Warm-starting quantum optimization

Daniel J. Egger¹, Jakub Mareček², and Stefan Woerner¹

¹IBM Quantum, IBM Research – Zurich, Säumerstrasse 4, 8803 Rüschlikon, Switzerland

²Czech Technical University, Karlovo nám. 13, Prague 2, the Czech Republic

There is an increasing interest in quantum algorithms for problems of integer programming and combinatorial optimization. Classical solvers for such problems employ relaxations, which replace binary variables with continuous ones, for instance in the form of higher-dimensional matrix-valued problems (semidefinite programming). Under the Unique Games Conjecture, these relaxations often provide the best performance ratios available classically in polynomial time. Here, we discuss how to warm-start quantum optimization with an initial state corresponding to the solution of a relaxation of a combinatorial optimization problem and how to analyze properties of the associated quantum algorithms. In particular, this allows the quantum algorithm to inherit the performance guarantees of the classical algorithm. We illustrate this in the context of portfolio optimization, where our results indicate that warm-starting the Quantum Approximate Optimization Algorithm (QAOA) is particularly beneficial at low depth. Likewise, Recursive QAOA for MAXCUT problems shows a systematic increase in the size of the obtained cut for fully connected graphs with random weights, when Goemans-Williamson randomized rounding is utilized in a warm start. It is straightforward to apply the same ideas to other randomized-rounding schemes and optimization problems.

1 Introduction

Gate-based quantum computers are expected to help solve problems in quantum chemistry [1–3], machine learning [4, 5], financial simulation [6–13] and combinatorial optimization [14, 15]. The quantum approximate optimization algorithm (QAOA) [16–18], inspired by a Trotterization of adiabatic quantum computing [19–21], runs on gate-based quantum computers [22, 23]. This algorithm encodes a combinatorial optimization problem in a Hamiltonian $\hat{H}_C$ whose

ground state is the optimum solution. The QAOA first creates an initial state which is the ground state of a mixer Hamiltonian $\hat{H}_M$. A common choice of $\hat{H}_M$ and initial state is $-\sum_{i=0}^{n-1} \hat{X}_i$ and $|+\rangle^{\otimes n}$, respectively. Next, in a QAOA with depth p , a quantum circuit applies $\exp(-i\beta_k \hat{H}_M) \exp(-i\gamma_k \hat{H}_C)$ at each layer $k = 1, \dots, p$ to create a trial state $|\psi(\beta, \gamma)\rangle$. A classical optimizer seeks the optimal values of β and γ to create a trial state which minimizes the energy of the Hamiltonian $\hat{H}_C$. This algorithm has lacked theoretical guarantees on its performance ratio and for certain problem instances of MAXCUT it cannot, with constant depth, outperform the classical Goemans-Williamson randomized rounding approximation [24, 25].

Recent work has improved the original QAOA, for instance, by aggregating only the best sampled candidate solutions [15] and carefully choosing the mixer operator to improve convergence [26–29], empirically. Reinforcement learning [30, 31], multi-start methods [32], and local optimization [33] help navigate the QAOA optimization landscape. Algorithms such as the Hamiltonian Variational Ansatz produce optimization landscapes that are easier to navigate [34]. Furthermore, optimal β and γ values concentrate on all typical instances generated by some reasonable distributions which may allow optimization strategies with fewer calls to the quantum computer [35]. Certain local classical algorithms match the performance of QAOA for Ising-like cost functions with multi-spin interactions [36] which has motivated the development of Recursive-QAOA (RQAOA) [24, 37]. RQAOA iteratively reduces the problem size and outperforms QAOA on certain forms of Ising Hamiltonians [24]. Implementing QAOA on noisy quantum hardware is challenging as the number of gates can be high for current gate fidelities [38, 39]. The circuits become especially deep when large p is required or when the native hardware connectivity does not match the problem structure, thence requiring SWAP gates [40]. Therefore, in the near term, quantum computers will most likely only run low-depth QAOA. Low-depth QAOA results are improved by robust control [41] and by mapping β and γ to parameters of the control pulses [42, 43], a method available to cloud-based quantum computers [44] with pulse-level control [45, 46].

