

Vanishing performance of the parity-encoded quantum approximate optimization algorithm applied to spin-glass models

Elisabeth Wybo and Martin Leib

IQM Quantum Computers, Georg-Brauchle-Ring 23-25, 80992 München, Germany

The parity mapping provides a geometrically local encoding of the Quantum Approximate Optimization Algorithm (QAOA), at the expense of having a quadratic qubit overhead for all-to-all connected problems. In this work, we benchmark the parity-encoded QAOA on spin-glass models. We address open questions in the scaling of this algorithm. In particular, we show that for fixed number of parity-encoded QAOA layers, the performance or the output energy, vanishes towards zero (the value achieved by random guessing) with problem size N as $N^{-1/2}$. Our results suggest that the parity-encoded QAOA does not have a promising scaling compared to the standard version of QAOA. We perform tensor-network calculations to confirm our results, and comment on the concentration of optimal QAOA parameters over problem instances.

1 Introduction

Variational quantum algorithms are considered valuable candidates for providing useful quantum-computational results on near-term devices [1, 2]. The Quantum Approximate Optimization Algorithm (QAOA), originally introduced by Farhi in Ref. [3], is such a variational algorithm for gate-based machines, that has as goal to solve combinatorial optimization problems. In some specific cases, QAOA has provable performance guarantees, and can even outperform the best known assumption-free classical algorithm, without the need of optimizing the variational parameters on an instance-by-instance basis [4, 5].

Most combinatorial optimization problems, however, are defined on arbitrary graph topologies, some even all-to-all connected. This implies that the topology of the problem rarely matches the topology of the actual device, and that rerouting using SWAP networks is necessary [6, 7]. A general way to circumvent the need for rerouting, is to employ different problem encodings that are native to the topology of the hardware, for instance a square grid. In this respect, the parity encoding [8] forms a general encoding scheme to map geometrically non-local optimization problems to local ones. The encoded ‘parity qubits’ represent products of logical qubits, at the expense of having a quadratic overhead and a set of local constraints that need to be satisfied.

However, apart from this clear advantage that qubit routing is not needed, it is unclear if such a strategy is actually promising in terms of algorithmic performance. In this work, we aim to address this question by investigating the scaling of the performance of parity-encoded QAOA as introduced in Ref. [9] for paradigmatic all-to-all connected problems, namely Sherrington-Kirkpatrick spin glasses. Our work addresses for the first time the open questions about the scaling of this algorithm, and also connects to the question of whether non-local optimization problems can be solved up to high precision by geometrically local algorithms. Thus also addressing of whether such algorithms can be useful for near-term hardware implementations. While formally not excluding a faster runtime, our results suggests that there are no clear scaling advantages of the parity-encoded algorithm compared to the standard QAOA algorithm, that would justify the quadratic qubit

Elisabeth Wybo: elisabeth.wybo@meetiqm.com

Martin Leib: martin.leib@meetiqm.com




Figure 1: Sketch of the parity mapping: the parity qubits (shown in blue) represent the parity between two logical qubits (shown in red) of the original problem graph, consisting of $N = 6$ qubits. A problem graph that is all-to-all connected, leads to a triangular layout on a square lattice of $K = N(N - 1)/2$ parity qubits, where the parity qubits need to satisfy plaquette constraints. Those constraints follow from a set of closed cycles that form a cycle basis of the original problem graph. An example of such a cycle is indicated in red.

overhead. As we will demonstrate, the main reason for this, is that the number of parity QAOA layers p needs to grow with the problem size in order to have non-vanishing performance in the thermodynamic limit. Here, we define performance as the average energy of decoded bitstrings according to ‘qubit lines’, in close correspondence to the decoding scheme proposed in Ref. [10]. No matter how the performance is exactly defined, it is likely that p will need to grow with growing system size, as each parity-encoded QAOA layer has fixed depth because of the geometric locality of the algorithm. This should be contrasted with the standard (‘vanilla’) QAOA where the depth of each QAOA layer scales with system size for all-to-all connected problems.

Additionally, we perform numerical experiments based on tensor-network techniques. These techniques have become standard to simulate quantum algorithms [11, 12, 13, 14, 15]. These experiments confirm our findings and illustrate that the algorithm does not give promising results. In practice, we even find that it is not outperforming a classical algorithm that simply sets each parity spin according to the sign of the interaction strength. Importantly, however, QAOA is considered to be a hybrid quantum-classical algorithm [16, 17] because the variational parameters in the ansatz circuit need to be optimized classically. This is known to be a tedious procedure, that is affected by various difficulties, e.g. Barren plateaus [18]. However, it has become clear that a parameter concentration phenomenon is present for vanilla QAOA, e.g. for MaxCut or Sherrington-Kirkpatrick spin-glass problems [4, 5, 19, 20, 21, 22, 23, 24]. This means that the optimal variational parameters for typical problem instances concentrate around specific values. If the optimal parameters do not depend much on system size (which is the case for Sherrington-Kirkpatrick vanilla QAOA), such a parameter concentration is of great practical importance, as expensive parameter optimization loops can be eliminated, even for large instances. Our numerical findings indicate that also for the parity-encoded QAOA the optimal parameters concentrate. For $p = 1$, we determine these optimal angles analytically. However, as our scaling results indicate, unlike vanilla QAOA for Sherrington-Kirkpatrick [4], fixed optimal angles do not lead to fixed performance for parity-encoded QAOA. This follows from the fact that parity-constraints ‘average out’ in the optimization procedure, preventing entanglement to build up. In the end, we will conclude that the parity-encoded QAOA does not show promising performance when compared to vanilla QAOA.

The structure of this paper is the following: in Sec. 2 we introduce the problem we want to solve and give an overview of parity-encoded QAOA. For determining the quality of the solution, we use a similar cost function as in Ref. [10]. Second, in Sec. 3 we discuss our analytical results, and give a vanishing bound on the performance of parity-encoded QAOA for fixed number of layers. In Sec. 4, we confirm our findings by tensor-network calculations of the circuit, and finally conclude in Sec. 5.

2 Parity encoded QAOA and spin-glass models

We consider a paradigmatic classical all-to-all connected spin-glass model, known as the Sherrington-Kirkpatrick model [25]

$$H_P = \sum_{i=0}^{N-2} \sum_{j=i+1}^{N-1} J_{ij} Z_i Z_j \quad (1)$$

defined on N spin variables. For simplicity we will usually take $J_{ij} \in \{-1, 1\}$, but our scaling results just rely on the fact that the couplings are taken from a symmetric distribution. This model is important in physics because it accurately describes glassy phases of matter. It is however also a model of strong mathematical interest because some exact results can be derived [26, 27]. However, finding the ground state of large spin-glass instances belongs to the class of NP-hard problems, meaning that there is no quantum or classical algorithm expected that can find the exact ground state in polynomial time. Recently however, under the assumption that this model has full replica-symmetry breaking, a very efficient approximate message-passing algorithm has been developed in Ref. [28]. This algorithm finds a solution that is $1 - \epsilon$ away from the optimum with high probability. On the other hand, conventional classical semi-definite programming methods have provable, and assumption-free, average performance guarantees [29, 30, 31]. Additionally, Ref. [32] proposed a ‘quantum-inspired’ classical algorithm based on a mean-field version of QAOA, and found good average performance. On the quantum side Refs. [4, 5] show that vanilla QAOA can outperform these assumption-free conventional methods at $p = 11$. When complemented with a ‘relax-and-round’ approach, this even happens for smaller-depth QAOA, see Ref. [33].

On the other hand, the approximate message-passing algorithm discussed in Ref. [28] with complexity $f(\epsilon)N^2$, leaves little to no room for a quantum computational advantage in the sense of improving the solution quality with a polynomial-time quantum algorithm. However, an interesting open question is whether a quantum polynomial speedup can be shown convincingly. For this reason, the Sherrington-Kirkpatrick model remains an extremely interesting model to investigate with quantum approaches.

The potential quantum speedup combined with the vast amount of knowledge about this model from many different fields, make this model ideal for benchmarking purposes. In addition, it is particularly suitable for benchmarking the parity-encoded QAOA because of its all-to-all connectivity. In this case, the parity qubits have a fixed geometric structure, as we explain next.

In the parity encoding, each parity qubit represents the product of two logical qubits $Z_i Z_j \leftrightarrow \sigma_{ij}^z$, hence an edge in the original problem graph. Thus the parity encoding of the problem consists out of $K = N(N - 1)/2$ parity qubits that we will label by the combined index ij with $i < j$ [see Fig. 1]. However, because of the redundancy of this encoding, the parity qubits need to satisfy a set of constraints. It can easily be seen that on closed cycles in the logical problem graph the parity operators σ_{ij}^z , expanded in their respective logical operators, multiply up to identity. In order to have a mapping between the low-energy states of the logical Hamiltonian (1), and the low-energy states of the encoded Hamiltonian, a subset of closed cycles needs to be added to the encoded Hamiltonian as penalty terms. The number of closed cycles in the problem graph is exponential $\mathcal{O}(N^N)$, but it is sufficient to pick a subset of cycles that forms a cycle basis. (Note that for an all-to-all connected N -node graph a cycle basis consist out of $P = (N - 1)(N - 2)/2$ elements.) A particular convenient choice of cycle basis can be represented by the triangular layout on a square lattice, where the square or triangular plaquettes form the elements of the cycle basis. This is illustrated in Fig. 1 for $N = 6$ logical qubits¹. This structure is also referred to as LHZ triangle, after the authors of the original work in which this mapping was introduced, see Ref. [8].

Under the parity mapping, the problem Hamiltonian (1) then takes the geometrically local Parity-Encoded (PE) form

$$H_{PE} = \sum_{i=0}^{N-2} \sum_{j=i+1}^{N-1} J_{ij} \sigma_{ij}^z - \sum_{i=0}^{N-3} \sum_{j=i+1}^{N-2} C_{ij} \sigma_{ij}^z \sigma_{ij+1}^z \sigma_{i+1j+1}^z [\sigma_{i+1j}^z] \quad (2)$$

The first term encodes the problem in the coupling strengths J_{ij} , the second term are energy penalties for unsatisfied constraints. While these can in principle be fine-tuned for each plaquette, we will simply put $C_{ij} \equiv C > 0$ in the remaining of this paper.

The structure of the QAOA circuit corresponding to this Hamiltonian is discussed in Ref. [9]. In particular, each layer of the QAOA contains three unitaries, as there are the constraints in addition to the conventional ‘problem’ and ‘mixer’ layers. The parity QAOA ansatz is given by

$$|\psi(\gamma, \beta, \Omega)\rangle = U_x(\beta_p) U_z(\gamma_p) U_c(\Omega_p) \dots U_x(\beta_1) U_z(\gamma_1) U_c(\Omega_1) |+\rangle \quad (3)$$

¹Note that one could find a more compact cycle basis, the so-called fundamental cycle basis. A cycle basis is defined as fundamental if and only if every basis cycle contains an edge that is not included in any other cycle of the basis. However, these fundamental cycles can not be arranged in a convenient lattice structure.




Figure 2: A ‘qubit line’ in the LHZ triangle is connecting parity qubits that each have one repeated index. Naturally, there are N such qubit lines, and each of them corresponds to a spanning tree of fixed topology in the logical picture. Here, as an example, we show two such qubit lines with their corresponding spanning trees in blue and orange.

where p is the number of layers and $|+\rangle = \bigotimes_{i=0}^{K-1} \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$. The unitaries are defined as

$$\begin{aligned} U_x(\beta) &= \prod_{i<j} e^{-i\beta\sigma_{ij}^x}, \\ U_z(\gamma) &= \prod_{i<j} e^{-i\gamma J_{ij}\sigma_{ij}^z}, \\ U_c(\Omega) &= \prod_{i<j} e^{-i\Omega\sigma_{ij}^z\sigma_{ij+1}^z\sigma_{i+1j+1}^z\sigma_{i+1j}^z}. \end{aligned} \quad (4)$$

We note here that both the mixer and problem layers are strictly local. The constraint layer contains geometrically local three- or four-body interactions, and generates entanglement. Conventionally, such a single four-body gate is decomposed as a depth-6 CNOT-ladder with an Rz rotation in the middle [9]. However, for a full layer of these operators applied to qubits arranged on a square lattice, smaller depth decompositions are possible, see Ref. [34].

