Circuit depth scaling for quantum approximate optimization

V. Akshay [1,*] H. Philathong,[1] E. Campos,[1] D. Rabinovich,[1] I. Zacharov,[1] Xiao-Ming Zhang,[2] and J. D. Biamonte [1]

[1]*Skolkovo Institute of Science and Technology, 30 Bolshoy Boulevard, Moscow 121205, Russian Federation*
[2]*Department of Physics, City University of Hong Kong, Tat Chee Avenue, Kowloon, Hong Kong SAR, China*

Variational quantum algorithms are the centerpiece of modern quantum programming. These algorithms involve training parametrized quantum circuits using a classical coprocessor, an approach adapted partly from classical machine learning. An important subclass of these algorithms, designed for combinatorial optimization on current quantum hardware, is the quantum approximate optimization algorithm (QAOA). Despite efforts to realize deeper circuits, experimental state-of-the-art implementations are limited to a fixed depth. However, it is known that *problem density*—a problem constraint to a variable ratio—induces underparametrization in fixed depth QAOA. Density-dependent performance has been reported in the literature, yet the circuit depth required to achieve fixed performance (henceforth called critical depth) remained unknown. Here, we propose a predictive model, based on a *logistic saturation conjecture* for critical depth scaling with respect to density. Focusing on random instances of MAX-2-SAT, we test our predictive model against simulated data with up to 15 qubits. We report the average critical depth, required to attain a success probability of 0.7, saturates at a value of 10 for densities beyond 4. We observe the predictive model to describe the simulated data within a 3σ confidence interval. Furthermore, based on the model, a linear trend for the critical depth with respect to problem size is recovered for the range of 5–15 qubits.

I. INTRODUCTION

The quantum approximate optimization algorithm (QAOA) was introduced to solve combinatorial optimization problems by Farhi *et al.* [1]. In this algorithm, the candidate solution is prepared using a fixed ansatz structure with tunable parameters that are variationally adjusted to minimize an objective function. QAOA was inspired by adiabatic quantum computation [2–4], where a system prepared in the ground state of an initial Hamiltonian evolves towards the ground state of a problem Hamiltonian, given a slow enough evolution. The circuit structure of QAOA shares a resemblance to adiabatic evolution via a finite depth Trotterization procedure, and thus allows for recovering the exact ground state of the problem Hamiltonian in the infinite depth limit. In contrast, a fixed p-depth QAOA circuit consists of alternating applications of the complex exponentiation of the problem Hamiltonian and a so-called mixer Hamiltonian, accounting for $2p$ variational parameters.

Although QAOA recovers adiabatic evolution in the infinite depth limit, short depth circuits have indeed been shown to offer some benefits in approximating solutions to low-density instances of combinatorial optimization problems [1,5,6]. Farhi *et al.* [1] studied the performance of QAOA applied to the Maximum-Cut (or MAX-CUT) problem on 3-regular graphs and reported that for $p = 1$, QAOA guarantees a cut that is at least 0.6924 of the optimal cut. Furthermore, the quality of the approximation is shown to improve with

increasing p. Beyond combinatorial optimization problems, Farhi and Harrow [7] reported that the split operator structure of QAOA might be a suitable candidate for establishing quantum supremacy, since the output distributions of such circuits may not be efficiently sampled classically. Moreover, in the works of Jiang *et al.* [8] and Morales *et al.* [9], the $O(\sqrt{N})$ query complexity of Grover's algorithm is recovered by a class of QAOA circuits, thus establishing quantum advantage within the QAOA framework. These, along with universality results [10–12], make QAOA a promising model for noisy intermediate-scale quantum (NISQ) era devices.

Although several advantages have been reported [1,5–9,13–15], the effects limiting the performance of QAOA have also been discovered [16–18]. Considering the energy error in the approximation as a performance metric, Akshay *et al.* [16,17] demonstrated *reachability deficits*, an effect limiting performance for fixed depth QAOA on high-density problem instances. Deeper circuits therefore become a necessity and hence come with overheads in classical outer-loop optimization. While approaches such as layerwise training aim to reduce this classical computational cost, Campos *et al.* [18] showed that such a strategy would not improve performance beyond some threshold depth. Other approaches aimed at reducing the classical computation cost exploit what are known as *concentration effects* [19–22]. Although these effects allow one to heuristically guess near-optimal parameters at fixed depth and for an increasing number of qubits, they fail to address the depth required to guarantee fixed performance.