Recently, there has been substantial progress [47] in the study of continuous relaxations of NP-Hard com-

Daniel J. Egger: deg@zurich.ibm.com

Jakub Mareček: jakub.marecek@fel.cvut.cz

Stefan Woerner: wor@zurich.ibm.com

binatorial optimization problems. The best-known continuous relaxations of MAXCUT and many other problems take the form of semidefinite programs [48]. These can be solved efficiently both in theoretical models of computation [49], where a real-number arithmetic operation can be performed in unit time, and in practice¹ [51]. Subsequently, the solution of a continuous relaxation of a combinatorial optimization problem is transformed into a good solution of the discrete-valued problem by randomized rounding [52]. For instance, the celebrated Goemans-Williamson (GW) random hyperplane rounding [53, 54] for MAXCUT finds cuts whose expected value is an α fraction of the global optimum, for $0.87856 < \alpha < 0.87857$, with the expectation over the randomization in the rounding procedure. The Unique Games Conjecture [55–57], introduced in Appendix A, suggests that GW randomized rounding has the best possible polynomial-time performance on MAXCUT.

Our work is motivated by the desire to provide at least as good guarantees for QAOA as there are for classical approximations. For example, we show that for MAXCUT a warm-start can preserve the GW approximation ratio at any depth p [58]. Further, if the Unique Games Conjecture were to be false even stronger guarantees may be available, improving upon those for randomized rounding. In simulations, our variant of QAOA consistently performs as well as the GW algorithm or better.

We discuss how to warm-start quantum optimization in Sec. 2.2. We explore warm-starting QAOA (WS-QAOA) numerically in Sec. 3 by relaxing Quadratic Unconstrained Binary Optimization problems to continuous ones which provide QAOA with a good initial solution. In Sec. 4 we use the GW algorithm [53] to warm-start RQAOA. We discuss our results and conclude in Sec. 5.

2 Warm-start Quantum Optimization

2.1 Preliminaries

Quadratic Unconstrained Binary Optimization (QUBO) has been studied in Combinatorial Optimization since the 1960s [59]. A common formulation is

$$\min_{x \in \{0,1\}^n} x^T \Sigma x + \mu^T x, \quad (\text{QUBO})$$

where x is a vector of n binary decision variables, $\Sigma \in \mathbb{R}^{n \times n}$ a symmetric matrix, and $\mu \in \mathbb{R}^n$ a vector. Since for binary variables $x_i^2 = x_i$, μ can be added to the diagonal of Σ , and in the following, we only add μ when it simplifies the notation in the given

¹Notice that the analysis of [50] shows the situation is less trivial in the Turing machine and one may need to consider the dimension or the number of constraints a constant, as it often is for a particular relaxation.

context. Considering that any mixed-integer linear program can be encoded in a QUBO [60], QUBO is NP-Hard. Indeed, even checking local optimality is NP-Hard [61], and hence only very special cases [62, 63] can be solved in polynomial time.

If Σ is positive semidefinite, the trivial continuous relaxation of QUBO

$$\min_{x \in [0,1]^n} x^T \Sigma x, \quad (\text{QP})$$

is a convex quadratic program and the optimal solution c^* of the continuous relaxation is easily obtainable with classical optimizers [64].

If Σ is not positive semidefinite, one can apply the well-known recipe [65] to obtain another continuous-valued relaxation, known as semidefinite programming (SDP):

$$\begin{aligned} \max_{Y \in \mathbb{S}^n} \text{tr}(\Sigma Y) \\ \text{diag}(Y) = e \\ Y \succeq 0, \end{aligned} \quad (\text{SDP})$$

where $\mathbb{S}^{n \times n}$ denotes the set of $n \times n$ symmetric matrices, e is an n -vector of ones, and $Y \succeq 0$ denotes that Y must be positive semidefinite. Given the optimal solution Y^* to (SDP), there exist several approaches to generating solutions of the corresponding (QUBO), often with approximation guarantees, as discussed later in this section and Appendix B. A classical laptop can solve instances of (SDP) relaxations of QUBO, where Σ has 10^{13} entries [51]. Furthermore, quantum computers offer the prospect of some speed-ups in solving SDPs [66, 67], although recent quantum-inspired algorithms for SDPs may reduce the potential speedup [68].