The parametric angles need to be classically optimized in such a way that the energy

$$E(\gamma, \beta, \Omega) = \langle \psi(\gamma, \beta, \Omega) | H_{PE} | \psi(\gamma, \beta, \Omega) \rangle$$

is minimized. The angles satisfy certain symmetries that are discussed in Appendix A. However, this energy does not directly relate to the energy of a decoded bitstring, i.e. we can not directly associate a logical ‘solution’ bitstring to a measured parity-encoded bitstring, as the parity-encoded bitstring will generally not satisfy all constraints. Indeed, we start from the $|+\rangle$ state which on average violates half of the constraints, and apply a QAOA circuit that in principle allows for the full exploration of the parity-encoded Hilbert space.

Applying a constraint-fixing procedure to explicitly correct broken constraints in the measured output of the algorithm, can potentially be computationally demanding. Furthermore, to understand the performance of the quantum algorithm, it is preferred not to add any expensive post-processing steps. Otherwise, they could dominate the scaling of the algorithm.

Therefore, to measure the performance of the algorithm, we will evaluate the cost function [10]

$$H_{ST} = \frac{1}{N} \sum_{t=0}^{N-1} \sum_{i<j} J_{ij} \sigma_{it}^z \sigma_{ij}^z, \quad (5)$$

on the output state. Here we use the notation $\sigma_{tj}^z \equiv \sigma_{jt}^z$, and $\sigma_{jj}^z = 1$. This cost function evaluates the mean logical energy of decoding along the ‘qubit lines’, which are the lines in the LHZ triangle with a common index. There are N such lines, and they correspond to spanning trees in the logical problem, as sketched in Fig. 2. Thus if we pick such a spanning tree t at random, we can associate a logical bitstring with a measured parity-encoded bitstring $\mathbf{z}$ in the following way: (i) We fix the logical qubit t , because of the $\mathbb{Z}_2$ symmetry of (1) this is an arbitrary choice. (ii) We fix the remainder of logical qubits $i \neq t$ according to $\sigma_{it}^z = \sigma_{ti}^z$ measured in the parity bitstring. For example, if we choose the logical qubit t to be in the ‘1’ state, and we have the measurement outcome $\langle \mathbf{z} | \sigma_{it}^z | \mathbf{z} \rangle = -1$ in the parity bitstring, we then have that the i th qubit is in the ‘-1’ state in the logical bitstring. In this way, we can associate N logical bitstrings, each corresponding to one of the qubit line spanning trees, to the parity bitstring $\mathbf{z}$. If no constraints are broken in $\mathbf{z}$, these N logical bitstrings are the same.




Figure 3: (A) The bulk reverse causal cone of the onsite operator. The thick lines represent the boundary of the reverse causal cones, respectively shown for $p = 1$ (black) and $p = 2$ (blue). These reverse causal cones do not depend on the coupling strengths (the J 's) associated with the qubits on the boundary, but only on those excluding the boundary. We indicate this reduced dependence by coloring the qubits for which the coupling strength is relevant. E.g. for $p = 1$ the σ_{ij}^z RCC depends only on the coupling strength of the parity qubit indicated in black, for $p = 2$ it depends on those indicated in black and blue, and so on. (B) The bulk reverse causal cone of the plaquette operator. (C) Sketch of the LHZ triangle for $N = 10$ with the $t = 5$ qubit line indicated in black. For this qubit line, the sites ij for which the RCC of $\sigma_{i5}^z \sigma_{5j}^z$ depends on J_{ij} are indicated by the color scheme as a function of p .

Following this procedure, Eq. (5) gives the average logical energy of decoding along N qubit lines. Broken constraints will lead to inconsistent logical bitstrings along different qubit lines. However, the exact ground state of Eq. (5) encodes the logical solution, and decoding along any spanning tree will lead to this solution.

The cost function (5) originates from Ref. [10], we refer to this work for further details. However, as we are interested in the average scaling of the algorithm, we will not explicitly sample bitstrings along qubit lines (or along more general spanning trees in the logical graph).

In this paper, we define the performance ratio of parity-encoded QAOA as

$$\mathcal{P}_{ST}(\gamma, \beta, \Omega) = - \frac{\langle \psi(\gamma, \beta, \Omega) | H_{ST} | \psi(\gamma, \beta, \Omega) \rangle}{N^{3/2}}. \quad (6)$$

This quantity corresponds to the negative energy density of the logical problem. The renormalized logical Hamiltonian $\tilde{H}_P = \frac{1}{\sqrt{N}} H_P$ is extensive (i.e. its ground state energy scales linearly with N), thus further dividing it by N gives the energy density. This explains the factor $N^{3/2}$ in the denominator. This quantity converges to the Parisi constant in the large N limit

$$\lim_{N \rightarrow \infty} \mathcal{P}_{ST}^0 = - \frac{\langle \psi_0 | H_{ST} | \psi_0 \rangle}{N^{3/2}} = \epsilon_P = 0.763 \dots, \quad (7)$$

for the ground state $|\psi_0\rangle$ of typical instances. Note that this is then also trivially true for the average $\mathbb{E}_J(\mathcal{P}_{ST}^0)$ over all instances. Throughout this paper, we will mostly be interested in computing $\mathbb{E}_J(\mathcal{P}_{ST}(\gamma, \beta, \Omega))$.

3 Bounding the performance of parity-encoded QAOA

In this section, we discuss our analytic results in evaluating the performance of parity-encoded QAOA. In the spirit of Ref. [4], we will evaluate the value of the cost function averaged over all

problem instances. Due to the lengthy calculations involved, we have only done this explicitly for $p = 1$. However, we can derive asymptotically tight upper bounds on the performance (6) of parity-encoded QAOA for arbitrary p .

Because of the geometric locality of the parity-encoded algorithm [see Eq. (4)], we can perform the calculations on the light cones or ‘Reverse Causal Cones’ (RCCs) of the observables in the Hamiltonian [see Fig. 3]. In case p is small enough compared to N , i.e. when $p < N - 1$, these RCCs do not span the whole system. This RCC reduction is due to unitary cancellations of gates that commute with the observables. The RCC concept was already included in the original QAOA paper [3], but is for standard QAOA applied on spin-glass models not very useful [4], as each RCC spans the whole system because of the all-to-all connectivity.

3.1 Results for $p = 1$

For $p = 1$ parity-encoded QAOA we can derive the following result

$$\mathbb{E}_J (\langle \psi(\gamma, \beta, \Omega) | H_{PE} | \psi(\gamma, \beta, \Omega) \rangle) \approx K \cos^4(2\Omega) \sin(2\beta) \sin(2\gamma) - 4 PC \cos(2\gamma) \sin(2\Omega) \cos^3(2\Omega) \sin(2\beta) \cos(2\beta) (\cos^2(2\beta) - \cos^2(2\gamma) \cos^2(2\Omega) \sin^2(2\beta)), \quad (8)$$

where $P = (N - 1)(N - 2)/2$ is the number of constraints, or plaquette terms in Eq. (2). The full derivation is included in Appendix B. Here, the approximation assumes that all reverse causal cones have the same square shape as shown in Fig. 3(A,B). This is true for the bulk of the reverse causal cones $\mathcal{O}(N^2)$. Boundary effects $\mathcal{O}(N)$ because of different shapes can be neglected. Concise expressions for $p = 1$ vanilla QAOA have also been derived for various problems including spin glasses [35].

Eq. (8) can be minimized as a function of the angles. In Fig. 4, we show the resulting optimal angles as a function of the constraint strength C , for large N . As expected, the constraint strength should be taken larger than the gap to obtain non-trivial angles. For $J_{ij} \in \{-1, 1\}$, the gap is 2 for non-degenerate problems. This corresponds to the energy difference induced by flipping spins on an edge in the ground state configuration. We also observe that if the constraint strength is taken very large, the problem angle becomes trivial. This indicates that the algorithm gradually forgets about the specific problem it is supposed to solve, and only tries to satisfy the constraints.

Next, we want to calculate the average optimal performance $\mathbb{E}_J(\mathcal{P}_{ST})$, which boils down to the calculation of $\mathbb{E}_J (\langle \psi(\gamma^*, \beta^*, \Omega^*) | H_{ST} | \psi(\gamma^*, \beta^*, \Omega^*) \rangle)$ with the optimal parameters γ^* , β^* and Ω^* . For this we start by rewriting H_{ST} (5) in the following form

$$H_{ST} = \frac{1}{N} \left(2 \sum_{i < j} J_{ij} \sigma_{ij}^z + \sum_t \sum_{\substack{i < j \\ i \neq t \\ j \neq t}} J_{ij} \sigma_{it}^z \sigma_{tj}^z \right), \quad (9)$$

such that the terms are explicitly separated in single-body terms and two-body terms. However, for $p = 1$ the gates that are present in the reverse causal cone of the two-body term $\sigma_{it}^z \sigma_{tj}^z$ are independent of J_{ij} . Indeed, the coupling strength J_{ij} does not enter in the $p = 1$ RCC of $\sigma_{it}^z \sigma_{tj}^z$, which follows from unitary cancellations, originating from the fact that $[\sigma_{it}^x, \sigma_{ij}^z] = 0$ when $t \neq j$ (and that $[\sigma_{it}^z, \sigma_{ij}^z] = 0, \forall j$). Therefore, we have that

$$\mathbb{E}_J (J_{ij} \langle \psi(\gamma, \beta, \Omega) | \sigma_{it}^z \sigma_{tj}^z | \psi(\gamma, \beta, \Omega) \rangle) = \mathbb{E}_J(J_{ij}) \mathbb{E}_J (\langle \psi(\gamma, \beta, \Omega) | \sigma_{it}^z \sigma_{tj}^z | \psi(\gamma, \beta, \Omega) \rangle) = 0,$$

as for spin glasses the couplings are taken from a symmetric distribution around zero. This implies that only the local part in Eq. (9) is contributing. These local terms are calculated in the first part of Appendix B. Therefore,

$$\mathbb{E}_J (\langle \psi(\gamma, \beta, \Omega) | H_{ST} | \psi(\gamma, \beta, \Omega) \rangle) \approx 2 \frac{K}{N} \cos^4(2\Omega) \sin(2\beta) \sin(2\gamma) \quad (10)$$

where the approximation means the same as before, i.e. that the topology of all the RCCs is assumed to be the one of the bulk, shown in Fig. 3(A) for $p = 1$ in blue. The optimal angles,




Figure 4: **Optimal parity-encoded QAOA angles** for $p = 1$. There is a transition from trivial to non-trivial angles as the constraint strength becomes larger than the gap. Here the couplings are $J_{ij} \in \{-1, 1\}$. In the large C limit, the angle that drives the problem Hamiltonian vanishes.

minimizing Eq. (10), are then the trivial angles $\beta^* = -\gamma^* = \pi/4$ and $\Omega^* = 0$. Notice that these angles are also optimal when minimizing the energy of the parity-encoded Hamiltonian, see Eq. (8), for small constraints strengths $C < 2$ and large N (see also Fig. 4). The latter implies that the algorithm does not generate any entanglement, but only performs local rotations. These local rotations map the state $|+\rangle$ back to a product state in the z -basis where the state of each qubit ij is set by J_{ij} (in case $J_{ij} \in \{-1, 1\}$). Concretely, $J_{ij} = 1 \rightarrow |0\rangle$ and $J_{ij} = -1 \rightarrow |1\rangle$. Because of the fact that there are no correlations generated, the algorithm has vanishing performance

$$\mathbb{E}_J(\mathcal{P}_{ST}) = \frac{1}{N^{1/2}} - \frac{1}{N^{3/2}} \quad \text{for } p = 1. \quad (11)$$

This result means that the performance of the $p = 1$ algorithm unavoidably drops to random sampling as a function of N . From this it can already be anticipated that p will necessarily need to grow with N in order to obtain a better performance.

3.2 Larger depth $p > 1$

Now, we will derive an upper bound on the performance for general fixed p . This upper bound will follow from counting the number of non-vanishing 2-body RCCs. In particular from (9), it follows that

$$\mathbb{E}_J(\mathcal{P}_{ST}) \leq \frac{1}{N^{1/2}} - \frac{1}{N^{3/2}} + \frac{\# \text{ of non-vanishing 2-body RCCs}}{N^{5/2}}, \quad (12)$$

with the number of non-vanishing reverse causal cones depending on p . Here the first two terms are equivalent to $2K/N^{5/2}$, which corresponds to the maximal weight of the onsite terms in Eq. (9). As we have shown in the previous paragraph, this is optimal when $p = 1$, as in this case *all* 2-body RCCs do not contribute. Here, we will show that for fixed p

$$\mathbb{E}_J(\mathcal{P}_{ST}) \leq \frac{1}{N^{1/2}} + \mathcal{O}(N^{-3/2}) \quad \forall p \text{ if } N > N^*(p), \quad (13)$$

hence that non vanishing 2-body RCCs only contribute to the subleading $\mathcal{O}(N^{-3/2})$ corrections. Concretely, in order to show (13), we will show that the number of non-vanishing 2-body RCCs only grows linearly with N for fixed p .

