.

References

- [1] *Reproducibility and Replicability in Science*. National Academies Press, September 2019. ISBN 9780309486163. DOI: [10.17226/25303](https://doi.org/10.17226/25303). URL <http://dx.doi.org/10.17226/25303>.
- [2] Scott Aaronson. Read the fine print. *Nature Physics*, 11(4):291–293, April 2015. ISSN 1745-2481. DOI: [10.1038/nphys3272](https://doi.org/10.1038/nphys3272). URL <http://dx.doi.org/10.1038/nphys3272>.
- [3] Dimitris Achlioptas, Amin Coja-Oghlan, and Federico Ricci-Tersenghi. On the solution-space geometry of random constraint satisfaction problems. *Random Structures and Algorithms*, 38(3):251–268, March 2011. ISSN 1098-2418. DOI: [10.1002/rsa.20323](https://doi.org/10.1002/rsa.20323). URL <http://dx.doi.org/10.1002/rsa.20323>.
- [4] Leonard M. Adleman, Jonathan DeMarrais, and Ming-Deh A. Huang. Quantum computability. *SIAM Journal on Computing*, 26(5):1524–1540, October 1997. ISSN 1095-7111. DOI: [10.1137/s0097539795293639](https://doi.org/10.1137/s0097539795293639). URL <http://dx.doi.org/10.1137/S0097539795293639>.
- [5] V. Akshay, D. Rabinovich, E. Campos, and J. Biamonte. Parameter concentrations in quantum approximate optimization. *Physical Review A*, 104(1), July 2021. ISSN 2469-9934. DOI: [10.1103/physreva.104.1010401](https://doi.org/10.1103/physreva.104.1010401). URL <http://dx.doi.org/10.1103/physreva.104.1010401>.
- [6] M. Sohaib Alam, Filip A. Wudarski, Matthew J. Reagor, James Sud, Shon Grabbe, Zhihui Wang, Mark Hodson, P. Aaron Lott, Eleanor G. Rieffel, and Davide Venturelli. Practical verification of quantum properties in quantum-approximate-optimization runs. *Physical Review Applied*, 17(2), 2022. ISSN 2331-7019. DOI: [10.1103/physrevapplied.17.024026](https://doi.org/10.1103/physrevapplied.17.024026). URL <http://dx.doi.org/10.1103/PhysRevApplied.17.024026>.
- [7] Md Zahangir Alom, Brian Van Essen, Adam T. Moody, David Peter Widemann, and Tarek M. Taha. Quadratic unconstrained binary optimization (qubo) on neuromorphic computing system. In *2017 International Joint Conference on Neural Networks (IJCNN)*, page 3922–3929. IEEE, May 2017. DOI: [10.1109/ijcnn.2017.7966350](https://doi.org/10.1109/ijcnn.2017.7966350). URL <http://dx.doi.org/10.1109/ijcnn.2017.7966350>.
- [8] Carlos Ansótegui, Jesús Giraldez-Cru, and Jordi Levy. The community structure of SAT formulas. In Alessandro Cimatti and Roberto Sebastiani, editors, *Theory and Applications of Satisfiability Testing - SAT 2012 - 15th International Conference, Trento, Italy, June 17-20, 2012. Proceedings*, volume 7317 of *Lecture Notes in Computer Science*, pages 410–423. Springer, 2012. DOI: [10.1007/978-3-642-31612-8_31](https://doi.org/10.1007/978-3-642-31612-8_31). URL https://doi.org/10.1007/978-3-642-31612-8_31.
- [9] Maliheh Aramon, Gili Rosenberg, Elisabetta Valiante, Toshiyuki Miyazawa, Hiro-taka Tamura, and Helmut G. Katzgraber. Physics-inspired optimization for quadratic unconstrained problems using a digital annealer. *Frontiers in Physics*, 7, 2019. ISSN 2296-424X. DOI: [10.3389/fphy.2019.00048](https://doi.org/10.3389/fphy.2019.00048). URL <http://dx.doi.org/10.3389/fphy.2019.00048>.
- [10] Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, April 2009. ISBN 9780511804090. DOI: [10.1017/cbo9780511804090](https://doi.org/10.1017/cbo9780511804090). URL <http://dx.doi.org/10.1017/CB09780511804090>.
- [11] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G. S. L. Brandao, David A. Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P. Harrigan, Michael J. Hartmann, Alan Ho, Markus Hoffmann, Trent Huang, Travis S. Humble, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul V. Klimov, Sergey Knysh, Alexander Korkotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R. McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C. Platt, Chris Quintana, Eleanor G. Rieffel, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Kevin J. Sung, Matthew D. Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M. Martinis. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779): 505–510, October 2019. ISSN 1476-4687. DOI: [10.1038/s41586-019-1666-5](https://doi.org/10.1038/s41586-019-1666-5). URL <http://dx.doi.org/10.1038/s41586-019-1666-5>.
- [12] Gilles Audemard and Laurent Simon. Predicting learnt clauses quality in modern sat solvers. *IJCAI’09*, page 399–404, San Francisco, CA, USA, 2009. Morgan Kaufmann Publishers Inc.