Here, we propose a *logistic saturation conjecture* which states that the minimum depth required for QAOA to achieve

---

*akshay.vishwanathan@skoltech.ru

fixed performance up to some tolerance (henceforth called critical depth) scales logistically with problem density. We numerically test this conjecture for the case of maximum satisfiability (specifically MAX-2-SAT) with up to 15 qubits and maximum density of 4, and observe (i) the predictive model describes the simulated data within a 3σ confidence interval and (ii) the critical depth scales linearly with problem size in this data set.

II. BACKGROUND

A. Quantum approximate optimization algorithm

The objective function of a combinatorial optimization problem is prescribed as a map from the assignment of variables to non-negative real numbers $\mathcal{C} : \{0, 1\}^{\times n} \to \mathbb{R}_+$. Under a vector space embedding [23,24], the objective function then induces a corresponding problem Hamiltonian $\mathcal{H}_c$, which encodes the solutions to the problem instance in its ground state space,

$$\mathcal{H}_c = \sum_{z \in \{0,1\}^{\times n}} \mathcal{C}(z)|z\rangle\langle z|, \tag{1}$$

where $|z\rangle$ are the computational basis vectors in the 2^n-dimensional Hilbert space of n qubits.

The optimization problem aims to find assignments such that the objective function $\mathcal{C}(z)$ is minimized. This is equivalent to finding a state $|z^*\rangle$ such that

$$|z^*\rangle \in \text{span}\left\{\arg\min_{z \in \{0,1\}^{\times n}}\langle z|\mathcal{H}_c|z\rangle\right\}. \tag{2}$$

Measuring the state $|z^*\rangle$ in the computational basis outputs a bit string which is a solution to the optimization problem $\mathcal{C}$.

To approximate solutions to an optimization problem instance using a p-depth QAOA circuit, we generate ansatz states $|\psi(\boldsymbol{\gamma}, \boldsymbol{\beta})\rangle$ on a quantum computer with tunable real parameters $\boldsymbol{\gamma}, \boldsymbol{\beta} \in [0, 2\pi)^{\times p}, [0, \pi)^{\times p}$. These states are prepared by applying an alternating sequence of $2p$ parametrized unitary gates on the initial state $|+\rangle^{\otimes n} = \frac{1}{\sqrt{2^n}}\sum_{z \in \{0,1\}^{\times n}}|z\rangle$ as

$$|\psi(\boldsymbol{\gamma}, \boldsymbol{\beta})\rangle = \prod_{k=1}^{p} e^{-i\beta_k \mathcal{H}_x} e^{-i\gamma_k \mathcal{H}_c}|+\rangle^{\otimes n}, \tag{3}$$

with the mixing Hamiltonian $\mathcal{H}_x = \sum_{l=1}^{n} X_l$. Here, X_l is the Pauli matrix applied to the lth qubit.

The ansatz state in Eq. (3) is iteratively adjusted via an outer-loop optimization routine to obtain an energy approximation as

$$\min_{\boldsymbol{\gamma}, \boldsymbol{\beta}} \langle\psi(\boldsymbol{\gamma}, \boldsymbol{\beta})|\mathcal{H}_c|\psi(\boldsymbol{\gamma}, \boldsymbol{\beta})\rangle \geqslant \min(\mathcal{H}_c). \tag{4}$$

The algorithmic performance of QAOA can be assessed under performance metrics such as (i) the energy error in the approximation and (ii) the ground state overlap [16–18]. Both these quantifiers feature shortcomings as energy approximations fail to indicate the closeness to optimal solutions while computing the ground state overlap requires prior knowledge of optimal solutions. However, it can be demonstrated that both metrics interrelate and therefore have some advantages as a theoretical tool for studying algorithmic performance.

Let $\mathcal{H}^\dagger = \mathcal{H} \geqslant 0$ be an optimization problem instance with d-degenerate ground state energy λ_0. In its eigenbasis $|\phi_j\rangle$, we can write $\mathcal{H} = \lambda_0 \sum_{j=1}^{d} |\phi_j\rangle\langle\phi_j| + \sum_{k>d} \lambda_k|\phi_k\rangle\langle\phi_k|$, with subsequent excited state energies, $\lambda_0 < \lambda_k \leqslant \lambda_{\max}$ for $k > d$, separated by a gap $\Delta = \min_{k>d} |\lambda_0 - \lambda_k|$. For any arbitrary state $|\psi\rangle$, the ground state overlap is calculated as

$$g(\psi) = \sum_{j=1}^{d} |\langle\phi_j|\psi\rangle|^2. \tag{5}$$

It is straightforward now to derive (see Appendix B) the bounds relating the energy error in approximations $\langle\psi|\mathcal{H}|\psi\rangle - \lambda_0$, and the ground state overlap $g(\psi)$:

$$1 - \frac{\langle\psi|\mathcal{H}|\psi\rangle - \lambda_0}{\Delta} \leqslant g(\psi) \leqslant 1 - \frac{\langle\psi|\mathcal{H}|\psi\rangle - \lambda_0}{\lambda_{\max} - \lambda_0}. \tag{6}$$

Note that the relation works when the error in approximation is less than the spectral gap Δ. These bounds are the extended version of the variational stability lemma in Ref. [10] to the degenerate case and allows one to recover an estimate on the ground state overlap from the measured energy.