On the other hand if p is taken to be at least of the order of N , the number of non-vanishing two-body reverse causal cones will scale as $\mathcal{O}(N^3)$.

For general p , a 2-body term ($i \neq t$ and $j \neq t$)

$$\mathbb{E}_J(J_{ij} \langle \psi(\gamma, \beta, \Omega) | \sigma_{it}^z \sigma_{tj}^z | \psi(\gamma, \beta, \Omega) \rangle)$$

will not contribute if, as before, J_{ij} is not part of the reverse causal cone of $\sigma_{it}^z \sigma_{tj}^z$. The number of terms for which J_{ij} is part of the reverse causal cone of $\sigma_{it}^z \sigma_{tj}^z$ scales as Np for fixed t . This is

illustrated in Fig. 3(C), where for the qubit line $t = 5$ the colored sites ij are inside the reverse causal cone of $\sigma_{it}^z \sigma_{tj}^z$ for a given p . Because of this argument the number of non vanishing terms (summing over all t 's as in Eq. (9)) scales at most as $N^2 p$. Therefore, this is already sufficient to conclude that $\mathbb{E}_J(\mathcal{P}_{ST})$ will decrease proportional to $N^{-1/2}$ for fixed p .

However, we can also fix the prefactor of the decay by realizing that for fixed p actually more 2-body RCCs are vanishing, thus showing that interaction terms do not play a role asymptotically. Even if J_{ij} is part of the reverse causal cone of $\sigma_{it}^z \sigma_{tj}^z$, it is likely (and in fact the bulk contribution, when $N \gg p$) that the reverse causal cones of σ_{it}^z and σ_{tj}^z are not mutually overlapping. This would imply that the expectation value factorizes

$$J_{ij} \langle \psi(\gamma, \beta, \Omega) | \sigma_{it}^z \sigma_{tj}^z | \psi(\gamma, \beta, \Omega) \rangle = J_{ij} \langle \psi(\gamma, \beta, \Omega) | \sigma_{it}^z | \psi(\gamma, \beta, \Omega) \rangle \langle \psi(\gamma, \beta, \Omega) | \sigma_{tj}^z | \psi(\gamma, \beta, \Omega) \rangle.$$

If this is the case, we can also show that $\mathbb{E}_J(J_{ij} \langle \psi(\gamma, \beta, \Omega) | \sigma_{it}^z | \psi(\gamma, \beta, \Omega) \rangle) = 0$, assuming the RCC of σ_{it}^z is dependent on J_{ij} . We elaborate on this in Appendix C.

In Fig. 5, we show the count of the total number of non-vanishing 2-body reverse causal cones. For fixed p , when $N > N^* = 3p - 1$ the number of those non-vanishing contributions starts growing only linearly with N (the value of N^* is a numerical observation from Fig. 5(A)). Therefore, the two-body terms in Eq. (9) do not even contribute to a prefactor in the decay in the large N limit, but only to the subleading $\mathcal{O}(N^{-3/2})$ corrections in the performance (13). Indeed, in Fig. 5(B) we plot the right-hand side of Eq. (12), illustrating the vanishing performance with $N^{-1/2}$ as N becomes large enough for fixed p . In this plot it can also be seen that when $N > N^* = 3p - 1$ the upper bound starts decreasing, but it only drops below the Parisi constant for $N \geq N_P$ where N_P depends on p .

We comment that the upper bound is based on a counting argument, where all non vanishing terms are counted, and thus all contribute with their maximal individual weight (i.e. 1). When N is roughly the same order as N^* , the upper bound is not physical as it is above the Parisi value. This is indicated by the shaded grey area in Fig. 5(B). Therefore, the bound should only be considered meaningful when it drops below the Parisi value ϵ_P for larger system sizes.

In Fig. 5(C), we show the values of N_P for each p and fit the exponent on how p should at least grow with problem size to allow for having maximal performance. From this, the necessary requirement to allow for maximal performance is that p should grow with system size, at least as N^q with $q \gtrsim 0.67$, i.e. as $\Omega(N^{0.67})$. At the same time, we have seen that all terms start contributing only when $p \geq N - 1$ [see Fig. 5(A)]. This suggest that it is likely that p will need to grow faster than $N^{0.67}$, e.g. linearly, in order to obtain fixed performance in practice.

3.3 Scaling of the variance

We note that the previous arguments so far only discussed the average over all instances. This does in principle not exclude the existence of an ‘easier’ subset of problem instances for which the performance is non vanishing in the large N limit for fixed p . To rule out such non-generic side effects, we would need to prove that the variance of the performance over problem instances is also vanishing. Similarly as for the average, we can construct a vanishing upper bound on the (scaling of the) variance

$$\text{Var}_J(\mathcal{P}_{ST}) = \mathbb{E}_J(\mathcal{P}_{ST}^2) - [\mathbb{E}_J(\mathcal{P}_{ST})]^2.$$

From the previous we already know that the last term is vanishing as $N^{-1} + \mathcal{O}(N^{-2})$, therefore we will focus on the first term. The leading contribution to understand the scaling is

$$\sum_t \sum_{\substack{i < j \\ i \neq t \\ j \neq t}} \sum_{t'} \sum_{\substack{i' < j' \\ i' \neq t' \\ j' \neq t'}} \mathbb{E}_J(J_{ij} \langle \psi | \sigma_{it}^z \sigma_{tj}^z | \psi \rangle J_{i'j'} \langle \psi | \sigma_{i't'}^z \sigma_{t'j'}^z | \psi \rangle)$$

By counting the number of terms, this is upper bounded by N^6 . However, similar as before, the term will be vanishing if the RCCs of $\sigma_{it}^z \sigma_{tj}^z$ and $\sigma_{i't'}^z \sigma_{t'j'}^z$ do not depend on J_{ij} or on $J_{i'j'}$. A counting argument, as shown in Fig. 3(C), however implies that the number of these non vanishing terms will only grow as $\sim N^4$ for fixed p , when N is large enough.




Figure 5: **(A) Number of non-vanishing two-body reverse causal cones** contributing to the second term on the right-hand side in Eq. (9). For, fixed p , when $N > N^* = 3p - 1$, the growth becomes exactly linear in N . The stars indicate the value of $N^* = 3p - 1$. The maximal number, as indicated by the black line $(N - 2)K$, is only reached when $p \geq N - 1$. This is when each reverse causal cone spans the whole LHZ triangle, as the short base of the triangle contains $N - 1$ qubits. **(B) Upper bound on the performance of parity-encoded QAOA**, as given in the right-hand side of Eq. (12). As indicated by the grey area, the bounds shown here should only be considered meaningful when they drop below the physical bound, which is the Parisi constant ϵ_P in the thermodynamic limit. The dots indicate the system size N_P for which the bound is the closest to ϵ_P for each p . **(C) The depth p as a function of N_P** . We make a numerical fit of the exponent indicating that p should grow with N at least as $N^{0.67}$ to allow for maximal performance.

	Parity-encoded QAOA	Vanilla QAOA
Qubits	$N(N-1)/2$	N
Number of gates / QAOA layer	$\sim N^2$	$\sim N^2$
Depth / QAOA layer	constant	$\sim N$
Number of QAOA layers	$\sim N^q, q \gtrsim 0.67$	constant
Total number of gates	$\sim N^{2+q}, q \gtrsim 0.67$	$\sim N^2$
Total depth	$\sim N^q, q \gtrsim 0.67$	$\sim N$

Table 1: Comparison of the scaling with N for the two QAOA variants to solve Sherrington-Kirkpatrick spin-glass models with fixed performance. Vanilla QAOA is the regular version of QAOA for spin-glass models as discussed in Ref. [4]. While for vanilla QAOA the scalings contained in this table are proven, for parity-encoded QAOA they should be interpreted as minimal (necessary) requirements, e.g. the total depth of the parity-encoded algorithm must grow at least as $N^{0.67}$.

Taking into account other effects such as non-overlapping reverse causal cones of the σ^z operators, could lead to more insight in the prefactor. However, this argument is sufficient to conclude that

$$\text{Var}_J(\mathcal{P}_{ST}) = \mathbb{E}_J(\mathcal{P}_{ST}^2) - [\mathbb{E}_J(\mathcal{P}_{ST})]^2 \sim \frac{1}{N} + \mathcal{O}(N^{-2})$$

for fixed p and N large enough.

3.4 Discussion and comparison with standard QAOA

From the bounds derived in the beginning of this section, we can conclude that the mapping from the non-local optimization problem to the geometrically local one comes with strong limitations on the performance of local algorithms attempting to solve the problem. Indeed, we derived that the minimum requirement to achieve a fixed performance is to run a non-local algorithm (in the sense that the depth needs to scale as a function of N) on the parity-encoded qubits. However, the non-local parity-encoded algorithm needs to act on $K = N(N-1)/2$ qubits. In contrast to vanilla QAOA, which is also a non-local algorithm but which acts on only N qubits, and furthermore has a fixed performance with the number of vanilla QAOA layers [4].

To summarize our findings, in Table. 1, we compare the scaling with N for parity-encoded QAOA with the standard vanilla QAOA to get a fixed performance ratio for solving spin-glass models. As demonstrated in Ref. [4], the standard QAOA implementation requires only fixed p to get fixed performance in the large N limit. The caveat, however, is that the physical depth of each QAOA layer will scale with N [6]. This is because of the assumption, which is true in all currently available hardware platforms, that we cannot perform all two-body gates in parallel. However, for parity-encoded QAOA the physical depth of each layer is constant. On the other hand, as we have discussed previously, the number of layers for parity-encoded QAOA must grow with system size in order to have non vanishing performance. From the upper bound on the performance that is based on the counting argument, we find an exponent of $q \approx 0.67$ to allow for maximal performance [see also Fig. 14(C)]. This formally leaves room for a speedup with respect to vanilla QAOA, in the sense of having the possibility to get the same average performance ratios with a polynomially smaller circuit. Although this is not guaranteed and it would still need to be proven that such performances can be achieved in practice.

As there is a quadratic qubit overhead in the parity-encoded QAOA, the total number of gates needed for this algorithm scales less favorably than for vanilla QAOA. However, we have not explicitly added noise to the circuit in order to investigate whether there could be some algorithmic robustness under noise.

We comment that for small system sizes the scalings presented in Table 1 might not be noticeable as such, because of the presence of finite-size effect. The SWAP-network routing that would be necessary for Vanilla QAOA [6], adds only a prefactor to the scaling of the depth of a QAOA layer, and therefore does not alter the scaling which is $\mathcal{O}(N)$.




Figure 6: **Numerical performance ratios with the angles optimized instance-by-instance** to minimize the spanning-tree cost function Eq. (5). The averages are taken over 200 instances, and the error bars show three times the standard error of the mean. Histograms over the found angles are shown in Fig. 7. The performance ratio remains very close to the optimal $p = 1$ analytic result [see Eq. (11)], indicated by the thick black line, even when increasing p . The exact solution is given by the dashed black line.

4 Tensor-network calculations

In this section, we study the parity-encoded QAOA algorithm numerically, and confirm the predictions from the previous section. We perform exact calculations by representing the state $\rho(\gamma, \beta, \Omega) = |\psi(\gamma, \beta, \Omega)\rangle \langle \psi(\gamma, \beta, \Omega)|$ as a tensor network using the QUIMB package [36]. In the tensor-network picture, all the gates are represented by tensors with ‘input’ and ‘output’ indices for each qubit it acts on [11]. Then, we calculate each observable that is present in H_{ST} or H_{PE} as a tensor-network contraction over the corresponding reverse causal cone of the Hamiltonian terms. For more details about the tensor-network calculations and for a pseudocode, we refer to Appendix D.