- [13] Abhishek Awasthi, Francesco Bär, Joseph Doetsch, Hans Ehm, Marvin Erdmann, Maximilian Hess, Johannes Klepsch, Peter A. Limacher, Andre Luckow, Christoph Niedermeier, Lilly Palackal, Ruben Pfeiffer, Philipp Ross, Hila Safi, Janik Schönmeier-Kromer, Oliver von Sicard, Yannick Wenger, Karen Winter-

- sperger, and Sheir Yarkoni. Quantum computing techniques for multi-knapsack problems. In *Intelligent Computing*, page 264–284. Springer Nature Switzerland, 2023. ISBN 9783031379635. DOI: [10.1007/978-3-031-37963-5_19](https://doi.org/10.1007/978-3-031-37963-5_19). URL http://dx.doi.org/10.1007/978-3-031-37963-5_19.
- [14] Andreas Bayerstadler, Guillaume Becquin, Julia Binder, Thierry Botter, Hans Ehm, Thomas Ehmer, Marvin Erdmann, Norbert Gaus, Philipp Harbach, Maximilian Hess, Johannes Klepsch, Martin Leib, Sebastian Luber, Andre Luckow, Maximilian Manský, Wolfgang Mauerer, Florian Neukart, Christoph Niedermeier, Lilly Palackal, Ruben Pfeiffer, Carsten Polenz, Johanna Sepulveda, Tammo Sievers, Brian Standen, Michael Streif, Thomas Strohm, Clemens Utschig-Utschig, Daniel Volz, Horst Weiss, and Fabian Winter. Industry quantum computing applications. *EPJ Quantum Technology*, 8(1), 11 2021. DOI: [10.1140/epjqt/s40507-021-00114-x](https://doi.org/10.1140/epjqt/s40507-021-00114-x). URL <https://epjquantumtechnology.springeropen.com/track/pdf/10.1140/epjqt/s40507-021-00114-x.pdf>.
- [15] Andreas Bengtsson, Pontus Vikstål, Christopher Warren, Marika Svensson, Xiu Gu, Anton Frisk Kockum, Philip Krantz, Christian Križan, Daryoush Shiri, Ida-Maria Svensson, Giovanna Tancredi, Göran Johansson, Per Delsing, Giulia Ferrini, and Jonas Bylander. Improved success probability with greater circuit depth for the quantum approximate optimization algorithm. *Physical Review Applied*, 14(3), September 2020. ISSN 2331-7019. DOI: [10.1103/PhysRevApplied.14.034010](https://doi.org/10.1103/PhysRevApplied.14.034010). URL <http://dx.doi.org/10.1103/PhysRevApplied.14.034010>.
- [16] Ethan Bernstein and Umesh Vazirani. Quantum complexity theory. *SIAM Journal on Computing*, 26(5):1411–1473, October 1997. ISSN 1095-7111. DOI: [10.1137/s0097539796300921](https://doi.org/10.1137/s0097539796300921). URL <http://dx.doi.org/10.1137/S0097539796300921>.
- [17] Kishor Bharti, Alba Cervera-Lierta, Thi Ha Kyaw, Tobias Haug, Sumner Alperin-Lea, Abhinav Anand, Matthias Degroote, Hermann Heimonen, Jakob S. Kottmann, Tim Menke, Wai-Keong Mok, Sukin Sim, Leong-Chuan Kwek, and Alán Aspuru-Guzik. Noisy intermediate-scale quantum algorithms. *Rev. Mod. Phys.*, 94:015004, Feb 2022. DOI: [10.1103/RevModPhys.94.015004](https://doi.org/10.1103/RevModPhys.94.015004). URL <https://link.aps.org/doi/10.1103/RevModPhys.94.015004>.
- [18] Lennart Bittel and Martin Kliesch. Training Variational Quantum Algorithms Is NP-Hard. *Physical Review Letters*, 127(12):120502, September 2021. DOI: [10.1103/PhysRevLett.127.120502](https://doi.org/10.1103/PhysRevLett.127.120502). URL <https://link.aps.org/doi/10.1103/PhysRevLett.127.120502>. Publisher: American Physical Society.
- [19] Kostas Blekos, Dean Brand, Andrea Ceschini, Chiao-Hui Chou, Rui-Hao Li, Komal Pandya, and Alessandro 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://www.sciencedirect.com/science/article/pii/S0370157324001078>.
- [20] Fernando G. S. L. Brandao, Michael Broughton, Edward Farhi, Sam Gutmann, and Hartmut Neven. For fixed control parameters the quantum approximate optimization algorithm’s objective function value concentrates for typical instances. 2018. DOI: [10.48550/ARXIV.1812.04170](https://doi.org/10.48550/ARXIV.1812.04170). URL <https://arxiv.org/abs/1812.04170>.
- [21] Sergey Bravyi, Alexander Kliesch, Robert Koenig, and Eugene Tang. Obstacles to variational quantum optimization from symmetry protection. *Physical Review Letters*, 125(26), December 2020. ISSN 1079-7114. DOI: [10.1103/physrevlett.125.260505](https://doi.org/10.1103/physrevlett.125.260505). URL <http://dx.doi.org/10.1103/physrevlett.125.260505>.
- [22] Andreas Bärttschi and Stephan Eidenbenz. Grover Mixers for QAOA: Shifting complexity from mixer design to state preparation. In *Proc. IEEE QCE*, pages 72–82, 2020. DOI: [10.1109/QCE49297.2020.00020](https://doi.org/10.1109/QCE49297.2020.00020).
- [23] M. Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C. Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R. McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, and Patrick J. Coles. Variational quantum algorithms. *Nature Reviews Physics*, 3(9):625–644, August 2021. ISSN 2522-5820. DOI: [10.1038/s42254-021-00348-9](https://doi.org/10.1038/s42254-021-00348-9). URL <http://dx.doi.org/10.1038/s42254-021-00348-9>.