B. Satisfiability

Boolean satisfiability, or SAT, is the problem of determining the satisfiability of Boolean formulas. More specifically, one aims to decide whether a given formula can be satisfied by assigning truth values to the variables. Limited to k variables or literals per clause, and expressed in the conjunctive normal form (CNF), k-SAT is *NP*-complete for $k \geqslant 3$ [25]. k-SAT clauses are randomly generated to form instances by uniformly selecting unique k-tuples from the union of a variable set (cardinality $n > k$) and its elementwise negation. Generated random instances admit an order parameter called a clause or problem density given by

$$\alpha = m/n, \tag{7}$$

where m is the number of clauses and n is the number of Boolean variables in the k-SAT instance.

The clause density plays an important role in the study of satisfiability phase transitions. It has been observed that for random k-SAT instances, an abrupt change in satisfiability occurs at some critical clause density α_c. Though this criticality has been studied only empirically for 3-SAT ($\alpha_c \approx 4.27$), it is proven for 2-SAT [26]. Furthermore, it is observed that the computational resources required for SAT solvers increase at the transition point with an easy-hard-easy pattern. This behavior suggests that the hard instances concentrate around the critical clause density [26–31].

In this paper we focus on the optimization version of k-SAT, otherwise called MAX-k-SAT. MAX-k-SAT is the canonical *NP*-hard optimization problem (for $k \geqslant 2$) [32] where one seeks to determine variable assignments that maximize the number of satisfied clauses in a given instance. Similar to decision k-SAT, MAX-k-SAT also features a phase transition with an easy-hard pattern [33]. Although the point of criticality differs for the optimization and decision variants of these problems, they coincide for $k = 2$ with $\alpha_c = 1$ [34].

With standard Ising embedding techniques [23,24], MAX-2-SAT instances can be mapped onto 2-local Ising

Hamiltonians of the form

$$\mathcal{H} = \sum_{i<j} \mathcal{J}_{ij} Z_i Z_j + \sum_i h_i Z_i, \qquad (8)$$

with appropriate coefficients and scale such that (i) for satisfiable assignments $|z\rangle$, $\langle z|\mathcal{H}|z\rangle = 0$ and (ii) for unsatisfiable assignments $\langle z|\mathcal{H}|z\rangle = m'$, where m' is the number of clauses violated by the assignment. A detailed construction of such Hamiltonians is shown in Appendix A. By construction, the ground state space of $\mathcal{H}$ is spanned by assignments that violate the least number of clauses in a given MAX-2-SAT instance.

III. LOGISTIC SATURATION CONJECTURE

The state space which can be prepared with a p-depth QAOA circuit is defined as

$$\Omega = \bigcup_{\gamma, \beta} |\psi(\gamma, \beta)\rangle, \qquad (9)$$

with $|\psi(\gamma, \beta)\rangle$ given by Eq. (3). In general, for a given instance size n, density α, and depth p, the variational state space may not cover the whole Hilbert space. Therefore, optimization over Ω may not converge to the exact ground state energy. The difference, or the energy error, in the approximation,

$$f(p, \alpha, n) = \min_{\psi \in \Omega \subseteq \mathbb{C}_2^{\otimes n}} \langle\psi|\mathcal{H}|\psi\rangle - \min_{\phi \in \mathbb{C}_2^{\otimes n}} \langle\phi|\mathcal{H}|\phi\rangle, \qquad (10)$$

quantifies the performance of p-depth QAOA. Whenever $f(p, \alpha, n) > 0$, QAOA suffers a *reachability deficit*. This behavior has been studied in Refs. [16,17], and it is observed that the onset of these deficits correspond to the phase transition criticality of the considered optimization problem. For the case of MAX-2-SAT, it was shown that $f(p, \alpha, n)$ grows with increasing clause density beyond $\alpha_c \approx 1$. Though better performance is achieved at the cost of increasing QAOA depth, a similar trend for deficits is reflected.