4.1 Optimized angles for H_{ST}

We start by optimizing all the QAOA angles on an instance-by-instance basis in order to minimize $\langle \psi(\gamma, \beta, \Omega) | H_{ST} | \psi(\gamma, \beta, \Omega) \rangle$, and thus maximize the performance (6). For the optimization, we use the ‘L-BFGS-B’ implementation of SciPy. This is an optimizer with bounded parameter intervals, therefore we restrict the angle search space to symmetry determined ranges [see Appendix A]. Within these ranges, we initialize the parameters uniformly random.

In Fig. 6, we plot the numerically determined performance ratio averaged over 200 sampled instances. We see that, in particular for $p = 1$, this matches exactly the result (11). However, also for higher $p > 1$ it does not improve much beyond this result, consistent with (13). However, also for very small system sizes for which all the RCCs could in principle contribute, we do not obtain better results. This is e.g. the case for $N = 4$ and $p = 3$. In this figure, we also show the exact solution (dashed black lines) for comparison. In the thermodynamic limit, this converges to the Parisi constant, see Eq. (7). However, because of the quadratic qubit overhead of the parity encoding and the underlying structure of the RCCs, we are numerically limited to very small problem sizes. For such small system sizes, there are strong finite size effects. For instance, one such effect is the seemingly different behavior for even and odd system sizes, as is clearly visible in the curve of the exact solution.

In Fig. 7, we show the histograms over the optimal angles found for $N = 7$ at various parity-encoded QAOA depths. Clearly, these angles concentrate around specific values. For $p = 1$, they exactly match the predicted angles ($\gamma = -\beta = \pi/4$ and $\Omega = 0$) even on an instance-by-instance basis. This is because of self-averaging over the coupling constants J_{ij} that define a problem instance. However, also for $p = 3$ the optimizer still finds values for all Ω ’s that are near zero for many instances. In those cases, the algorithm does not generate any entanglement, and is thus strictly local. This also explains why the performance on average is not much better than




Figure 7: **Histograms over the optimal QAOA angles, minimizing the cost function (5) for $N = 7$ and 200 uniformly sampled instances.** (A) $p = 1$, the angles of the first layer are indicated in blue, (B) $p = 2$, angles of the second layer shown in orange, and (C) $p = 3$, angles of the third layer shown in green.

$N^{-1/2} - N^{-3/2}$, which is the best average performance when only local rotations are allowed ($p = 1$).

By Eq. (13), we expect that the performance will drop as in the $p = 1$ case for any p , as long as N is large enough. (And that therefore all optimal Ω 's will cluster around zero.) However, Fig. 6 already shows this behavior for the smallest N , for which all RCCs could contribute non trivially. This is somewhat surprising, but however suggest that the parameter landscape has a broad (and perhaps only local) minimum for trivial Ω 's, that forms a large basin of attraction for the optimizer.

4.2 Trotter angles

In order to check that we can do better than the $p = 1$ result in a consistent way, we now take the QAOA angles according to the following Trotter scheme

$$\begin{aligned}\gamma_i &= \frac{i - 1/2}{p} \frac{t_{\max}}{p} \\ \beta_i &= - \left(1 - \frac{i - 1/2}{p} \right) \frac{t_{\max}}{p} \\ \Omega_i &= - \frac{i - 1/2}{p} \frac{t_{\max}}{p}.\end{aligned}\tag{14}$$

Similar schemes have been considered for standard QAOA to initialize the optimizers [23]. Here, $t_{\max}$ is a parameter that we tuned for each p (and N) in order to maximize the spanning-tree performance. However, we found fairly flat curves in the performance when varying this parameter. Hence, we do not believe that the presented results are very sensitive with respect to this parameter. Note that the assumed sign difference between γ_i and Ω_i originates from the fact that unsatisfied constraints give energy penalties [see Eq. (2)].

In Fig. 8, we show the performance of this scheme averaged over some randomly sampled instances. Of course, as the choices of the angles are not optimal here, the performance lies below the $p = 1$ result derived in the previous section [see Eq. (11)]. However, this scheme is guaranteed to




Figure 8: **Numerical performance ratios with angles taken according to the Trotter scheme (14).** The parameter $t_{\max}$ has been roughly optimized to $[0.8, 1.4, 2.0, 2.6, 3.2, 3.8, 4.4, 5.0, 5.6]$ in units of J for $p = 1, 2, \dots, 9$ and $N = 6$ (with slight variations for different N). These ratios are compared to (i) the optimal $p = 1$ analytical result (thick black line), and (ii) the ratio of the exact solution (dashed black line). The averages are taken over 200 instances, and the error bars show three times the standard error of the mean.

improve the performance as a function of p . We however found that we need to increase to around $p = 9$ to see a first improvement for small instances with respect to the $p = 1$ algorithm, which is a trivial classical algorithm as it performs only local rotations and generates no entanglement.

4.3 Optimized angles for H_{PE}

Finally, we also show some numerically optimized angles that minimize the energy of the Hamiltonian (2), which contains explicit energy penalties for broken constraints. For $p = 1$, we can find the optimal angles from minimizing Eq. (8) [see also Fig. 4], or find them on an instance-by-instance basis by numerical evaluation. Again we use the ‘L-BFGS-B’ algorithm from SciPy for this with the angle intervals restricted as in Appendix A and random initialization. For $p = 2$, we show the histograms over the numerically found angles in Fig. 9 for the constraint strength $C = 3$. Here, we observe that the clusters of optimal angles grow stronger as the problem size increases, and that some finite-size satellite peaks are disappearing. In particular, the optimal Ω_1 , shown in blue, clearly shifts away from 0 to a non-trivial value. This is expected for two reasons. First, for larger system sizes the probability of sampling an ‘easy’ problem (with no, or very few frustration) decreases. Second, contrary to the spanning-tree cost function Eq. (5) for which many two-body terms vanish, the multi-body RCCs do not vanish now. This leads to a QAOA circuit that generates entanglement.

However, in terms of performance in the spanning tree picture, these angles perform on average worse than $N^{-1/2} - N^{-3/2}$ for small p . We show this in Fig. 10 up to $p = 3$ for small system sizes. For $p = 1$, this behavior is expected as the trivial angles are optimal and lead to $N^{-1/2} - N^{-3/2}$ when maximizing the spanning tree performance [see Eq. (10)]. However, because we minimized the energy of Eq. (2) with $C = 3$, we find non-trivial angles that thus lead to a performance lower than $N^{-1/2} - N^{-3/2}$. Upon increasing p , improvement is guaranteed (see also the previous paragraph). However, we still do not improve on $N^{-1/2} - N^{-3/2}$, even for $p = 3$. This result shows that after a small number of layers, the state did not satisfy enough constraints, and did not enter in a sufficiently low-energy regime for which the spectra of both Hamiltonians (2) and (5) become compatible.

4.4 Comparison to vanilla QAOA

For completeness, we also report the results of vanilla QAOA for the same small system sizes. Here we thus optimize the ansatz defined on N qubits,

$$|\psi(\beta', \gamma')\rangle = U_M(\beta'_p)U_P(\gamma'_p) \dots U_M(\beta'_1)U_P(\gamma'_1)|+\rangle$$




Figure 9: **Histograms over the optimal QAOA angles, minimizing the cost function (2), for $p = 2$ QAOA layers.** The constraint strength is taken $C = 3$. (A-C) Logical system sizes $N = 6, 9, 12$ and $N = 12$, corresponding respectively to $K = 15, 36, 66$ parity qubits. The peaks grow stronger as the system size increases.




Figure 10: **Numerical performance ratios with the angles optimized instance-by-instance, to minimize the parity-encoded cost function (2).** The constraint strength in this cost function $C = 3$. The averages are taken over 100 spin-glass instances, and the errorbars show three times the standard error of the mean.




Figure 11: **Vanilla QAOA performance ratios.** These ratios are fixed as $N \rightarrow \infty$, there are only minor finite-size effects.

with $U_M(\beta') = e^{-i\beta' \sum_{i=0}^{N-1} X_i}$, $U_P(\gamma') = e^{-i\gamma' H_P}$ with H_P given by Eq. (1), and $|+\rangle = \bigotimes_{i=0}^{N-1} \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$. In this case, the RCC concept is not very useful as the QAOA, by the problem definition, ‘sees already the whole graph’ at any layer [37, 38].

In Fig. 11 we show the corresponding performance metric

$$\mathcal{P}(\gamma', \beta') = -\frac{\langle \psi(\gamma', \beta') | H_P | \psi(\gamma', \beta') \rangle}{N^{3/2}},$$

averaged over some instances. These results can be directly compared to the parity results shown in Figs. 6, 8 and 10. For comparison we also plot $N^{-1/2} - N^{-3/2}$ here. While the parity results in the aforementioned figures struggle to overcome $N^{-1/2} - N^{-3/2}$, the vanilla QAOA performs better than this for $p' = 2$ for all N and for $p = 1$ for $N \gtrsim 9$.

Vanilla QAOA has proven performance guarantees for fixed p' in the large N limit. For $p' = 1, 2, 3$ these asymptotic values are respectively 0.30, 0.41, 0.47, and can be found in Ref. [4]. This is thus very close to what we found numerically for small system sizes, and allows us to conclude that there are not much finite-size effects in the performance of vanilla QAOA. This should be contrasted with the parity-encoded algorithm for which we have seen that the performance drops asymptotically as $N^{-1/2}$ for any fixed p .

5 Conclusion and outlook

In this paper, we have investigated the scaling of parity-encoded QAOA for paradigmatic spin-glass models. To understand the quality of the algorithm, we have calculated the average logical energy of decoding along a qubit-line spanning tree. For fixed depth $p = 1$, the performance vanishes exactly as $N^{-1/2} - N^{-3/2}$. By counting the number of non-vanishing 2-body RCCs, we show that the performance is also asymptotically vanishing as $N^{-1/2}$ for arbitrary p . This leads to the conclusion that p has to scale up with system size in order to get fixed performance. Our upper bound based on these RCCs counting arguments [see Eq. (12)], sets the requirement that p needs to scale at least as N^q with $q \gtrsim 0.67$ in order to allow for maximal performance.

Formally, this result does not exclude a speed-up of parity-encoded QAOA with respect to vanilla QAOA. This is because the total depth of the algorithm could potentially be more favorable [see Table 1]. However, this would require that increasing p actually leads to increased performance. We performed small-scale numerical experiments, and found it in practice hard to obtain better performances than the $p = 1$ algorithm [see Figs. 6, 8, 10], which performs local rotations and is therefore a trivial and entirely classical algorithm. This follows from the fact that the entangling gate drops out in the parity-encoded QAOA circuit in the $p = 1$ case (4). Therefore we are left with only single-qubit rotations that in this case rotate the parity qubits to the sign of their associated

local field J_{ij} (or the interaction between the logical qubits i and j). We however find numerically that the parity-encoded algorithm struggles to even outperform this trivial algorithm when p is increased.

As for all variational algorithms, it is uncertain in the general case that the numerical optimization procedure always yields the globally optimal parameters. However, for some cases we considered here, we could at least approximately compare to analytical formulas (e.g. in the $p = 1$ case). In these cases, we found good correspondence between the parameters found by numerical optimization (on an instance-by-instance basis) and the parameters obtained from minimizing formulas that analytically average over all instances. Such a correspondence can also be understood from the fact that spin glasses are ‘self-averaging’ in the sense averaging over the energy over instances is equivalent to averaging different energy contributions within a single ‘typical’ instance. This makes us confident that our numerical results presented in Secs. 4.1 and 4.3 are not dependent on the chosen numerical optimization procedure.

In this work, we have always considered average performances based on energy expectation values, and we have thus not sampled candidate solutions explicitly. Typically, sampling is also not straightforward from a tensor-network state that is not a matrix-product state. When candidate solutions are however sampled by e.g. state-vector simulation or actual device runs, in principle a post-selection procedure could be applied to only return the candidate solution with the lowest energy. Checking the energy of a candidate solution in the logical picture comes with a classical computational cost of $\mathcal{O}(N^2)$, corresponding to a matrix-vector inner product. Although this may seem like a small computational cost, it is non negligible and, in fact, even overrules the cost scaling (or the depth) of the quantum algorithms considered, see Table 1. In the parity-encoded picture, the cost of post-selecting on the best logical solution increases to $\mathcal{O}(N^3)$ because we consider the decoding according to N spanning trees leading to N candidate solutions in the logical picture that need to be checked for every sample from the parity-encoded circuit. We however believe that due to the concentration results of both vanilla QAOA [see Ref. [4]] and parity-encoded QAOA [see 3.3], the best solution will not deviate much from the mean, at least for larger system sizes. For small system sizes, there might be finite-size effects due to post-selection that go however beyond the scope of our tensor-network investigation. We conclude that the parity-encoded QAOA is not showing promising results when compared to vanilla QAOA.