- [24] Barry A Cipra. The ising model is np-complete. *SIAM News*, 33(6):1–3, 2000.
- [25] Maxime Dupont, Nicolas Didier, Mark J. Hodson, Joel E. Moore, and Matthew J. Reagor. Entanglement perspective on the quantum approximate optimization algorithm. *Physical Review A*, 106(2), August 2022. ISSN 2469-9934. DOI: [10.1103/physreva.106.022423](https://doi.org/10.1103/physreva.106.022423). URL <http://dx.doi.org/10.1103/PhysRevA.106.022423>.
- [26] Martin E. Dyer. Approximate counting by dynamic programming. In Lawrence L. Larmore and Michel X. Goemans, editors, *Proceedings of the 35th Annual ACM Symposium on Theory of Computing, June 9-11, 2003, San Diego, CA, USA*, pages 693–699. ACM, 2003.

DOI: 10.1145/780542.780643. URL <https://doi.org/10.1145/780542.780643>.

- [27] Pablo Díez-Valle, Diego Porras, and Juan José García-Ripoll. Connection between single-layer quantum approximate optimization algorithm interferometry and thermal distribution sampling. *Frontiers in Quantum Science and Technology*, 3, February 2024. ISSN 2813-2181. DOI: 10.3389/frqst.2024.1321264. URL <http://dx.doi.org/10.3389/frqst.2024.1321264>.
- [28] Daniel J. Egger, Jakub Mareček, and Stefan Woerner. Warm-starting quantum optimization. *Quantum*, 5:479, June 2021. ISSN 2521-327X. DOI: 10.22331/q-2021-06-17-479. URL <https://doi.org/10.22331/q-2021-06-17-479>.
- [29] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm. 2014. DOI: 10.48550/ARXIV.1411.4028. URL <https://arxiv.org/abs/1411.4028>.
- [30] Edward Farhi, David Gamarnik, and Sam Gutmann. The quantum approximate optimization algorithm needs to see the whole graph: Worst case examples. *arXiv preprint arXiv:2005.08747*, 2020. DOI: 10.48550/arXiv.2005.08747. URL <https://doi.org/10.48550/arXiv.2005.08747>.
- [31] Edward Farhi, David Gamarnik, and Sam Gutmann. The quantum approximate optimization algorithm needs to see the whole graph: A typical case. *arXiv preprint arXiv:2004.09002*, 2020. DOI: 10.48550/arXiv.2004.09002. URL <https://doi.org/10.48550/arXiv.2004.09002>.
- [32] Mario Fernández-Pendás, Elías F. Combarro, Sofía Vallecorsa, José Ranilla, and Ignacio F. Rúa. A study of the performance of classical minimizers in the quantum approximate optimization algorithm. *Journal of Computational and Applied Mathematics*, 404: 113388, April 2022. ISSN 0377-0427. DOI: 10.1016/j.cam.2021.113388. URL <http://dx.doi.org/10.1016/j.cam.2021.113388>.
- [33] Mario Fernández-Pendás, Elías F. Combarro, Sofía Vallecorsa, José Ranilla, and Ignacio F. Rúa. A study of the performance of classical minimizers in the Quantum Approximate Optimization Algorithm. *Journal of Computational and Applied Mathematics*, 404:113388, April 2022. ISSN 0377-0427. DOI: 10.1016/j.cam.2021.113388. URL <https://www.sciencedirect.com/science/article/pii/S0377042721000078>.
- [34] Jernej Rudi Finžgar, Aron Kerschbaumer, Martin J.A. Schuetz, Christian B. Mendl, and Helmut G. Katzgraber. Quantum-informed recursive optimization algorithms. *PRX Quantum*, 5(2), May 2024. ISSN 2691-3399. DOI: 10.1103/prxquantum.5.020327. URL <http://dx.doi.org/10.1103/PRXQuantum.5.020327>.
- [35] Lance Fortnow and John Rogers. Complexity limitations on quantum computation. *Journal of Computer and System Sciences*, 59(2): 240–252, October 1999. ISSN 0022-0000. DOI: 10.1006/jcss.1999.1651. URL <http://dx.doi.org/10.1006/jcss.1999.1651>.
- [36] Thomas Gabor, Sebastian Zielinski, Sebastian Feld, Christoph Roch, Christian Seidel, Florian Neukart, Isabella Galter, Wolfgang Mauere, and Claudia Linnhoff-Popien. Assessing solution quality of 3sat on a quantum annealing platform. In *Quantum Technology and Optimization Problems*, page 23–35. Springer International Publishing, 2019. ISBN 9783030140823. DOI: 10.1007/978-3-030-14082-3_3. URL http://dx.doi.org/10.1007/978-3-030-14082-3_3.
- [37] Alexey Galda, Eesh Gupta, Jose Falla, Xiaoyuan Liu, Danylo Lykov, Yuri Alexeev, and Ilya Safro. Similarity-based parameter transferability in the quantum approximate optimization algorithm. *Frontiers in Quantum Science and Technology*, 2, July 2023. ISSN 2813-2181. DOI: 10.3389/frqst.2023.1200975. URL <http://dx.doi.org/10.3389/frqst.2023.1200975>.