Motivated by these findings, we propose a *logistic saturation conjecture*: Given some fixed tolerance on performance, $\epsilon > 0$. For QAOA on instances characterized by density α, the critical depth p^*, required to guarantee approximations that fall within an ϵ tolerance, is given by

$$p^*(\alpha) = \min\{p \mid f(p, \alpha, n) \leqslant \epsilon\}. \qquad (11)$$

We conjecture the critical depth to scale with density as

$$p^*(\alpha) \approx \frac{p_{\max}}{1 + e^{-\kappa(\alpha - \alpha_c)}}, \qquad (12)$$

where α_c is the critical density, κ the logistic growth rate, and $p_{\max}$ the saturation value.

IV. RESULTS

We first recover the findings of Refs. [16,17] in Fig. 1 by studying the performance of fixed depth QAOA on 300 randomly generated MAX-2-SAT instances on 15 variables and clauses ranging from 1 to 60 ($\alpha \in [0.06, 4.0]$).

Numerically, it is seen that the average performance worsens monotonically with increasing clause density. This is also evident (see Fig. 2) for the case of ground state overlap as described by stability bounds in Eq. (6).
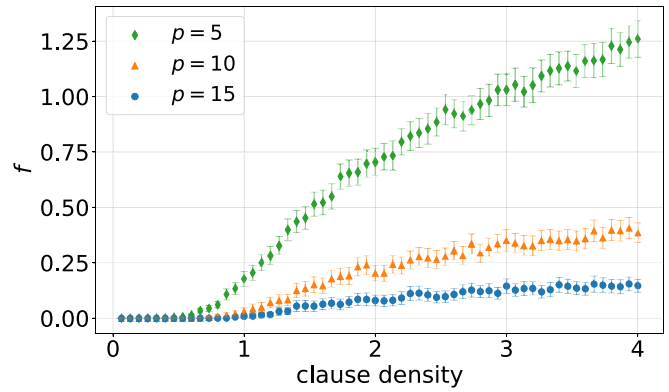

FIG. 1. Average energy error in approximation, f in Eq. (10), vs clause density for MAX-2-SAT instances on 15 variables. Data points represent the average value obtained over statistics of 300 random instances with error bars representing 3σ for the estimated mean. Colors indicate varying QAOA depths $p = \{5, 10, 15\}$, illustrating improved performance at higher depths.

It becomes apparent that for guaranteeing fixed performance, the QAOA depth must necessarily increase with increasing clause density.

To test the logistic saturation conjecture, we numerically calculate the depth required to attain performance within an ϵ tolerance. This is done by incrementing depth, until the energy error in the approximations falls within acceptance, $f(p, \alpha, n) \leqslant \epsilon$. For each instance this process is repeated and the average critical depth is calculated. Figure 3 illustrates scaling of the average critical depth with respect to clause density and compares the numerical data against the predictive model in Eq. (12). We see that the model is able to match the numerical data well within a 3σ confidence interval for the estimated mean.

We further investigate the reliability of the predictive model by treating the critical density as a fitting parameter in Eq. (12) and compare against the theoretical value of $\alpha_c \approx 1$. In Fig. 4 fitting parameters are recovered for $n = 5$ up to 15 and indeed we see the critical density recovered from the predictive model matches the theoretical value up to finite-size variations.
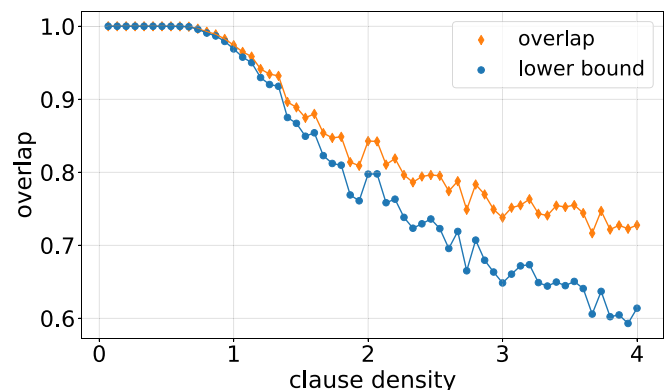