2.2 Continuous warm-start QAOA

The solutions of either continuous-valued relaxation (QP or SDP) can be used to initialize quantum-classical hybrid algorithms, which is known as warm-starting them [69]. In particular, we focus on warm-starting QAOA.

In QAOA, each decision variable x_i of the discrete optimization problem corresponds to a qubit by the relation $x_i = (1 - z_i)/2$. Each z_i is replaced by a spin operator $\hat{Z}_i$ to transform the cost function to a cost Hamiltonian $\hat{H}_C$ [70, 71]. Note that the final measurement in QAOA can be considered as a randomized rounding. In the simplest variant of WS-QAOA, we replace the initial equal superposition state $|+\rangle^{\otimes n}$ with a state

$$|\phi^*\rangle = \bigotimes_{i=0}^{n-1} \hat{R}_Y(\theta_i) |0\rangle_n, \quad (1)$$

which corresponds to the solution c^* of the relaxed Problem (QP). Here, $\hat{R}_Y(\theta_i)$ is a rotation around the

Figure 1: Initial state $\hat{R}_Y(\theta_i)|0\rangle$ (red point) on the Bloch sphere. The dashed blue line from $|0\rangle$ to $|1\rangle$ is the interval $[0, 1]$. The red line is the solution c_i^* to the relaxed problem. The regularization parameter ε restricts the range of allowed rotation angles. The orange path shows the evolution of the quantum state starting at $|0\rangle$ for $c_i^* = 0$, $\varepsilon = 0.25$ and $\beta = \pi/2$ with the rotations $\hat{R}_Y(-\theta_i)\hat{R}_Z(-2\beta)\hat{R}_Y(2\theta_i)$.

Y-axis of qubit i with angle $\theta_i = 2 \arcsin(\sqrt{c_i^*})$ and $c_i^* \in [0, 1]$ is the i -th coordinate of the optimum of the continuous-valued relaxation (QP). The probability to measure $|1\rangle$ in qubit i is thus c_i^* , see the geometric representation in Fig. 1.

We also replace the mixer Hamiltonian $\hat{H}_M = -\sum_{i=0}^{n-1} \hat{X}_i$ with $\hat{H}_M^{(ws)} = \sum_{i=0}^{n-1} \hat{H}_{M,i}^{(ws)}$ where

$$\hat{H}_{M,i}^{(ws)} = \begin{pmatrix} 2c_i^* - 1 & -2\sqrt{c_i^*(1-c_i^*)} \\ -2\sqrt{c_i^*(1-c_i^*)} & 1 - 2c_i^* \end{pmatrix} \quad (2)$$

and has $\hat{R}_Y(\theta_i)|0\rangle$ as ground state with eigenvalue of -1 . The ground state of $\hat{H}_M^{(ws)}$ is thus $|\phi^*\rangle$ with eigenvalue $-n$ [16]. Therefore, WS-QAOA applies at layer k a mixing gate which is given by the time-evolved mixing Hamiltonian $\exp(-i\beta_k \hat{H}_M^{(ws)})$, see Fig. 2. Since $\hat{H}_{M,i}^{(ws)} = -\sin(\theta_i)\hat{X} - \cos(\theta_i)\hat{Z}$ the time-evolved mixing Hamiltonian is a rotation around the axis $\vec{n} = [-\sin(\theta_i), 0, -\cos(\theta_i)]$ on the Bloch-sphere of qubit i and can be implemented using the single-qubit rotations $\hat{R}_Y(\theta_i)\hat{R}_Z(-2\beta)\hat{R}_Y(-\theta_i)$.