- [38] Vijay Ganesh and Moshe Y. Vardi. *On the Unreasonable Effectiveness of SAT Solvers*, page 547–566. Cambridge University Press, December 2020. DOI: 10.1017/9781108637435.032. URL <http://dx.doi.org/10.1017/9781108637435.032>.
- [39] Martin Gogeißl, Hila Safi, and Wolfgang Mauere. Quantum data encoding patterns and their consequences. In *Proceedings of the Workshop on Quantum Computing and Quantum-Inspired Technology for Data-Intensive Systems and Applications*, QDSM '24, 08 2024. DOI: 10.1145/3665225.3665446. URL <https://doi.org/10.1145/3665225.3665446>.
- [40] Felix Greiwe, Tom Krüger, and Wolfgang Mauere. Effects of imperfections on quantum algorithms: A software engineering perspective. In *2023 IEEE International Conference on Quantum Software (QSW)*, page 31–42. IEEE, July 2023. DOI: 10.1109/qsw59989.2023.00014. URL <http://dx.doi.org/10.1109/qsw59989.2023.00014>.
- [41] G. G. Guerreschi and A. Y. Matsuura. Qaoa for max-cut requires hundreds of qubits for quantum speed-up. *Scientific Reports*, 9(1), May 2019. ISSN 2045-2322. DOI: 10.1038/s41598-019-43176-9. URL <http://dx.doi.org/10.1038/s41598-019-43176-9>.
- [42] Stuart Hadfield, Zhihui Wang, Bryan O’Gorman, Eleanor G. Rieffel, Davide Venturelli, and Rupak Biswas. From the

- quantum approximate optimization algorithm to a quantum alternating operator ansatz. *Algorithms*, 12(2), 2019. ISSN 1999-4893. DOI: [10.3390/a12020034](https://doi.org/10.3390/a12020034). URL <https://www.mdpi.com/1999-4893/12/2/34>.
- [43] Matthew P. Harrigan, Kevin J. Sung, Matthew Neeley, Kevin J. Satzinger, Frank Arute, Kunal Arya, Juan Atalaya, Joseph C. Bardin, Rami Barends, Sergio Boixo, Michael Broughton, Bob B. Buckley, David A. Buell, Brian Burkett, Nicholas Bushnell, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Sean Demura, Andrew Dunsworth, Daniel Eppens, Austin Fowler, Brooks Foxen, Craig Gidney, Marissa Giustina, Rob Graff, Steve Habegger, Alan Ho, Sabrina Hong, Trent Huang, L. B. Ioffe, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Cody Jones, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Seon Kim, Paul V. Klimov, Alexander N. Korotkov, Fedor Kostritsa, David Landhuis, Pavel Laptev, Mike Lindmark, Martin Leib, Orion Martin, John M. Martinis, Jarrod R. McClean, Matt McEwen, Anthony Megrant, Xiao Mi, Masoud Mohseni, Wojciech Mruczkiewicz, Josh Mutus, Ofer Naaman, Charles Neill, Florian Neukart, Murphy Yuezhen Niu, Thomas E. O’Brien, Bryan O’Gorman, Eric Ostby, Andre Petukhov, Harald Putterman, Chris Quintana, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Andrea Skolik, Vadim Smelyanskiy, Doug Strain, Michael Streif, Marco Szalay, Amit Vainsencher, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Leo Zhou, Hartmut Neven, Dave Bacon, Erik Lucero, Edward Farhi, and Ryan Babbush. Quantum approximate optimization of non-planar graph problems on a planar superconducting processor. *Nature Physics*, 17(3):332–336, February 2021. ISSN 1745-2481. DOI: [10.1038/s41567-020-01105-y](https://doi.org/10.1038/s41567-020-01105-y). URL <http://dx.doi.org/10.1038/s41567-020-01105-y>.
- [44] Kyle Henke, Elijah Pelofske, Georg Hahn, and Garrett Kenyon. Sampling binary sparse coding qubo models using a spiking neuromorphic processor. In *Proceedings of the 2023 International Conference on Neuromorphic Systems, ICONS ’23*, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400701757. DOI: [10.1145/3589737.3606003](https://doi.org/10.1145/3589737.3606003). URL <https://doi.org/10.1145/3589737.3606003>.
- [45] Rebekah Herrman, Lorna Treffert, James Ostrowski, Phillip C. Lotshaw, Travis S. Humble, and George Siopsis. Globally optimizing qaoa circuit depth for constrained optimization problems. *Algorithms*, 14(10):294, 2021. ISSN 1999-4893. DOI: [10.3390/a14100294](https://doi.org/10.3390/a14100294). URL <http://dx.doi.org/10.3390/a14100294>.
- [46] Marijn JH Heule, Markus Iser, Matti Järvisalo, and Martin Suda. Proceedings of sat competition 2024: Solver, benchmark and proof checker descriptions, 2024.
- [47] Tad Hogg. Refining the phase transition in combinatorial search. *Artificial Intelligence*, 81(1–2):127–154, March 1996. ISSN 0004-3702. DOI: [10.1016/0004-3702\(95\)00050-x](https://doi.org/10.1016/0004-3702(95)00050-x). URL [http://dx.doi.org/10.1016/0004-3702\(95\)00050-x](http://dx.doi.org/10.1016/0004-3702(95)00050-x).
- [48] Nishant Jain, Brian Coyle, Elham Kashefi, and Niraj Kumar. Graph neural network initialisation of quantum approximate optimisation. *Quantum*, 6:861, November 2022. ISSN 2521-327X. DOI: [10.22331/q-2022-11-17-861](https://doi.org/10.22331/q-2022-11-17-861). URL <https://doi.org/10.22331/q-2022-11-17-861>.