We leave open if the parity-encoding could work better for different problem classes. However, our work indicates that there is little hope of finding good quality solutions of an originally non-local optimization problem by solving it with a geometrically local reformulation and geometrically local algorithm. We expect that our arguments of RCCs that factorize and eventually drop out, can also be applied on different problems and could also be used to derive performance bounds for other algorithms.

6 Acknowledgements

We thank Glen Bigan Mbeng, Michael Fellner, Martin Lanthaler, Maximilian Lutz, Pedro Parrado, PV Sriluckshmy and Anita Weidinger for useful discussions. This project is supported by the Federal Ministry for Economic Affairs and Climate Action on the basis of a decision by the German Bundestag through the project Quantum-enabling Services and Tools for Industrial Applications (QuaST). QuaST aims to facilitate the access to quantum-based solutions for optimization problems.




Figure 12: (A) The $p = 1$ reverse causal cone of the onsite operator. (B) The $p = 1$ reverse causal cone of the plaquette operator. The elements of a constraint vector, e.g. k_m , should be seen as a product of the elements of a z -vector as given by the constraints, e.g. $k_m^5 = z_m^i z_m^j z_m^k z_m^l$ (we use single case site labels here).

A Angle symmetries

A angle symmetry $(\beta, \gamma, \Omega) \rightarrow (\beta', \gamma', \Omega')$ in the QAOA algorithm leaves the value of the cost function unchanged, i.e.

$$E(\beta, \gamma, \Omega) = E(\beta', \gamma', \Omega')$$

with $E(\beta, \gamma, \Omega) = \langle \psi(\beta, \gamma, \Omega) | H | \psi(\beta, \gamma, \Omega) \rangle$, where H can be either H_{PE} as in Eq. (2), or H_{ST} as in Eq. (5). The most obvious symmetry is a global ‘time-reversal’ symmetry implying that $E(\beta, \gamma, \Omega) = E(-\beta, -\gamma, -\Omega)$.

Another very general (and local) QAOA symmetry, arises from the standard form of the mixing Hamiltonian, we have that

$$e^{i(\beta_i + \pi)H_x} = \prod_{i < j} (-\cos(\beta_i) - \sin(\beta_i)\sigma_{ij}^x) = (-1)^K e^{i\beta_i H_x}. \quad (15)$$

The possible minus sign is irrelevant for the evaluation of the cost function. Therefore, the size of the range of each β_i can be restricted to half the unit circle, e.g. $(-\pi/2, \pi/2]$.

Next, the plaquette operator

$$H_c = \sum_{i=0}^{N-2} \sum_{j=i+1}^{N-2} \sigma_{ij}^z \sigma_{ij+1}^z \sigma_{i+1j+1}^z [\sigma_{i+1j}^z]$$

has eigenvalues that are either even or odd, depending on whether the number of plaquettes is even or odd. If it is odd, we have that

$$e^{i(\Omega_i + \pi)H_c} = -e^{i\Omega_i H_c},$$

and if it is even

$$e^{i(\Omega_i + \pi)H_c} = e^{i\Omega_i H_c}.$$

Therefore, we can also restrict each $\Omega_i \in (-\pi/2, \pi/2]$.

Finally, the symmetries on the γ -angles are problem specific. However with our choice of $J_{ij} \in \{-1, 1\}$, exactly the same argument as in (15) applies. So in conclusion the search range for the optimal angles can be bounded to $(-\pi/2, \pi/2]$ for every angle except one, for which the sign can be fixed in addition.

We finally note that further geometric simplifications could lead to even more restricted ranges. For example, if our original problem would map to a parity configuration that does not simultaneously contain square and triangular plaquettes. Indeed, if it would only contain squares, the plaquette operator would be $\mathbb{Z}_2$ symmetric. While if it would only contain triangles, a global spin flip leads to a global minus sign. This can then in both cases be used to restrict the angles further.

B Average expectation values for $p = 1$

B.1 Onsite term

Our goal is to calculate the term $\mathbb{E}_J \langle \psi(\gamma_1, \beta_1, \Omega_1) | J_{ij} \sigma_{ij}^z | \psi(\gamma_1, \beta_1, \Omega_1) \rangle$. Only the gates that are part of the RCC of the operators σ_{ij}^z contribute to this expectation value, and we will only consider

the bulk (i.e. square) RCC [see Fig. 12(A)]. We start by inserting identities and by grouping the phases. By $\mathbf{z}_1, \mathbf{z}_{-1}, \mathbf{z}_m$ we denote bitstrings of length $(2p+1)^2 = 9$ describing a configuration on the RCC, $z_1^{ij} \in \{-1, +1\}$, and by $\sum_{\{\mathbf{z}\}}$ we mean summing over all $(2^{3 \times 9})$ variable configurations.

$$\begin{aligned}
& \mathbb{E}_J \left(\langle \psi(\gamma_1, \beta_1, \Omega_1) | J_{ij} \sigma_{ij}^z | \psi(\gamma_1, \beta_1, \Omega_1) \rangle \right) \\
&= \mathbb{E}_J \left(\langle x | e^{iJ_{ij}\gamma_1\sigma_{ij}^z} e^{i\Omega_1 H_c} e^{i\beta_1\sigma_{ij}^x} J_{ij} \sigma_{ij}^z e^{-i\beta_1\sigma_{ij}^x} e^{-iJ_{ij}\Omega_1 H_c} e^{-iJ_{ij}\gamma_1\sigma_{ij}^z} | x \rangle \right) \\
&= \frac{1}{2^9} \mathbb{E}_J \left(\sum_{\{\mathbf{z}\}} \langle \mathbf{z}_1 | e^{iJ_{ij}\gamma_1\sigma_{ij}^z} | \mathbf{z}_1 \rangle \langle \mathbf{z}_1 | e^{i\Omega_1 H_c} | \mathbf{z}_1 \rangle \langle \mathbf{z}_1 | e^{i\beta_1\sigma_{ij}^x} | \mathbf{z}_m \rangle \langle \mathbf{z}_m | J_{ij} \sigma_{ij}^z | \mathbf{z}_m \rangle \langle \mathbf{z}_m | e^{-i\beta_1\sigma_{ij}^x} | \mathbf{z}_{-1} \rangle \right. \\
&\quad \left. \langle \mathbf{z}_{-1} | -e^{i\Omega_1 H_c} | \mathbf{z}_{-1} \rangle \langle \mathbf{z}_{-1} | e^{-iJ_{ij}\gamma_1\sigma_{ij}^z} | \mathbf{z}_{-1} \rangle \right) \\
&= \frac{1}{2^9} \sum_{\{\mathbf{z}\}} e^{i\Omega_1(H_c(\mathbf{z}_1) - H_c(\mathbf{z}_{-1}))} \langle \mathbf{z}_1 | e^{i\beta_1\sigma_{ij}^x} | \mathbf{z}_m \rangle \langle \mathbf{z}_m | e^{-i\beta_1\sigma_{ij}^x} | \mathbf{z}_{-1} \rangle \mathbb{E}_J \left(J_{ij} z_m^{ij} e^{iJ_{ij}\gamma_1(z_1^{ij} - z_{-1}^{ij})} \right)
\end{aligned}$$

At this point, we can easily perform the average $\mathbb{E}_J$, and in a second step redefine $\mathbf{z}_1 \rightarrow \mathbf{z}_1 \mathbf{z}_m$, and $\mathbf{z}_{-1} \rightarrow \mathbf{z}_{-1} \mathbf{z}_m$.

$$\begin{aligned}
& \mathbb{E}_J \left(\langle \psi(\gamma_1, \beta_1, \Omega_1) | J_{ij} \sigma_{ij}^z | \psi(\gamma_1, \beta_1, \Omega_1) \rangle \right) \\
&= \frac{1}{2^9} \sum_{\{\mathbf{z}\}} e^{i\Omega_1(H_c(\mathbf{z}_1) - H_c(\mathbf{z}_{-1}))} \langle \mathbf{z}_1 | e^{i\beta_1\sigma_{ij}^x} | \mathbf{z}_m \rangle \langle \mathbf{z}_m | e^{-i\beta_1\sigma_{ij}^x} | \mathbf{z}_{-1} \rangle i z_m^{ij} \sin \left(\gamma_1 (z_1^{ij} - z_{-1}^{ij}) \right) \\
&= \frac{1}{2^9} \sum_{\{\mathbf{z}\}} e^{i\Omega_1(H_c(\mathbf{z}_1 \mathbf{z}_m) - H_c(\mathbf{z}_{-1} \mathbf{z}_m))} \langle \mathbf{z}_1 | e^{i\beta_1\sigma_{ij}^x} | \mathbf{1} \rangle \langle \mathbf{1} | e^{-i\beta_1\sigma_{ij}^x} | \mathbf{z}_{-1} \rangle i z_m^{ij} \sin \left(\gamma_1 (z_1^{ij} - z_{-1}^{ij}) z_m^{ij} \right)
\end{aligned}$$

Where we used the fact that $\langle \mathbf{z}_1 \mathbf{z}_m | e^{i\beta_1\sigma_{ij}^x} | \mathbf{z}_m \rangle = \langle \mathbf{z}_1 | e^{i\beta_1\sigma_{ij}^x} | \mathbf{1} \rangle$. Therefore all the vector elements of $\mathbf{z}_1$ and $\mathbf{z}_{-1}$ must be 1, except on position (ij) . The dependence on z_m^{ij} vanishes in the last part, such that only the constraints are dependent on the $\mathbf{z}_m$ bitstrings.

$$\begin{aligned}
& \mathbb{E}_J \left(\langle \psi(\gamma_1, \beta_1, \Omega_1) | J_{ij} \sigma_{ij}^z | \psi(\gamma_1, \beta_1, \Omega_1) \rangle \right) \\
&= \frac{1}{2^9} \sum_{\mathbf{z}_1, \mathbf{z}_{-1}} \left(\sum_{\mathbf{z}_m} e^{i\Omega_1(H_c(\mathbf{z}_1 \mathbf{z}_m) - H_c(\mathbf{z}_{-1} \mathbf{z}_m))} \right) \langle \mathbf{z}_1 | e^{i\beta_1\sigma_{ij}^x} | \mathbf{1} \rangle \langle \mathbf{1} | e^{-i\beta_1\sigma_{ij}^x} | \mathbf{z}_{-1} \rangle i \sin \left(\gamma_1 (z_1^{ij} - z_{-1}^{ij}) \right)
\end{aligned}$$

Now we can perform the summation over $\mathbf{z}_m$. For this, we can go to the eigenbasis of the constraint Hamiltonian. We will denote $\mathbf{k}_1$ as the constraint bitstring (in this case of length 4) corresponding to $\mathbf{z}_1$, i.e. $k_1^i = 1$ (-1) means that the i th constraint is satisfied (broken).

$$\begin{aligned}
\sum_{\mathbf{z}_m} e^{i\Omega_1(H_c(\mathbf{z}_1 \mathbf{z}_m) - H_c(\mathbf{z}_{-1} \mathbf{z}_m))} &= 2^5 \sum_{\mathbf{k}_m} \exp \left(i\Omega_1 (z_1^{ij} - z_{-1}^{ij}) \sum_{i=0}^3 k_m^i \right) \\
&= 2^5 \sum_{\mathbf{k}_m} \prod_{i=0}^3 \exp \left(i\Omega_1 (z_1^{ij} - z_{-1}^{ij}) k_m^i \right) \\
&= 2^9 \cos^4(\Omega_1 (z_1^{ij} - z_{-1}^{ij}))
\end{aligned}$$

Plugging it back and performing the remaining summations yields,

$$\begin{aligned}
& \mathbb{E}_J \left(\langle \psi(\gamma_1, \beta_1, \Omega_1) | J_{ij} \sigma_{ij}^z | \psi(\gamma_1, \beta_1, \Omega_1) \rangle \right) \\
&= \sum_{\mathbf{z}_1, \mathbf{z}_{-1}} \cos^4(\Omega_1(z_1^{ij} - z_{-1}^{ij})) \langle \mathbf{z}_1 | e^{i\beta_1 \sigma_{ij}^x} | \mathbf{1} \rangle \langle \mathbf{1} | e^{-i\beta_1 \sigma_{ij}^x} | \mathbf{z}_{-1} \rangle i \sin(\gamma_1(z_1^{ij} - z_{-1}^{ij})) \\
&= \cos^4(2\Omega_1) \sin(2\beta_1) \sin(2\gamma_1)
\end{aligned}$$