FIG. 2. Ground state overlap vs clause density for depth $p = 10$ QAOA on MAX-2-SAT instances with 15 variables. Stability lower bound is calculated by considering average energy error in approximation obtained over statistics of 300 random instances. Points in orange indicate the corresponding average ground state overlap.
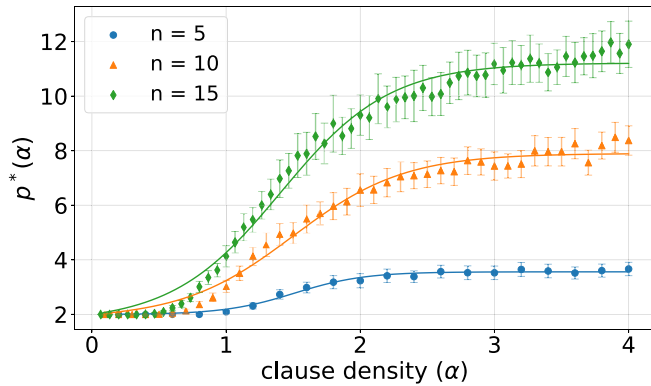
FIG. 3. Average critical depth p^* vs clause density. Data points represent the average p^* calculated according to Eq. (11) across the problem density for MAX-2-SAT with tolerance $\epsilon = 0.3$. Colors correspond to different problem sizes $n = \{5, 10, 15\}$ with error bars representing 3σ for the estimated mean. Solid curves represent the least-squares fit of the data to the predictive model as described in Eq. (12).
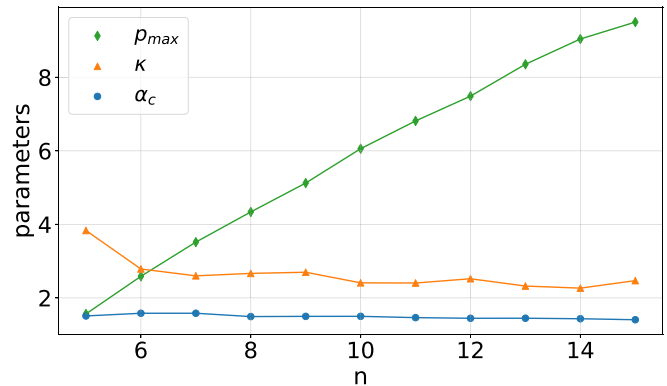

FIG. 4. Fitting parameters vs number of qubits for QAOA on MAX-2-SAT with tolerance $\epsilon = 0.3$. Note that the recovered critical density $\alpha_c \approx 1$, and growth rate κ, show very little variability with respect to the number of qubits. In contrast, the saturation value $p_{\max}$ illustrates a linear trend in the considered range of 5–15 qubits.

V. NUMERICAL DETAILS

The computationally intensive calculations were performed on the KunLun supercomputer with a maximum CPU frequency of 2.1 GHz. The code was written in PYTHON and C++ and we employ a full state-vector emulation scheme with the unitaries appearing in Eq. (3) implemented as array manipulations. The PYTHON code is responsible for providing easy-to-use interfaces and C++ code is responsible for the actual simulation tasks as the underlying simulation code. The code is executed in parallel with a resource of 1500 cores. The calculations were performed for 300 random instances for each respective problem density and problem sizes.

We used a heuristic optimisation strategy, motivated by layerwise training [35,36], to find good initial parametrization for the ansatz such that gradient-based classical optimizers may avoid most local minima and saturation effects [18,37].

In order to find optimal parameters for depth p, we start by optimizing a depth one ansatz (single layer). Then keeping the previous parameters fixed, a second layer is added and trained. We then use this layerwise trained circuit as a starting point for simultaneously optimizing all the circuit parameters before considering the third layer. This simultaneous optimization step allows for the optimizer to explore paths in the optimization landscape which do not saturate [18]. Furthermore, the approach can recursively be repeated until the desired depth p is attained. In our implementation of this heuristic, at each training step, the best result out of 25 random seed runs of the L-BFGS-B optimizer is considered to escape potential local minima.

VI. CONCLUSION

QAOA is among the most studied gate-based approaches for combinatorial optimization in NISQ era devices. Although analytical techniques address QAOA either at low depth or in the adiabatic limit, $p \to \infty$, understanding the performance at intermediate depths remains largely open. Furthermore, problem density induces underparametrization and limits the performance of fixed depth QAOA. An open question in this regard is the required depth for QAOA to succeed. Additionally, by treating the circuit depth of QAOA as a computational resource, it becomes paramount in complexity studies to know the scaling of depth with respect to the problem size. Understanding these open problems paves the way towards identifying the quantum advantage in QAOA.

In this paper we introduce a methodology based on a predictive model that aims to address these open problems. We first relate QAOA energy approximations and success probability (or ground state overlap) via Eq. (6). This relation then determines the success/failure of the algorithm up to some ϵ tolerance. Based on this acceptance criterion, we numerically recovered the critical depth needed for QAOA to achieve a success probability of at least 0.7 for the case of MAX-2-SAT. Considering up to 15 variables and a clause density up to 4 we indeed observe that our proposed model describes the data within a 3σ confidence interval. Furthermore, we recovered critical depth scaling with respect to the problem size to illustrate a linear trend for our considered range of problem sizes. Though this finite range is insufficient to assert quantum advantage, we anticipate future works to test the accuracy of the presented model over comprehensive ranges of densities and problem sizes.