- [49] Zhang Jiang, Eleanor G. Rieffel, and Zhihui Wang. Near-optimal quantum circuit for grover’s unstructured search using a transverse field. *Physical Review A*, 95(6), June 2017. ISSN 2469-9934. DOI: [10.1103/PhysRevA.95.062317](https://doi.org/10.1103/PhysRevA.95.062317). URL <http://dx.doi.org/10.1103/PhysRevA.95.062317>.
- [50] Zhang Jiang, Eleanor G. Rieffel, and Zhihui Wang. Near-optimal quantum circuit for grover’s unstructured search using a transverse field. *Phys. Rev. A*, 95, Jun 2017. DOI: [10.1103/PhysRevA.95.062317](https://doi.org/10.1103/PhysRevA.95.062317). URL <https://link.aps.org/doi/10.1103/PhysRevA.95.062317>.
- [51] Richard M. Karp. *Reducibility among Combinatorial Problems*, page 85–103. Springer US, 1972. ISBN 9781468420012. DOI: [10.1007/978-1-4684-2001-2_9](https://doi.org/10.1007/978-1-4684-2001-2_9). URL http://dx.doi.org/10.1007/978-1-4684-2001-2_9.
- [52] Gary Kochenberger, Jin-Kao Hao, Fred Glover, Mark Lewis, Zhipeng Lü, Haibo Wang, and Yang Wang. The unconstrained binary quadratic programming problem: a survey. *Journal of Combinatorial Optimization*, 28(1): 58–81, April 2014. ISSN 1573-2886. DOI: [10.1007/s10878-014-9734-0](https://doi.org/10.1007/s10878-014-9734-0). URL <http://dx.doi.org/10.1007/s10878-014-9734-0>.
- [53] Tom Krüger and Wolfgang Maurer. Quantum annealing-based software components: An experimental case study with SAT solving. In *Proceedings of the IEEE/ACM 42nd International Conference on Software Engineering Workshops, ICSEW’20*, page 445–450, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450379632. DOI: [10.1145/3387940.3391472](https://doi.org/10.1145/3387940.3391472). URL <https://doi.org/10.1145/3387940.3391472>.
- [54] Wolfgang Lechner. Quantum approximate optimization with parallelizable gates. *IEEE Transactions on Quantum Engineering*, 1:1–6, 2020. ISSN 2689-1808. DOI:

- 10.1109/tqe.2020.3034798. URL <http://dx.doi.org/10.1109/TQE.2020.3034798>.
- [55] Juneseo Lee, Alicia B. Magann, Herschel A. Rabitz, and Christian Arenz. Progress toward favorable landscapes in quantum combinatorial optimization. *Physical Review A*, 104(3), 2021. ISSN 2469-9934. DOI: 10.1103/physrev.a.104.032401. URL <http://dx.doi.org/10.1103/PhysRevA.104.032401>.
- [56] Phillip C. Lotshaw, Travis S. Humble, Rebekah Herrman, James Ostrowski, and George Siopsis. Empirical performance bounds for quantum approximate optimization. *Quantum Information Processing*, 20(12):403, December 2021. ISSN 1570-0755, 1573-1332. DOI: 10.1007/s11128-021-03342-3. URL <http://arxiv.org/abs/2102.06813>. arXiv:2102.06813 [physics, physics:quant-ph].
- [57] Andrew Lucas. Ising formulations of many np problems. *Frontiers in Physics*, 2, 2014. ISSN 2296-424X. DOI: 10.3389/fphy.2014.00005. URL <http://dx.doi.org/10.3389/fphy.2014.00005>.
- [58] Wolfgang Mauerer and Stefanie Scherzinger. 1-2-3 reproducibility for quantum software experiments. In *2022 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*, page 1247–1248. IEEE, March 2022. DOI: 10.1109/saner53432.2022.00148. URL <http://dx.doi.org/10.1109/saner53432.2022.00148>.
- [59] Catherine C. McGeoch and Pau Farré. Milestones on the quantum utility highway: Quantum annealing case study. *ACM Transactions on Quantum Computing*, 5(1), dec 2023. DOI: 10.1145/3625307. URL <https://doi.org/10.1145/3625307>.
- [60] Matija Medvidović and Giuseppe Carleo. Classical variational simulation of the quantum approximate optimization algorithm. *npj Quantum Information*, 7(1), June 2021. ISSN 2056-6387. DOI: 10.1038/s41534-021-00440-z. URL <http://dx.doi.org/10.1038/s41534-021-00440-z>.
- [61] Rémi Monasson, Riccardo Zecchina, Scott Kirkpatrick, Bart Selman, and Lidror Troyansky. Determining computational complexity from characteristic ‘phase transitions’. *Nature*, 400(6740): 133–137, July 1999. ISSN 1476-4687. DOI: 10.1038/22055. URL <http://dx.doi.org/10.1038/22055>.
- [62] Ashley Montanaro. Quantum algorithms: an overview. *npj Quantum Information*, 2(1), January 2016. ISSN 2056-6387. DOI: 10.1038/npjqi.2015.23. URL <http://dx.doi.org/10.1038/npjqi.2015.23>.