B.2 Plaquette term

Here the goal is to evaluate the term $\mathbb{E}_J \langle \psi(\gamma_1, \beta_1, \Omega_1) | \sigma_i^z \sigma_j^z \sigma_k^z \sigma_l^z | \psi(\gamma_1, \beta_1, \Omega_1) \rangle$. In order to make the notation slightly lighter, we will use single case labels here [see Fig. 12(B)]. Again, only the gates that are part of the RCC of $\sigma_i^z \sigma_j^z \sigma_k^z \sigma_l^z$ contribute to this expectation value. We start by performing similar steps as before, now the bitstrings $\mathbf{z}_1, \mathbf{z}_{-1}, \mathbf{z}_m$ have length $(2p+2)^2 = 16$.

$$\begin{aligned}
& \mathbb{E}_J \left(\langle \psi(\gamma_1, \beta_1, \Omega_1) | \sigma_i^z \sigma_j^z \sigma_k^z \sigma_l^z | \psi(\gamma_1, \beta_1, \Omega_1) \rangle \right) \\
&= \mathbb{E}_J \left(\langle x | e^{iJ_{ij}\gamma_1 \sum_{n \in \square} \sigma_n^z} e^{i\Omega_1 H_c} e^{i\beta_1 \sum_{n \in \square} \sigma_n^x} \sigma_i^z \sigma_j^z \sigma_k^z \sigma_l^z e^{-i\beta_1 \sum_{n \in \square} \sigma_n^x} e^{-iJ_{ij}\gamma_1 \sum_{n \in \square} \sigma_n^z} | x \rangle \right) \\
&= \frac{1}{2^{16}} \mathbb{E}_J \left(\sum_{\{\mathbf{z}\}} \langle \mathbf{z}_1 | e^{iJ_{ij}\gamma_1 \sum_{n \in \square} \sigma_n^z} | \mathbf{z}_1 \rangle \langle \mathbf{z}_1 | e^{i\Omega_1 H_c} | \mathbf{z}_1 \rangle \langle \mathbf{z}_1 | e^{i\beta_1 \sum_{n \in \square} \sigma_n^x} | \mathbf{z}_m \rangle \langle \mathbf{z}_m | \sigma_i^z \sigma_j^z \sigma_k^z \sigma_l^z | \mathbf{z}_m \rangle \right. \\
&\quad \left. \langle \mathbf{z}_m | e^{-i\beta_1 \sum_{n \in \square} \sigma_n^x} | \mathbf{z}_{-1} \rangle \langle \mathbf{z}_{-1} | e^{i\Omega_1 H_c} | \mathbf{z}_{-1} \rangle \langle \mathbf{z}_{-1} | e^{-iJ_{ij}\gamma_1 \sum_{n \in \square} \sigma_n^z} | \mathbf{z}_{-1} \rangle \right) \\
&= \frac{1}{2^{16}} \sum_{\{\mathbf{z}\}} z_m^i z_m^j z_m^k z_m^l e^{i\Omega_1(H_c(\mathbf{z}_1) - H_c(\mathbf{z}_{-1}))} \langle \mathbf{z}_1 | e^{i\beta_1 \sum_{n \in \square} \sigma_n^x} | \mathbf{z}_m \rangle \langle \mathbf{z}_m | e^{-i\beta_1 \sum_{n \in \square} \sigma_n^x} | \mathbf{z}_{-1} \rangle \\
&\quad \mathbb{E}_J \left(e^{i\gamma_1 \sum_{n \in \square} J_n(z_1^n - z_{-1}^n)} \right) \\
&= \frac{1}{2^{16}} \sum_{\{\mathbf{z}\}} z_m^i z_m^j z_m^k z_m^l e^{i\Omega_1(H_c(\mathbf{z}_1) - H_c(\mathbf{z}_{-1}))} \langle \mathbf{z}_1 | e^{i\beta_1 \sum_{n \in \square} \sigma_n^x} | \mathbf{z}_m \rangle \langle \mathbf{z}_m | e^{-i\beta_1 \sum_{n \in \square} \sigma_n^x} | \mathbf{z}_{-1} \rangle \\
&\quad \prod_{n \in \square} \cos(\gamma_1(z_1^n - z_{-1}^n)) \\
&= \frac{1}{2^{16}} \sum_{\mathbf{z}_1, \mathbf{z}_{-1}} \left(\sum_{\mathbf{z}_m} z_m^i z_m^j z_m^k z_m^l e^{i\Omega_1(H_c(\mathbf{z}_1 \mathbf{z}_m) - H_c(\mathbf{z}_{-1} \mathbf{z}_m))} \right) \langle \mathbf{z}_1 | e^{i\beta_1 \sum_{n \in \square} \sigma_n^x} | \mathbf{1} \rangle \langle \mathbf{1} | e^{-i\beta_1 \sum_{n \in \square} \sigma_n^x} | \mathbf{z}_{-1} \rangle \\
&\quad \prod_{n \in \square} \cos(\gamma_1(z_1^n - z_{-1}^n))
\end{aligned}$$

Now we can again perform the summation over $\mathbf{z}_m$ first, we have that $z_m^i z_m^j z_m^k z_m^l = k_m^5$

$$\begin{aligned}
\sum_{\mathbf{z}_m} z_m^i z_m^j z_m^k z_m^l e^{i\Omega_1(H_c(\mathbf{z}_1 \mathbf{z}_m) - H_c(\mathbf{z}_{-1} \mathbf{z}_m))} &= 2^7 \sum_{\mathbf{k}_m} k_m^5 e^{i\Omega_1 \sum_{n=0}^8 (k_1^n - k_{-1}^n) k_m^n} \\
&= 2^7 \sum_{\mathbf{k}_m} \prod_{n=0}^8 k_m^5 e^{i\Omega_1 (k_1^n - k_{-1}^n) k_m^n} \\
&= 2^{16} i \sin(\Omega_1(k_1^5 - k_{-1}^5)) \prod_{n \neq 5} \cos(\Omega_1(k_1^n - k_{-1}^n))
\end{aligned}$$

Therefore we have that

$$\begin{aligned} & \mathbb{E}_J (\langle \psi(\gamma_1, \beta_1, \Omega_1) | \sigma_i^z \sigma_l^z \sigma_k^z \sigma_l^z | \psi(\gamma_1, \beta_1, \Omega_1) \rangle) \\ &= \sum_{\mathbf{z}_1, \mathbf{z}_{-1}} i \sin(\Omega_1(k_1^5 - k_{-1}^5)) \prod_{n \neq 5} \cos(\Omega_1(k_1^n - k_{-1}^n)) \langle \mathbf{z}_1 | e^{i\beta_1 \sum_{n \in \square} \sigma_n^x} | \mathbf{1} \rangle \langle \mathbf{1} | e^{-i\beta_1 \sum_{n \in \square} \sigma_n^x} | \mathbf{z}_{-1} \rangle \\ & \quad \prod_{n \in \square} \cos(\gamma_1(z_1^n - z_{-1}^n)) \end{aligned}$$

We see that $\mathbf{z}_1$ and $\mathbf{z}_{-1}$ run each over $2^4 = 16$ spin configurations (256 configurations in total). However, the $\sin(\Omega_1(k_1^5 - k_{-1}^5))$ implies that an even (odd) number of flips in $\mathbf{z}_1$ must combine with an odd (even) number of flips in $\mathbf{z}_{-1}$, to be non vanishing. We list the non-vanishing contribution in the table below.

# of flips in $\mathbf{z}_1$	# of flips in $\mathbf{z}_{-1}$	Multiplicity	Term
0	1	4	$\cos^7 \beta_1 \sin \beta_1 \cos(2\gamma_1) \sin(2\Omega_1) \cos^3(2\Omega_1)$
0	3	4	$-\cos^5 \beta_1 \sin^3 \beta_1 \cos^3(2\gamma_1) \sin(2\Omega_1) \cos^5(2\Omega_1)$
1	2	12	$-\cos^5 \beta_1 \sin^3 \beta_1 \cos(2\gamma_1) \sin(2\Omega_1) \cos^3(2\Omega_1)$
1	2	12	$-\cos^5 \beta_1 \sin^3 \beta_1 \cos^3(2\gamma_1) \sin(2\Omega_1) \cos^5(2\Omega_1)$
1	4	4	$\cos^3 \beta_1 \sin^5 \beta_1 \cos^3(2\gamma_1) \sin(2\Omega_1) \cos^5(2\Omega_1)$
2	3	12	$\cos^3 \beta_1 \sin^5 \beta_1 \cos^3(2\gamma_1) \sin(2\Omega_1) \cos^5(2\Omega_1)$
2	3	12	$\cos^3 \beta_1 \sin^5 \beta_1 \cos(2\gamma_1) \sin(2\Omega_1) \cos^3(2\Omega_1)$
3	4	4	$-\cos \beta_1 \sin^7 \beta_1 \cos(2\gamma_1) \sin(2\Omega_1) \cos^3(2\Omega_1)$

Table 2: Non-vanishing combinations. The terms are symmetric upon interchanging $\mathbf{z}_1 \leftrightarrow \mathbf{z}_{-1}$.

First, combining the terms with $\cos(2\gamma_1) \cos^3(2\Omega_1)$ gives

$$\begin{aligned} & 2 \cdot 4 \cos(2\gamma_1) \sin(2\Omega_1) \cos^3(2\Omega_1) (\cos^7 \beta_1 \sin \beta_1 - 3 \cos^5 \beta_1 \sin^3 \beta_1 + 3 \cos^3 \beta_1 \sin^5 \beta_1 - \cos \beta_1 \sin^7 \beta_1) \\ &= 4 \cos(2\gamma_1) \sin(2\Omega_1) \cos^3(2\Omega_1) \sin(2\beta_1) (\cos^6 \beta_1 - 3 \cos^4 \beta_1 \sin^2 \beta_1 + 3 \cos^2 \beta_1 \sin^4 \beta_1 - \sin^6 \beta_1) \\ &= 4 \cos(2\gamma_1) \sin(2\Omega_1) \cos^3(2\Omega_1) \sin(2\beta_1) \cos^3(2\beta_1). \end{aligned}$$

Secondly, combining the terms with $\cos^3(2\gamma_1) \cos^5(2\Omega_1)$ gives

$$\begin{aligned} & 2 \cdot 16 \cos^3(2\gamma_1) \sin(2\Omega_1) \cos^5(2\Omega_1) (\cos^3 \beta_1 \sin^5 \beta_1 - \cos^5 \beta_1 \sin^3 \beta_1) \\ &= -4 \cos^3(2\gamma_1) \sin(2\Omega_1) \cos^5(2\Omega_1) \sin^3(2\beta_1) \cos(2\beta_1) \end{aligned}$$

Combining both finally gives

$$\begin{aligned} & \mathbb{E}_J (\langle \psi(\gamma_1, \beta_1, \Omega_1) | \sigma_i^z \sigma_l^z \sigma_k^z \sigma_l^z | \psi(\gamma_1, \beta_1, \Omega_1) \rangle) \\ &= 4 \cos(2\gamma_1) \sin(2\Omega_1) \cos^3(2\Omega_1) \sin(2\beta_1) \cos(2\beta_1) (\cos^2(2\beta_1) - \cos^2(2\gamma_1) \cos^2(2\Omega_1) \sin^2(2\beta_1)) \end{aligned}$$

B.3 Result $p = 1$

Combining the two previous sections yields

$$\begin{aligned} & \mathbb{E}_J (\langle \psi(\gamma_1, \beta_1, \Omega_1) | H_{PE} | \psi(\gamma_1, \beta_1, \Omega_1) \rangle) \\ & \approx K \cos^4(2\Omega_1) \sin(2\beta_1) \sin(2\gamma_1) \\ & \quad - 4 \cdot P \cdot C \cos(2\gamma_1) \sin(2\Omega_1) \cos^3(2\Omega_1) \sin(2\beta_1) \cos(2\beta_1) (\cos^2(2\beta_1) - \cos^2(2\gamma_1) \cos^2(2\Omega_1) \sin^2(2\beta_1)). \end{aligned}$$

Where we have assumed that all RCC are square. One could easily repeat the calculations for different RCC topologies, but even for small $N \gtrsim 6$, this approximation turns out to perform well.