The data that support the findings of this study are available within the paper. The code for generating the data will be available on reasonable request.

APPENDIX A: ISING EMBEDDING OF SAT INSTANCES

We start by mapping logical bits $\{0, 1\}$ to states $\{|0\rangle, |1\rangle\} \in \mathbb{C}_2$, respectively. The objective function to be minimized is equivalent to minimizing a Hamiltonian $\mathcal{H}_{\mathrm{SAT}}$, constructed by the real linear extension of $\{P_0, P_1, \mathbb{1}\}$ as

$$x_j \longrightarrow P_j^0, \quad \neg x_j \longrightarrow P_j^1, \tag{A1}$$

and

$$\wedge \longrightarrow +, \quad \vee \longrightarrow \otimes. \tag{A2}$$

Here, $P_j^1 = |1\rangle\langle 1|_j$ and $P_j^0 = |0\rangle\langle 0|_j$ are rank-one projectors acting on the jth qubit. The construction ensures each clause in a given SAT instance is mapped onto projectors, $P_{ij\cdots k}^{\alpha\beta\cdots\gamma} = P_i^\alpha \otimes P_j^\beta \otimes \cdots \otimes P_k^\gamma$, which penalize unsatisfiable assignments $|x\rangle \in \mathbb{C}_2^{\otimes n}$ with a penalty $\langle x|P_{ij\cdots k}^{\alpha\beta\cdots\gamma}|x\rangle = 1$.

The full Hamiltonian $\mathcal{H}_{\mathrm{SAT}}$ is then constructed by summing over all the clauses in an instance,

$$\mathcal{H}_{\mathrm{SAT}} = \sum_{l=1}^{m} \mathcal{C}_l\{P_{ij\cdots k}^{\alpha\beta\cdots\gamma}\}, \tag{A3}$$

where $\mathcal{C}_l$ assigns the value of $\alpha, \beta, \ldots, \gamma \in \{0, 1\}$, corresponding to the negation of the literals indexed by $i, j, \ldots, k \in \{1, 2, \ldots, n\}$ in the lth clause. It is straightforward to see that the ground state space of this SAT Hamiltonian is spanned by assignments that violate the minimal number of clauses.

Decomposing the projectors onto the Z Pauli basis as

$$P_j^\alpha = \frac{1}{2}[\mathbb{1} + (-1)^\alpha Z_j], \tag{A4}$$

Eq. (A3) can be rewritten in the form of a generalized Ising Hamiltonian with at most a k-body interaction. In case of 2-SAT, the Hamiltonian $\mathcal{H}_{\mathrm{2SAT}}$ contains at most two-body interactions as

$$\begin{aligned}
\mathcal{H}_{\mathrm{2SAT}} &= \sum_{l=1}^{m} \mathcal{C}_l\{P_j^\alpha \otimes P_k^\beta\} \\
&= \frac{1}{4}\left(m\mathbb{1} + \sum_j h_j Z_j + \sum_{j<k} \mathcal{J}_{jk} Z_j Z_k\right),
\end{aligned} \tag{A5}$$

with appropriate coefficients h_j and $\mathcal{J}_{jk}$.

APPENDIX B: DERIVATION OF ENERGY-OVERLAP BOUNDS

Once a variational state $|\psi(\boldsymbol{\theta})\rangle$ has been prepared on a quantum computer as a function of a set of tunable parameters $\boldsymbol{\theta}$, we would like to know how close this state is to the solution space of the problem being solved. We will adapt and extend the Lemma 1 (variational stability) from Ref. [10] to the degenerate case.

Let $\mathcal{H} \geqslant 0$ be some Hamiltonian with d-degenerate ground state energy $\lambda_0 \geqslant 0$. The Hamiltonian in its eigenbasis can be expressed as

$$\mathcal{H} = \lambda_0 \sum_{j=1}^{d} |\phi_j\rangle\langle\phi_j| + \sum_{k>d} \lambda_k |\phi_k\rangle\langle\phi_k|. \tag{B1}$$

Here, we reorder the eigenbasis so that $\lambda_0 < \lambda_1 \leqslant \cdots \lambda_k \cdots \leqslant \lambda_{\max}$. Consider the expectation value of the Hamiltonian on an arbitrary state $|\psi\rangle$,

$$\langle\psi|\mathcal{H}|\psi\rangle = \lambda_0 \sum_{j=1}^{d} |\langle\phi_j|\psi\rangle|^2 + \sum_{k>d} \lambda_k |\langle\phi_k|\psi\rangle|^2. \tag{B2}$$