- [63] J. A. Montañez-Barrera and Kristel Michiels. Toward a linear-ramp qaoa protocol: evidence of a scaling advantage in solving some combinatorial optimization problems. *npj Quantum Information*, 11(1), August 2025. ISSN 2056-6387. DOI: 10.1038/s41534-025-01082-1. URL <http://dx.doi.org/10.1038/s41534-025-01082-1>.
- [64] M. E. S. Morales, J. D. Biamonte, and Z. Zimborás. On the universality of the quantum approximate optimization algorithm. *Quantum Information Processing*, 19(9), August 2020. ISSN 1573-1332. DOI: 10.1007/s11128-020-02748-9. URL <http://dx.doi.org/10.1007/s11128-020-02748-9>.
- [65] Asier Ozaeta, Wim van Dam, and Peter L McMahon. Expectation values from the single-layer quantum approximate optimization algorithm on ising problems. *Quantum Science and Technology*, 7(4):045036, September 2022. ISSN 2058-9565. DOI: 10.1088/2058-9565/ac9013. URL <http://dx.doi.org/10.1088/2058-9565/ac9013>.
- [66] Guido Pagano, Aniruddha Bapat, Patrick Becker, Katherine S. Collins, Arinjoy De, Paul W. Hess, Harvey B. Kaplan, Antonis Kyprianidis, Wen Lin Tan, Christopher Baldwin, Lucas T. Brady, Abhinav Deshpande, Fangli Liu, Stephen Jordan, Alexey V. Gorshkov, and Christopher Monroe. Quantum approximate optimization of the long-range ising model with a trapped-ion quantum simulator. *Proceedings of the National Academy of Sciences*, 117(41):25396–25401, October 2020. ISSN 1091-6490. DOI: 10.1073/pnas.2006373117. URL <http://dx.doi.org/10.1073/pnas.2006373117>.
- [67] Yu Pan, Yifan Tong, Shibe Xue, and Guofeng Zhang. Efficient depth selection for the implementation of noisy quantum approximate optimization algorithm. *Journal of the Franklin Institute*, 359(18): 11273–11287, December 2022. ISSN 0016-0032. DOI: 10.1016/j.jfranklin.2022.10.027. URL <http://dx.doi.org/10.1016/j.jfranklin.2022.10.027>.
- [68] Pushkin R. Pari, Jane Lin, Lin Yuan, and Gang Qu. Generating ‘random’ 3-sat instances with specific solution space structure. In Deborah L. McGuinness and George Ferguson, editors, *Proceedings of the Nineteenth National Conference on Artificial Intelligence, Sixteenth Conference on Innovative Applications of Artificial Intelligence, July 25-29, 2004, San Jose, California, USA*, pages 960–961. AAAI Press / The MIT Press, 2004. URL <http://www.aaai.org/Library/AAAI/2004/aaai04-129.php>.
- [69] Aidan Pellow-Jarman, Shane McFarthing, Ilya Sinayskiy, Daniel K. Park, Anban Pillay, and

- Francesco Petruccione. The effect of classical optimizers and ansatz depth on qaoa performance in noisy devices. *Scientific Reports*, 14(1), July 2024. ISSN 2045-2322. DOI: [10.1038/s41598-024-66625-6](https://doi.org/10.1038/s41598-024-66625-6). URL <http://dx.doi.org/10.1038/s41598-024-66625-6>.
- [70] Elijah Pelofske, Andreas Bärttschi, and Stephan Eidenbenz. Short-depth qaoa circuits and quantum annealing on higher-order ising models. *npj Quantum Information*, 10(1), March 2024. ISSN 2056-6387. DOI: [10.1038/s41534-024-00825-w](https://doi.org/10.1038/s41534-024-00825-w). URL <http://dx.doi.org/10.1038/s41534-024-00825-w>.
- [71] Maniraman Periyasamy, Axel Plinge, Christopher Mutschler, Daniel D. Scherer, and Wolfgang Mauerer. Guided-spsa: Simultaneous perturbation stochastic approximation assisted by the parameter shift rule. In *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)*, page 1504–1515. IEEE, September 2024. DOI: [10.1109/qce60285.2024.00177](https://doi.org/10.1109/qce60285.2024.00177). URL <http://dx.doi.org/10.1109/qce60285.2024.00177>.
- [72] Niklas Pirnay, Vincent Ulitzsch, Frederik Wilde, Jens Eisert, and Jean-Pierre Seifert. An in-principle super-polynomial quantum advantage for approximating combinatorial optimization problems via computational learning theory. *Science Advances*, 10(11), March 2024. ISSN 2375-2548. DOI: [10.1126/sciadv.adj5170](https://doi.org/10.1126/sciadv.adj5170). URL <http://dx.doi.org/10.1126/sciadv.adj5170>.
- [73] Stefan H. Sack and Maksym Serbyn. Quantum annealing initialization of the quantum approximate optimization algorithm. *Quantum*, 5:491, July 2021. ISSN 2521-327X. DOI: [10.22331/q-2021-07-01-491](https://doi.org/10.22331/q-2021-07-01-491). URL <http://dx.doi.org/10.22331/q-2021-07-01-491>.
- [74] Hila Safi, Karen Wintersperger, and Wolfgang Mauerer. Influence of hw-sw-co-design on quantum computing scalability. In *IEEE International Conference on Quantum Software (QSW)*, pages 104–115, 2023. DOI: [10.1109/QSW59989.2023.00022](https://doi.org/10.1109/QSW59989.2023.00022).