Figure 13: (A) Sketch of the LHZ triangle for $N = 10$ with the $t = 7$ qubit line indicated. The boundary of the $p = 2$ RCC of $\langle \psi(\gamma, \beta, \Omega) | \sigma_{27}^z \sigma_{78}^z | \psi(\gamma, \beta, \Omega) \rangle$ is indicated in blue. These RCC consist of two isolated parts. Therefore the expectation value factorizes. The expectation value $\langle \psi(\gamma, \beta, \Omega) | \sigma_{27}^z | \psi(\gamma, \beta, \Omega) \rangle$ depends on J_{28} . However Eq. (16) is zero because $z_m^{27} z_m^{28}$ cannot be completed to a product of constraints (inside the RCC). (B) When $p = 3$, $\langle \psi(\gamma, \beta, \Omega) | \sigma_{27}^z \sigma_{78}^z | \psi(\gamma, \beta, \Omega) \rangle$ does not factorize and the product $z_m^{27} z_m^{28} z_m^{78}$ can be completed as indicated by the path of green dots.

In the large N limit ($N = 800$), and for $C = 3$, we find the following optimal angles

$$\begin{aligned}\beta_1 &= 0.219 \\ \gamma_1 &= -0.167 \\ \Omega_1 &= 0.263.\end{aligned}$$

C Vanishing contribution when the 2-body RCCs of H_{ST} do not overlap

In this appendix, we show why

$$\mathbb{E}_J (J_{ij} \langle \psi(\gamma, \beta, \Omega) | \sigma_{it}^z | \psi(\gamma, \beta, \Omega) \rangle \langle \psi(\gamma, \beta, \Omega) | \sigma_{tj}^z | \psi(\gamma, \beta, \Omega) \rangle) = 0$$

where we assumed that the RCCs of σ_{it}^z and σ_{tj}^z do not overlap, i.e. that

$$J_{ij} \langle \psi(\gamma, \beta, \Omega) | \sigma_{it}^z \sigma_{tj}^z | \psi(\gamma, \beta, \Omega) \rangle = J_{ij} \langle \psi(\gamma, \beta, \Omega) | \sigma_{it}^z | \psi(\gamma, \beta, \Omega) \rangle \langle \psi(\gamma, \beta, \Omega) | \sigma_{tj}^z | \psi(\gamma, \beta, \Omega) \rangle.$$

Additionally, we also assume the it and tj are chosen such that either $\langle \psi(\gamma, \beta, \Omega) | \sigma_{it}^z | \psi(\gamma, \beta, \Omega) \rangle$ or $\langle \psi(\gamma, \beta, \Omega) | \sigma_{tj}^z | \psi(\gamma, \beta, \Omega) \rangle$ is dependent on J_{ij} . Otherwise we would anyway have a vanishing contribution, simply because $\mathbb{E}_J(J_{ij}) = 0$. An example of such a case is sketched in Fig. 13(A). Here $t = 7$, $i = 2$ and $j = 8$. When $p = 2$ the RCCs do not overlap, but $\langle \psi(\gamma, \beta, \Omega) | \sigma_{27}^z | \psi(\gamma, \beta, \Omega) \rangle$ depends on J_{28} . Therefore

$$\mathbb{E}_J (J_{28} \langle \psi(\gamma, \beta, \Omega) | \sigma_{27}^z | \psi(\gamma, \beta, \Omega) \rangle) \quad (16)$$

does not factorize. We can however realize that this still vanishes. Making a similar basis expansion as in the previous section, leads to

$$\begin{aligned}\mathbb{E}_J (J_{28} \langle \psi(\gamma, \beta, \Omega) | \sigma_{27}^z | \psi(\gamma, \beta, \Omega) \rangle) &\sim \sum_{\{z\}} z_m^{27} e^{i\Omega_1(H_c(z_1 z_2 z_m) - H_c(z_{-1} z_{-2} z_m))} e^{i\Omega_2(H_c(z_1 z_m) - H_c(z_{-1} z_m))} \sin(\gamma_1(z_1^{28} z_2^{28} z_m^{28} - z_{-1}^{28} z_{-2}^{28} z_m^{28})) \\ &\times G(\gamma_1, \gamma_2, \beta_1, \beta_2, z_1, z_{-1}, z_2, z_{-2}) \\ &= \sum_{z_1, z_{-1}, z_2, z_{-2}} \left[\sum_{z_m} z_m^{27} z_m^{28} e^{i\Omega_1(H_c(z_1 z_2 z_m) - H_c(z_{-1} z_{-2} z_m))} e^{i\Omega_2(H_c(z_1 z_m) - H_c(z_{-1} z_m))} \right] \\ &\times G(\gamma_1, \gamma_2, \beta_1, \beta_2, z_1, z_{-1}, z_2, z_{-2}) \sin(\gamma_1(z_1^{28} z_2^{28} - z_{-1}^{28} z_{-2}^{28})),\end{aligned}$$

here $G(\dots)$ should just be read as being a function of many variables, but that crucially has no dependence on z_m . Now we can see that the factor in square brackets is always zero. To work out $H_c(z_1 z_2 z_m)$, we would typically go to the eigenbasis of this Hamiltonian. (This has been described




Figure 14: The maximal number of qubits that occur in the RCCs for different QAOA algorithms. Importantly, the contraction complexity does not scale with these numbers for Maxcut QAOA, as the structure of the RCCs here are mostly trees. On the other hand, for parity-encoded QAOA, the structure of the RCCs is square lattice. Therefore, the contraction complexity does scale with quadratic qubit count in this case.

by the $\mathbf{k}$ vectors before.) However, we can never complete the product of $z_m^{27} z_m^{28}$ into a product of k_m 's, e.g. $k_m^{27} = z_m^{27} z_m^{28} z_m^{38} z_m^{38}$ (we use the convention that a plaquette is labelled by its leftmost qubit). Therefore each k_m^{ab} that appears in the exponents comes with $z_m^{27} z_m^{28} = 1$ in half of the cases, and $z_m^{27} z_m^{28} = -1$ in the other half of the cases. This then sums to zero. This is true for any p , as long as these RCCs do not overlap.

In case the RCCs overlap, as is the case for $p = 3$ in our example [see Fig. 13(B)]. Here, we see that the product $z_m^{27} z_m^{28} z_m^{78}$ can be completed, as indicated by the green dots in Fig. 13(B). This then implies that the factor will not be vanishing. Concretely, for our example we would have that

$$\begin{aligned}
z_m^{27} z_m^{28} z_m^{78} &= z_m^{27} z_m^{28} z_m^{37} z_m^{37} z_m^{38} z_m^{38} z_m^{47} z_m^{47} z_m^{48} z_m^{48} z_m^{57} z_m^{57} z_m^{58} z_m^{58} z_m^{67} z_m^{67} z_m^{68} z_m^{68} z_m^{78} \\
&= z_m^{27} z_m^{28} z_m^{38} z_m^{37} z_m^{37} z_m^{38} z_m^{48} z_m^{47} z_m^{47} z_m^{48} z_m^{58} z_m^{57} z_m^{57} z_m^{58} z_m^{68} z_m^{67} z_m^{67} z_m^{68} z_m^{78} \\
&= k_m^{27} k_m^{37} k_m^{47} k_m^{57} k_m^{67}.
\end{aligned}$$

D Tensor-network calculations and reverse causal cones

In this appendix, we discuss further details about the tensor-network calculations performed in Sec. 4 of the main text. All the calculations presented in the main text are exact, thus we did not perform any truncation in the “bond” dimensions of the tensor network. For these calculations we have used the QUIMB package [36], which allowed for an implementation of the parity-encoded QAOA circuit (4) and the computation of expectation values on the RCCs. We provide pseudocode in Algorithm 1. To reproduce our parity-encoded results, one needs to set up the problem graph as the triangle [see e.g. Fig. 1], the Hamiltonian as Eq. (2) or Eq. (5) and the QAOA ansatz as Eq. (3). Then the angle optimization is carried out in order to minimize the energy of the chosen Hamiltonian. As described in the main text, we have used ‘L-BFGS-B’ from SciPy for this optimization where we restrict the parameter ranges to the ones of Appendix A. To evaluate local expectation values (i.e. individual terms in the Hamiltonians), QUIMB only contracts tensors (gates) that are within the reverse causal cone of the operators. So the function *LocalExpectationValue* of our pseudocode is already native in QUIMB.

The advantage of the tensor-network picture is that it allows to calculate local expectation values of local quantum circuits for a larger number of qubits than a naive state-vector calculation, on the condition that the depth p is small. This follows from the fact that each local term that is present in the Hamiltonian as a tensor-network contraction over the tensors that are present in the lightcone (or RCC) of this term. The cost of this procedure in general is dependent on the structure of the RCCs, and on the exact contraction path that is found. The disadvantage of the tensor-network calculation compared to a state-vector simulation is that it is less straightforward to sample. However, in this paper we are only interested in computing expectation values.

Algorithm 1: Parity-encoded QAOA

Input: N, p

```

1 Interactions =  $\emptyset$ 
2 for  $j$  in range( $N - 1$ ) do
3   for  $0 \leq i < j$  do
4     Sample  $J_{ij}$ 
5     Add  $J_{ij}$  to Interactions
6 Initialize  $\beta, \gamma$  and  $\Omega$  as arrays of length  $p$ 
7 Function LocalExpectationValue( $\rho, O$ )
8   RCC  $\leftarrow$  Select all gates in  $\rho$  within the light cone of  $O$ 
9   Find efficient contraction path to compute  $\text{Tr}(\rho O)$  restricted to RCC
10  Perform contraction according to the contraction path found
11  return  $\text{Tr}(\rho O)$ 
12 Function MinimizeEnergy(Interactions,  $\beta, \gamma, \Omega$ )
13   Function CalculateEnergy( $\beta, \gamma, \Omega$ )
14      $\rho \leftarrow$  Build the parity encoded QAOA circuit  $|\psi(\gamma, \beta, \Omega)\rangle \langle \psi(\gamma, \beta, \Omega)|$  according to Eq. (3)
15     energy = 0
16     for  $O$  in Hamiltonian do
17       energy += LocalExpectationValue( $\rho, O$ )
18     return energy
19   energy,  $\gamma^*, \beta^*, \Omega^* \leftarrow$  Minimize the energy of the Hamiltonians (2) or (5)
20   return energy,  $\gamma^*, \beta^*, \Omega^*$ 
21 energy,  $\gamma^*, \beta^*, \Omega^* \leftarrow$  MinimizeEnergy(Interactions,  $\beta, \gamma, \Omega$ )
Output: energy,  $\gamma^*, \beta^*, \Omega^*$ 

```

To make a connection with the literature on simulating QAOA circuits with tensor networks [20], we point out some differences with respect to paradigmatic Maxcut QAOA on D -regular graphs $G(N, D)$. However, once again we should note that the fixed p algorithms are not directly comparable, for this we should first fix a performance ratio and then adjust p in the parity-encoded algorithm such that fixed performance is reached.

It is a well known fact that in the large N limit for Maxcut QAOA, the bulk of the RCC subgraphs are trees, followed by a first order contribution of single cycle subgraphs [3]. Tree structures can be efficiently classically simulated. Indeed, it can be demonstrated that the contraction cost is at most scaling as $\mathcal{O}(\exp(p \cdot \text{tw}(G')))$, where $\text{tw}(G')$ denotes the treewidth of the subgraph $G' \subset G$. For tree subgraphs, by default $\text{tw}(G') = 1$, hence the classical scaling is only exponential in p , which allowed for efficient classical evaluations of the Maxcut QAOA cost function for up to $p = 11$ for $D = 3$, see e.g. Ref. [20]. (A tree RCC corresponding to this depth contains maximally $N' = 8190$ qubits, as given by $2 \cdot \frac{(D-1)^{p+1}-1}{D-1}$. We plot these numbers in Fig. 14 as a function of p for $D = 3, 4$.)

On the other hand, for parity-encoded QAOA, the RCCs all have the same square lattice topology (see Fig. 3 up to boundary effects), due to the locality of the algorithm. This means that the number of qubits will only grow quadratically with p , as maximally $(2p+2)^2$ (see Fig. 14). However this reduced scaling in the number of qubits inside the RCC, does not lead to any simplification in the calculation. On the contrary, as the square lattice structure of the RCCs is now inherently loopy. (The treewidth of a square lattice graph is just the lattice width.) The cost of doing exact calculations now scales exponentially in this qubit count. The size of the largest tensor in the contraction path is then also maximally $2^{(2p+2)^2}$. Exponents up to ~ 30 can be simulated with a standard laptop. This strongly limits the exact calculation capabilities, either to very low p , or to a very small system sizes N that does not span complete RCCs. For instance, the parity encoding of the $N = 6$ graph shown in Fig. 1 does not span a full reverse causal cone, even not at lowest depth $p = 1$ (see Fig. 3).