Assuming the gap $\Delta = \lambda_1 - \lambda_0$ is known,

$$\begin{aligned}
\langle\psi|\mathcal{H}|\psi\rangle &= \lambda_0 \sum_{j=1}^{d} |\langle\phi_j|\psi\rangle|^2 + \sum_{k>d} \lambda_k |\langle\phi_k|\psi\rangle|^2 \\
&\geqslant \lambda_0 \left(\sum_{j=1}^{d} |\langle\phi_j|\psi\rangle|^2\right) + (\lambda_0 + \Delta)\left(\sum_{k>d} |\langle\phi_k|\psi\rangle|^2\right) \\
&= \lambda_0 \left(\sum_{j=1}^{d} |\langle\phi_j|\psi\rangle|^2\right) + (\lambda_0 + \Delta)\left(1 - \sum_{j=1}^{d} |\langle\phi_j|\psi\rangle|^2\right) \\
&= \lambda_0 + \Delta\left(1 - \sum_{j=1}^{d} |\langle\phi_j|\psi\rangle|^2\right).
\end{aligned}$$

Rearranging the terms now gives the required upper bound as:

$$\sum_{j=1}^{d} |\langle\phi_j|\psi\rangle|^2 \geqslant 1 - \frac{\langle\psi|\mathcal{H}|\psi\rangle - \lambda_0}{\lambda_{\max} - \lambda_0}. \tag{B3}$$

For the upper bound, we substitute for each excited state energies λ_k in Eq. (B2) by $\lambda_{\max}$,

$$\begin{aligned}
\langle\psi|\mathcal{H}|\psi\rangle &\leqslant \lambda_0 \sum_{j=1}^{d} |\langle\phi_j|\psi\rangle|^2 + \lambda_{\max} \sum_{k>d} |\langle\phi_k|\psi\rangle|^2 \\
&\leqslant \lambda_0 \left(\sum_{j=1}^{d} |\langle\phi_j|\psi\rangle|^2\right) \\
&\quad + \lambda_{\max}\left(1 - \sum_{j=1}^{d} |\langle\phi_j|\psi\rangle|^2\right).
\end{aligned} \tag{B4}$$

Rearranging the terms now gives the required upper bound as

$$\sum_{j=1}^{d} |\langle\phi_j|\psi\rangle|^2 \leqslant 1 - \frac{\langle\psi|\mathcal{H}|\psi\rangle - \lambda_0}{\lambda_{\max} - \lambda_0}. \tag{B5}$$

By combining both bounds, we obtain the stability theorem that relates energy error in approximation with the success probability (or the ground state overlap) for the degenerated case:

$$1 - \frac{f}{\Delta} \leqslant \sum_{i=1}^{d} |\langle\phi_i|\psi\rangle|^2 \leqslant 1 - \frac{f}{\lambda_{\max} - \lambda_0}, \tag{B6}$$

where $f = \langle\psi|\mathcal{H}|\psi\rangle - \lambda_0$.

Notice that for the case of MAX-2-SAT instances, $\lambda_{\max} \leqslant m$, where m the number of clauses in the instance and $\Delta \geqslant 1$.

[1] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, arXiv:1411.4028.

[2] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse Ising model, Phys. Rev. E **58**, 5355 (1998).

[3] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, A quantum adiabatic evolution algorithm applied to random instances of an *NP*-complete problem, Science **292**, 472 (2001).

[4] S. Boixo, T. F. Rønnow, S. V. Isakov, Z. Wang, D. Wecker, D. A. Lidar, J. M. Martinis, and M. Troyer, Evidence for quantum annealing with more than one hundred qubits, Nat. Phys. **10**, 218 (2014).

[5] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm applied to a bounded occurrence constraint problem, arXiv:1412.6062.

[6] Z. Wang, S. Hadfield, Z. Jiang, and E. G. Rieffel, Quantum approximate optimization algorithm for MaxCut: A fermionic view, Phys. Rev. A **97**, 022304 (2018).

[7] E. Farhi and A. W. Harrow, Quantum supremacy through the quantum approximate optimization algorithm, arXiv:1602.07674.

[8] Z. Jiang, E. G. Rieffel, and Z. Wang, Near-optimal quantum circuit for Grover's unstructured search using a transverse field, Phys. Rev. A **95**, 062317 (2017).

[9] M. E. S. Morales, T. Tlyachev, and J. Biamonte, Variational learning of Grover's quantum search algorithm, Phys. Rev. A **98**, 062333 (2018).

[10] J. Biamonte, Universal variational quantum computation, Phys. Rev. A **103**, L030401 (2021).

[11] M. E. S. Morales, J. D. Biamonte, and Z. Zimborás, On the universality of the quantum approximate optimization algorithm, Quantum Inf. Process. **19**, 291 (2020).