If a coordinate in the optimal solution of a continuous relaxation is $c_i^* = 0$ or $c_i^* = 1$, qubit i would be initialized in state $|0\rangle$ or $|1\rangle$, respectively. In such cases, the qubit will remain in its initial state throughout the QAOA optimization when $\hat{H}_C$ contains only $\hat{Z}_i\hat{Z}_j$ and identity spin-operators. This creates a reachability issue when the optimal continuous and discrete solutions do not overlap, i.e., $d_i^* = 1$ and $c_i^* = 0$ or $d_i^* = 0$ and $c_i^* = 1$, where d^* is the solution to the (QUBO).

Figure 2: Quantum circuit for WS-QAOA. The first $\hat{R}_Y$ rotations prepare the initial state $|\phi^*\rangle$. The mixer operator, i.e. $\hat{R}_Y(\theta_i)\hat{R}_Z(-2\beta_k)\hat{R}_Y(-\theta_i)$, is applied after the time-evolved problem Hamiltonian $\hat{H}_C$.

To mitigate this effect, we introduce a variant of WS-QAOA that utilizes a regularization parameter $\varepsilon \in [0, 0.5]$ and changes the rotation angle creating the initial state according to

$$\begin{aligned} \theta_i &= 2 \arcsin(\sqrt{c_i^*}) & \text{if } c_i^* \in [\varepsilon, 1 - \varepsilon], \\ \theta_i &= 2 \arcsin(\sqrt{\varepsilon}) & \text{if } c_i^* \leq \varepsilon, \\ \theta_i &= 2 \arcsin(\sqrt{1 - \varepsilon}) & \text{if } c_i^* \geq 1 - \varepsilon. \end{aligned}$$

The mixer Hamiltonian is adjusted accordingly. The parameter ε provides a continuous mapping between WS-QAOA and standard QAOA since at $\varepsilon = 0.5$ the initial state is the equal superposition state and the mixer Hamiltonian is the X operator. If all $c_i^* \in (0, 1)$ or $\varepsilon > 0$, WS-QAOA converges to the optimal solution of (QUBO) as the depth p approaches infinity as does standard QAOA [16]. This directly follows from the adiabatic theorem and the fact that we start in the ground state of the mixer which overlaps with all computational basis states including the optimal solution. For large enough p , (WS-)QAOA therefore reproduces the adiabatic evolution transforming the ground state of the mixer into the ground state of $\hat{H}_C$. The speed of the adiabatic evolution is limited by the spectral gap of the intermediate Hamiltonians. If the evolution is too fast transitions to excited states occur which may not result in the optimal solution. The speed of the evolution can be related to p , where a slow evolution, i.e., longer total evolution time, implies a larger p . The idea of (WS-)QAOA is to speed-up this evolution by optimizing the parameters instead of following a fixed annealing schedule.

2.3 Rounded warm-start QAOA

Further variants of WS-QAOA randomly round the optimum of the continuous-valued relaxation before using it as the initial state. This is appealing to quantum hardware with limited qubit numbers as even for convex relaxations in dimensions that scale super-linearly with the number n of binary variables in (QUBO), such as (SDP) with dimension $n(n+1)/2$, the representation of the rounded solution to (QUBO)

Table 1: An overview of design choices in warm-starting a quantum optimization algorithm for (QUBO). Under “What to round?”, columns are ordered left to right to suggest the increasing strength of the relaxations, although this is necessarily fraught in the case of hierarchies of relaxations [72–74], where one column represents a potentially infinite number of relaxations. Similarly, under “How to round?”, we order the options approximately by their performance.

Variant	What to round?	When to round?	How to round?
	<p>(QP) Relaxation</p> <p>SOCP Relaxations [74]</p> <p>(SDP) Relaxation [53, 65, 75]</p> <p>Entropy-penalized SDP [76]</p> <p>Sparse Moment SDP [73]</p> <p>Moment SDP [72]</p>	<p>Classical pre-processing</p> <p>Quantum circuit (QC)</p>	<p>0.5-Approximation [77, 78]</p> <p>Random hyperplane [53, 75]</p> <p>Iterative [79, 80]</p> <p>Iterative [81]</p> <p>Measurement in a QC</p>
WS-QAOA	✓		
WS-RQAOA	✓	✓	✓

requires only $O(n)$ qubits. Two notable examples are the random-hyperplane rounding of SDP relaxations for MAXCUT [53], see Appendix B, and iterative rounding of SDP relaxations for a wider variety of problems, see Appendix D. Both of these examples provide initial states that already have the best approximation guarantee available classically in polynomial time.