References

- [1] 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, 625–644 (2021), 3(9):625–644, aug 2020. DOI: [10.1038/s42254-021-00348-9](https://doi.org/10.1038/s42254-021-00348-9).
- [2] Kishor Bharti, Alba Cervera-Liarta, 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 (nisq) algorithms. *Rev. Mod. Phys.* 94, 015004 (2022), 94(1):015004, feb 2021. DOI: [10.1103/revmodphys.94.015004](https://doi.org/10.1103/revmodphys.94.015004).
- [3] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm. November 2014. DOI: [10.48550/ARXIV.1411.4028](https://doi.org/10.48550/ARXIV.1411.4028).
- [4] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Leo Zhou. The quantum approximate optimization algorithm and the sherrington-kirkpatrick model at infinite size. *Quantum*, 6:759, July 2022. DOI: [10.22331/q-2022-07-07-759](https://doi.org/10.22331/q-2022-07-07-759). URL <https://doi.org/10.22331/q-2022-07-07-759>.
- [5] Joao Basso, Edward Farhi, Kunal Marwaha, Benjamin Villalonga, and Leo Zhou. The quantum approximate optimization algorithm at high depth for maxcut on large-girth regular graphs and the sherrington-kirkpatrick model. In *Proceedings of the 17th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC '22)*, 7:1–7:21, (2022), October 2021. DOI: [10.4230/LIPICS.TQC.2022.7](https://doi.org/10.4230/LIPICS.TQC.2022.7).
- [6] Bryan O’Gorman, William J. Huggins, Eleanor G. Rieffel, and K. Birgitta Whaley. Generalized swap networks for near-term quantum computing. May 2019. DOI: [10.48550/ARXIV.1905.05118](https://doi.org/10.48550/ARXIV.1905.05118).
- [7] Phillip C. Lotshaw, Thien Nguyen, Anthony Santana, Alexander McCaskey, Rebekah Herman, James Ostrowski, George Siopsis, and Travis S. Humble. Scaling quantum approximate optimization on near-term hardware. *Scientific Reports*, 12(1), jul 2022. DOI: [10.1038/s41598-022-14767-w](https://doi.org/10.1038/s41598-022-14767-w).
- [8] Wolfgang Lechner, Philipp Hauke, and Peter Zoller. A quantum annealing architecture with all-to-all connectivity from local interactions. *Science Advances*, 1(9):e1500838, 2015. DOI: [10.1126/sciadv.1500838](https://doi.org/10.1126/sciadv.1500838). URL <https://www.science.org/doi/abs/10.1126/sciadv.1500838>.
- [9] Wolfgang Lechner. Quantum approximate optimization with parallelizable gates. February 2018. DOI: [10.48550/ARXIV.1802.01157](https://doi.org/10.48550/ARXIV.1802.01157).
- [10] Anita Weidinger, Glen Bigan Mbeng, and Wolfgang Lechner. Error mitigation for quantum approximate optimization. January 2023. DOI: [10.48550/ARXIV.2301.05042](https://doi.org/10.48550/ARXIV.2301.05042).
- [11] Igor L. Markov and Yaoyun Shi. Simulating quantum computation by contracting tensor networks. *SIAM Journal on Computing*, 38(3):963–981, 2008. DOI: [10.1137/050644756](https://doi.org/10.1137/050644756). URL <https://doi.org/10.1137/050644756>.
- [12] Cupjin Huang, Fang Zhang, Michael Newman, Xiaotong Ni, Dawei Ding, Junjie Cai, Xun Gao, Tenghui Wang, Feng Wu, Gengyan Zhang, Hsiang-Sheng Ku, Zhengxiong Tian, Junyin Wu, Haihong Xu, Huanjun Yu, Bo Yuan, Mario Szegedy, Yaoyun Shi, Hui-Hai Zhao, Chunqing Deng, and Jianxin Chen. Efficient parallelization of tensor network contraction for simulating quantum computation. *Nature Computational Science*, 1(9):578–587, September 2021. DOI: [10.1038/s43588-021-00119-7](https://doi.org/10.1038/s43588-021-00119-7). URL <https://doi.org/10.1038/s43588-021-00119-7>.
- [13] Johnnie Gray and Stefanos Kourtis. Hyper-optimized tensor network contraction. *Quantum* 5, 410 (2021), 5:410, mar 2020. DOI: [10.22331/q-2021-03-15-410](https://doi.org/10.22331/q-2021-03-15-410).
- [14] Yong Liu, Yaojian Chen, Chu Guo, Jiawei Song, Xinmin Shi, Lin Gan, Wenzhao Wu, Wei Wu, Haohuan Fu, Xin Liu, Dexun Chen, Guangwen Yang, and Jiangang Gao. Validating quantum-supremacy experiments with exact and fast tensor network contraction. December 2022. DOI: [10.48550/ARXIV.2212.04749](https://doi.org/10.48550/ARXIV.2212.04749).
- [15] Johnnie Gray and Garnet Kin-Lic Chan. Hyper-optimized compressed contraction of tensor networks with arbitrary geometry. June 2022. DOI: [10.48550/ARXIV.2206.07044](https://doi.org/10.48550/ARXIV.2206.07044).
- [16] Jarrod R. McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. The theory

- of variational hybrid quantum-classical algorithms. *New J. Phys.* 18 (2016) 023023, 18(2): 023023, feb 2015. DOI: [10.1088/1367-2630/18/2/023023](https://doi.org/10.1088/1367-2630/18/2/023023).
- [17] Suguru Endo, Zhenyu Cai, Simon C. Benjamin, and Xiao Yuan. Hybrid quantum-classical algorithms and quantum error mitigation. *J. Phys. Soc. Jpn.*, Vol.90, No.3, Article ID: 032001 (2021), 90(3):032001, mar 2020. DOI: [10.7566/jpsj.90.032001](https://doi.org/10.7566/jpsj.90.032001).
 - [18] M. Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J. Coles. Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Nature Communications* 12, 1791 (2021), 12(1), mar 2020. DOI: [10.1038/s41467-021-21728-w](https://doi.org/10.1038/s41467-021-21728-w).
 - [19] 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. December 2018. DOI: [10.48550/ARXIV.1812.04170](https://doi.org/10.48550/ARXIV.1812.04170).
 - [20] 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. DOI: [10.1088/2058-9565/ab8c2b](https://doi.org/10.1088/2058-9565/ab8c2b). URL <https://doi.org/10.1088/2058-9565/ab8c2b>.
 - [21] V. Akshay, D. Rabinovich, E. Campos, and J. Biamonte. Parameter concentrations in quantum approximate optimization. *Phys. Rev. A*, 104:L010401, Jul 2021. DOI: [10.1103/PhysRevA.104.L010401](https://doi.org/10.1103/PhysRevA.104.L010401). URL <https://link.aps.org/doi/10.1103/PhysRevA.104.L010401>.
 - [22] Alexey Galda, Xiaoyuan Liu, Danylo Lykov, Yuri Alexeev, and Ilya Safro. Transferability of optimal qaoa parameters between random graphs. June 2021. DOI: [10.48550/ARXIV.2106.07531](https://doi.org/10.48550/ARXIV.2106.07531).
 - [23] Stefan H. Sack and Maksym Serbyn. Quantum annealing initialization of the quantum approximate optimization algorithm. *Quantum* 5, 491 (2021), 5:491, jul 2021. DOI: [10.22331/q-2021-07-01-491](https://doi.org/10.22331/q-2021-07-01-491).
 - [24] 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. July 2023. DOI: [10.48550/ARXIV.2307.05420](https://doi.org/10.48550/ARXIV.2307.05420).
 - [25] David Sherrington and Scott Kirkpatrick. Solvable model of a spin-glass. *Phys. Rev. Lett.*, 35:1792–1796, Dec 1975. DOI: [10.1103/PhysRevLett.35.1792](https://doi.org/10.1103/PhysRevLett.35.1792). URL <https://link.aps.org/doi/10.1103/PhysRevLett.35.1792>.
 - [26] Michel Talagrand. *Mean Field Models for Spin Glasses: Volume I: Basic Examples*. Springer Berlin Heidelberg, 2011. ISBN 9783642152023. DOI: [10.1007/978-3-642-15202-3](https://doi.org/10.1007/978-3-642-15202-3). URL <http://dx.doi.org/10.1007/978-3-642-15202-3>.
 - [27] Michel Talagrand. *Mean Field Models for Spin Glasses: Volume II: Advanced Replica-Symmetry and Low Temperature*. Springer Berlin Heidelberg, 2011. ISBN 9783642222535. DOI: [10.1007/978-3-642-22253-5](https://doi.org/10.1007/978-3-642-22253-5). URL <http://dx.doi.org/10.1007/978-3-642-22253-5>.
 - [28] Andrea Montanari. Optimization of the sherrington-kirkpatrick hamiltonian. December 2018. DOI: [10.48550/ARXIV.1812.10897](https://doi.org/10.48550/ARXIV.1812.10897).
 - [29] M. Aizenman, J. L. Lebowitz, and D. Ruelle. Some rigorous results on the Sherrington-Kirkpatrick spin glass model. *Communications in Mathematical Physics*, 112(1):3 – 20, March 1987. ISSN 1432-0916. DOI: [10.1007/bf01217677](https://doi.org/10.1007/bf01217677).
 - [30] Andrea Montanari and Subhabrata Sen. Semidefinite programs on sparse random graphs and their application to community detection. April 2015. DOI: [10.48550/ARXIV.1504.05910](https://doi.org/10.48550/ARXIV.1504.05910).
 - [31] Afonso S. Bandeira, Dmitriy Kunisky, and Alexander S. Wein. Computational hardness of certifying bounds on constrained pca problems. February 2019. DOI: [10.48550/ARXIV.1902.07324](https://doi.org/10.48550/ARXIV.1902.07324).
 - [32] Aditi Misra-Spielfenner, Tim Bode, Peter K. Schuhmacher, Tobias Stollenwerk, Dmitry Bagrets, and Frank K. Wilhelm. Mean-field approximate optimization algorithm. March 2023. DOI: [10.48550/ARXIV.2303.00329](https://doi.org/10.48550/ARXIV.2303.00329).
 - [33] Maxime Dupont and Bhuvanesh Sundar. Quantum relax-and-round algorithm for combinatorial optimization. July 2023. DOI: [10.48550/ARXIV.2307.05821](https://doi.org/10.48550/ARXIV.2307.05821).
 - [34] P V Sriluckshmy, Vicente Pina-Canelles, Mario Ponce, Manuel G Algaba, Fedor Šimkovic IV, and Martin Leib. Optimal, hardware native decomposition of parameterized multi-qubit pauli gates. *Quantum Science and Technology*, 8(4):045029, sep 2023. DOI: [10.1088/2058-9565/acfa20](https://doi.org/10.1088/2058-9565/acfa20). URL <https://dx.doi.org/10.1088/2058-9565/acfa20>.

- [35] Asier Ozaeta, Wim van Dam, and Peter L. McMahon. Expectation values from the single-layer quantum approximate optimization algorithm on ising problems. *Quantum Sci. Technol.* 7 045036 (2022), 7(4):045036, September 2020. ISSN 2058-9565. DOI: [10.1088/2058-9565/ac9013](https://doi.org/10.1088/2058-9565/ac9013).
- [36] Johnnie Gray. quimb: a python library for quantum information and many-body calculations. *Journal of Open Source Software*, 3(29):819, 2018. DOI: [10.21105/joss.00819](https://doi.org/10.21105/joss.00819).
- [37] Edward Farhi, David Gamarnik, and Sam Gutmann. The quantum approximate optimization algorithm needs to see the whole graph: Worst case examples. May 2020. DOI: [10.48550/ARXIV.2005.08747](https://doi.org/10.48550/ARXIV.2005.08747).
- [38] Edward Farhi, David Gamarnik, and Sam Gutmann. The quantum approximate optimization algorithm needs to see the whole graph: A typical case. April 2020. DOI: [10.48550/ARXIV.2004.09002](https://doi.org/10.48550/ARXIV.2004.09002).