- [75] Irmi Sax, Sebastian Feld, Sebastian Zielinski, Thomas Gabor, Claudia Linnhoff-Popien, and Wolfgang Mauerer. Approximate approximation on a quantum annealer. In *Proceedings of the 17th ACM International Conference on Computing Frontiers*, page 108–117, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450379564. DOI: [10.1145/3387902.3392635](https://doi.org/10.1145/3387902.3392635). URL <https://arxiv.org/pdf/2004.09267>.
- [76] Lukas Schmidbauer, Karen Wintersperger, Elisabeth Lobe, and Wolfgang Mauerer. Polynomial reduction methods and their impact on QAOA circuits. In *IEEE International Conference on Quantum Software, QSW 2024, Shenzhen, China, July 7-13, 2024*, pages 35–45. IEEE, 2024. DOI: [10.1109/QSW62656.2024.00018](https://doi.org/10.1109/QSW62656.2024.00018). URL <https://doi.org/10.1109/QSW62656.2024.00018>.
- [77] Manuel Schönberger, Immanuel Trummer, and Wolfgang Mauerer. Quantum-inspired digital annealing for join ordering. *Proc. VLDB Endow.*, 17(3):511–524, nov 2023. ISSN 2150-8097. DOI: [10.14778/3632093.3632112](https://doi.org/10.14778/3632093.3632112). URL <https://doi.org/10.14778/3632093.3632112>.
- [78] Changpeng Shao, Yang Li, and Hongbo Li. Quantum algorithm design: Techniques and applications. *Journal of Systems Science and Complexity*, 32(1):375–452, February 2019. ISSN 1559-7067. DOI: [10.1007/s11424-019-9008-0](https://doi.org/10.1007/s11424-019-9008-0). URL <http://dx.doi.org/10.1007/s11424-019-9008-0>.
- [79] Ruslan Shaydulin, Changhao Li, Shouvanik Chakrabarti, Matthew DeCross, Dylan Herman, Niraj Kumar, Jeffrey Larson, Danylo Lykov, Pierre Minssen, Yue Sun, Yuri Alexeev, Joan M. Dreiling, John P. Gaebler, Thomas M. Gatterman, Justin A. Gerber, Kevin Gilmore, Dan Gresh, Nathan Hewitt, Chandler V. Horst, Shaohan Hu, Jacob Johansen, Mitchell Matheny, Tanner Mengle, Michael Mills, Steven A. Moses, Brian Neyenhuis, Peter Siegfried, Romina Yalovetzky, and Marco Pistoia. Evidence of scaling advantage for the quantum approximate optimization algorithm on a classically intractable problem. *Science Advances*, 10(22), May 2024. ISSN 2375-2548. DOI: [10.1126/sciadv.adm6761](https://doi.org/10.1126/sciadv.adm6761). URL <http://dx.doi.org/10.1126/sciadv.adm6761>.
- [80] Sanyam Singhal, Vandit Srivastava, P Rohith, Prateek Jain, and Debanjan Bhowmik. Performance comparison for quantum approximate optimization algorithm (qaoa) across noiseless simulation, experimentally benchmarked noisy simulation, and experimental hardware platforms. In *2024 8th IEEE Electron Devices Technology and Manufacturing Conference (EDTM)*, page 1–3. IEEE, March 2024. DOI: [10.1109/edtm58488.2024.10512142](https://doi.org/10.1109/edtm58488.2024.10512142). URL <http://dx.doi.org/10.1109/edtm58488.2024.10512142>.
- [81] Daniel Stilck França and Raul García-Patrón. Limitations of optimization algorithms on noisy quantum devices. *Nature Physics*, 17(11):1221–1227, October 2021. ISSN 1745-2481. DOI: [10.1038/s41567-021-01356-3](https://doi.org/10.1038/s41567-021-01356-3). URL <http://dx.doi.org/10.1038/s41567-021-01356-3>.
- [82] Larry J. Stockmeyer. The complexity of approximate counting (preliminary version). In David S. Johnson, Ronald Fagin, Michael L. Fredman, David Harel, Richard M. Karp,

- Nancy A. Lynch, Christos H. Papadimitriou, Ronald L. Rivest, Walter L. Ruzzo, and Joel I. Seiferas, editors, *Proceedings of the 15th Annual ACM Symposium on Theory of Computing, 25-27 April, 1983, Boston, Massachusetts, USA*, pages 118–126. ACM, 1983. DOI: 10.1145/800061.808740. URL <https://doi.org/10.1145/800061.808740>.
- [83] Michael Streif and Martin Leib. Comparison of qaoa with quantum and simulated annealing. 2019. DOI: 10.48550/ARXIV.1901.01903. URL <https://arxiv.org/abs/1901.01903>.
- [84] Michael Streif and Martin Leib. Training the quantum approximate optimization algorithm without access to a quantum processing unit. *Quantum Science and Technology*, 5(3):034008, May 2020. ISSN 2058-9565. DOI: 10.1088/2058-9565/ab8c2b. URL <http://dx.doi.org/10.1088/2058-9565/ab8c2b>.
- [85] James Sud, Stuart Hadfield, Eleanor Rieffel, Norm Tubman, and Tad Hogg. Parameter-setting heuristic for the quantum alternating operator ansatz. *Phys. Rev. Res.*, 6, May 2024. DOI: 10.1103/PhysRevResearch.6.023171. URL <https://link.aps.org/doi/10.1103/PhysRevResearch.6.023171>.