[12] S. Lloyd, Quantum approximate optimization is computationally universal, arXiv:1812.11075.

[13] E. Anschuetz, J. Olson, A. Aspuru-Guzik, and Y. Cao, Variational quantum factoring, in *International Workshop on Quantum Technology and Optimization Problems* (Springer, Berlin, 2019), pp. 74–85.

[14] W. W. Ho and T. H. Hsieh, Efficient variational simulation of non-trivial quantum states, SciPost Phys. **6**, 029 (2019).

[15] Z.-C. Yang, A. Rahmani, A. Shabani, H. Neven, and C. Chamon, Optimizing Variational Quantum Algorithms Using Pontryagin's Minimum Principle, Phys. Rev. X **7**, 021027 (2017).

[16] V. Akshay, H. Philathong, M. E. S. Morales, and J. D. Biamonte, Reachability Deficits in Quantum Approximate Optimization, Phys. Rev. Lett. **124**, 090504 (2020).

[17] V. Akshay, H. Philathong, I. Zacharov, and J. Biamonte, Reachability deficits in quantum approximate optimization of graph problems, Quantum **5**, 532 (2021).

[18] E. Campos, D. Rabinovich, V. Akshay, and J. Biamonte, Training saturation in layerwise quantum approximate optimization, Phys. Rev. A **104**, L030401 (2021).

[19] E. Farhi, J. Goldstone, S. Gutmann, and L. Zhou, The quantum approximate optimization algorithm and the Sherrington-Kirkpatrick model at infinite size, Quantum **6**, 759 (2022).

[20] M. Streif and M. Leib, Training the quantum approximate optimization algorithm without access to a quantum processing unit, Quantum Sci. Technol. **5**, 034008 (2020).

[21] V. Akshay, D. Rabinovich, E. Campos, and J. Biamonte, Parameter concentrations in quantum approximate optimization, Phys. Rev. A **104**, L010401 (2021).

[22] D. Wecker, M. B. Hastings, and M. Troyer, Training a quantum optimizer, Phys. Rev. A **94**, 022309 (2016).

[23] J. D. Biamonte, Nonperturbative *k*-body to two-body commuting conversion Hamiltonians and embedding problem instances into Ising spins, Phys. Rev. A **77**, 052331 (2008).

[24] A. Lucas, Ising formulations of many *NP* problems, Front. Phys. **2**, 5 (2014).

[25] S. A. Cook, The complexity of theorem-proving procedures, in *Proceedings of the Third Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, 1971), pp. 151–158.

[26] A. Goerdt, A threshold for unsatisfiability, J. Comput. Syst. Sci. **53**, 469 (1996).

[27] J. M. Crawford and L. D. Auton, Experimental results on the crossover point in random 3-SAT, Artif. Intell. **81**, 31 (1996).

[28] E. Friedgut, J. Bourgain *et al.*, Sharp thresholds of graph properties, and the *k*-SAT problem, J. Am. Math. Soc. **12**, 1017 (1999).

[29] B. Selman and S. Kirkpatrick, Critical behavior in the computational cost of satisfiability testing, Artif. Intell. **81**, 273 (1996).

[30] B. Selman, D. G. Mitchell, and H. J. Levesque, Generating hard satisfiability problems, Artif. Intell. **81**, 17 (1996).

[31] V. Chvátal and B. Reed, Mick gets some (the odds are on his side)(satisfiability), in *Proceedings of the 33rd Annual Symposium on Foundations of Computer Science* (IEEE Computer Society, Washington, DC, 1992), pp. 620–627.

[32] F. Barahona, On the computational complexity of Ising spin glass models, J. Phys. A: Math. Gen. **15**, 3241 (1982).

[33] W. Zhang, Phase transitions and backbones of 3-SAT and maximum 3-SAT, in *International Conference on Principles and Practice of Constraint Programming* (Springer, Berlin, 2001), pp. 153–167.

[34] D. Coppersmith, D. Gamarnik, M. T. Hajiaghayi, and G. B. Sorkin, Random MAX SAT, random MAX CUT, and their phase transitions, Random Struct. Alg. **24**, 502 (2004).

[35] A. Skolik, J. R. McClean, M. Mohseni, P. van der Smagt, and M. Leib, Layerwise learning for quantum neural networks, Quantum Mach. Intell. **3**, 5 (2021).

[36] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, Quantum Approximate Optimization Algorithm: Performance, Mechanism, and Implementation on Near-Term Devices, Phys. Rev. X **10**, 021067 (2020).

[37] M. Cain, E. Farhi, S. Gutmann, D. Ranard, and E. Tang, The QAOA gets stuck starting from a good classical string, arXiv:2207.05089.