- [86] Reuben Tate, Jai Moondra, Bryan Gard, Greg Mohler, and Swati Gupta. Warm-Started QAOA with Custom Mixers Provably Converges and Computationally Beats Goemans-Williamson’s Max-Cut at Low Circuit Depths. *Quantum*, 7:1121, September 2023. ISSN 2521-327X. DOI: 10.22331/q-2023-09-26-1121. URL <https://doi.org/10.22331/q-2023-09-26-1121>.
- [87] Simon Thelen, Hila Safi, and Wolfgang Mauerer. Approximating under the influence of quantum noise and compute power. In *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)*, page 274–279. IEEE, September 2024. DOI: 10.1109/qce60285.2024.10291.
- [88] Wolfgang Mauerer Tom Krüger. Reproduction package: Out of the loop: Structural approximation of optimisation landscapes and non-iterative quantum optimisation, 2025. URL <https://doi.org/10.5281/zenodo.13286242>.
- [89] V Vijendran, Aritra Das, Dax Enshan Koh, Syed M Assad, and Ping Koy Lam. An expressive ansatz for low-depth quantum approximate optimisation. *Quantum Science and Technology*, 9(2):025010, feb 2024. DOI: 10.1088/2058-9565/ad200a. URL <https://dx.doi.org/10.1088/2058-9565/ad200a>.
- [90] Zhihui Wang, Nicholas C. Rubin, Jason M. Dominy, and Eleanor G. Rieffel. Xy mixers: Analytical and numerical results for the quantum alternating operator ansatz. *Physical Review A*, 101(1), January 2020. ISSN 2469-9934. DOI: 10.1103/physreva.101.012320. URL <http://dx.doi.org/10.1103/PhysRevA.101.012320>.
- [91] Dave Wecker, Matthew B. Hastings, and Matthias Troyer. Training a quantum optimizer. *Physical Review A*, 94(2), August 2016. ISSN 2469-9934. DOI: 10.1103/physreva.94.022309. URL <http://dx.doi.org/10.1103/PhysRevA.94.022309>.
- [92] Wei Wei and Bart Selman. A new approach to model counting. In Fahiem Bacchus and Toby Walsh, editors, *Theory and Applications of Satisfiability Testing, 8th International Conference, SAT 2005, St. Andrews, UK, June 19-23, 2005, Proceedings*, volume 3569 of *Lecture Notes in Computer Science*, pages 324–339. Springer, 2005. DOI: 10.1007/11499107_24. URL https://doi.org/10.1007/11499107_24.
- [93] Johannes Weidenfeller, Lucia C. Valor, Julien Gacon, Caroline Tornow, Luciano Bello, Stefan Woerner, and Daniel 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. URL <http://dx.doi.org/10.22331/q-2022-12-07-870>.
- [94] Dennis Willsch, Madita Willsch, Fengping Jin, Kristel Michielsen, and Hans De Raedt. Gpu-accelerated simulations of quantum annealing and the quantum approximate optimization algorithm. *Computer Physics Communications*, 278:108411, September 2022. ISSN 0010-4655. DOI: 10.1016/j.cpc.2022.108411. URL <http://dx.doi.org/10.1016/j.cpc.2022.108411>.
- [95] Madita Willsch, Dennis Willsch, Fengping Jin, Hans De Raedt, and Kristel Michielsen. Benchmarking the quantum approximate optimization algorithm. *Quantum Information Processing*, 19(7), June 2020. ISSN 1573-1332. DOI: 10.1007/s11128-020-02692-8. URL <http://dx.doi.org/10.1007/s11128-020-02692-8>.
- [96] Karen Wintersperger, Hila Safi, and Wolfgang Mauerer. *QPU-System Co-design for Quantum HPC Accelerators*, page 100–114. Springer International Publishing, 2022. ISBN 9783031218675. DOI: 10.1007/978-3-031-21867-5_7. URL http://dx.doi.org/10.1007/978-3-031-21867-5_7.
- [97] Elisabeth Wybo and Martin Leib. Missing puzzle pieces in the performance landscape of the quantum approximate optimization algorithm. 2024. DOI: 10.48550/ARXIV.2406.14618. URL <https://arxiv.org/abs/2406.14618>.
- [98] Leo Zhou, Sheng-Tao Wang, Soonwon Choi, Hannes Pichler, and Mikhail D. Lukin. Quantum approximate optimization algorithm: Per-

- formance, mechanism, and implementation on near-term devices. *Physical Review X*, 10(2), June 2020. ISSN 2160-3308. DOI: [10.1103/physrevx.10.021067](https://doi.org/10.1103/physrevx.10.021067). URL <http://dx.doi.org/10.1103/physrevx.10.021067>.
- [99] Sebastian Zielinski, Jonas Nüklein, Jonas Stein, Thomas Gabor, Claudia Linnhoff-Popien, and Sebastian Feld. Pattern qubos: Algorithmic construction of 3sat-to-qubo transformations. *Electronics*, 12(16):3492, August 2023. ISSN 2079-9292. DOI: [10.3390/electronics12163492](https://doi.org/10.3390/electronics12163492). URL <http://dx.doi.org/10.3390/electronics12163492>.
- [100] Oylum Şeker, Neda Tanoumand, and Merve Bodur. Digital annealer for quadratic unconstrained binary optimization: A comparative performance analysis. *Applied Soft Computing*, 127:109367, 2022. ISSN 1568-4946. DOI: <https://doi.org/10.1016/j.asoc.2022.109367>. URL <https://www.sciencedirect.com/science/article/pii/S156849462200521X>.