We now elaborate on the example of Goemans-Williamson random-hyperplane rounding of (SDP). For a given GW cut we generate an initial state using Y -rotations with a $\varepsilon \in (0, 0.5)$, as discussed in Sec. 2.2. To warm-start QAOA such that we can retain the GW bound on MAXCUT, we wish create a quantum circuit that can both represent solutions of the random-hyperplane rounding as well as deviate from them. We therefore modify the mixer such that its time-evolution is $\hat{R}_Y(-\theta_i)\hat{R}_Z(-2\beta)\hat{R}_Y(\theta_i)$ instead of $\hat{R}_Y(\theta_i)\hat{R}_Z(-2\beta)\hat{R}_Y(-\theta_i)$, i.e., we multiply the off-diagonal elements in (2) by -1 . With this modification, the value of the regularization parameter ε can be set to 0.25 to generate states that differ from the GW rounding as well as retain it by choosing $\beta_1 = \pi/2$ and $\gamma_1 = 0$. At these values the depth-one variational form reduces to

$$\hat{R}_Y(-\theta_i)\hat{R}_Z(-\pi)\hat{R}_Y(2\theta_i)|0\rangle, \quad (3)$$

for each qubit, and creates the states $-i|1\rangle$ and $-i|0\rangle$ when $c_i^* = \varepsilon$ and $1 - \varepsilon$, respectively. Thus, the variational form can recover the solution given by the GW rounding, considering that z and $1 - z$ represent the same cut, see the orange path in Fig. 1 as example. Therefore, WS-QAOA is at least as good as GW rounding. This adjustment also comes with a drawback. Since the prepared initial state is no longer an eigenstate of the mixer (otherwise we would not be able to deviate from it) we cannot use the same arguments as in [16] to derive the convergence of the al-

gorithm to the global optimum with increasing depth p . We will analyze this numerically in Sec. 4.

Notice that measuring an initial state provided by a randomized rounding of the semidefinite programming relaxation (SDP) yields the best approximation guarantees available classically in polynomial time under the Unique Games Conjecture [55–57]. Therefore, any quantum circuit that preserves or improves the performance ratio would preserve or improve the overall performance guarantees.

Rounding in the classical pre-processing readily leads to the warm-started recursive QAOA (WS-RQAOA), illustrated in Fig. 3 and demonstrated in Sec. 4. For MAXCUT of a graph G_n , we leverage a GW pre-solver $GW(G_n, N, M)$ to generate N good cuts of which we retain the $M < N$ best unique cuts. These M cuts therefore initialize M WS-QAOA optimizers with $\varepsilon \in (0, 0.5)$. Each QAOA solver produces an optimized variational state $|\psi^*\rangle_l = \sum_{i=0}^{2^n-1} \alpha_{il} |i\rangle$ for $l = 1, \dots, M$. We then aggregate these M variational states by averaging the probability of sampling each bit-string $|i\rangle$, i.e. $\bar{p}_i = M^{-1} \sum_{l=1}^M |\alpha_{il}|^2$, and use these average probabilities to create the correlation matrix $\mathcal{M}$ needed by RQAOA [24, 37], see Appendix E and F. At each iteration, RQAOA removes one decision variable z_i from the problem by replacing it with $\text{sign}(\mathcal{M}_{ij})z_j$, where $(i, j) = \arg \max_{(i,j)} |\mathcal{M}_{ij}|$. This generates a new MAXCUT problem with a new graph G_{n-1} , see Appendix G, for which we repeat this procedure, illustrated in Fig. 3, until the reduced graph reaches a certain size n_{stop} . The graph $G_{n_{\text{stop}}}$ is solved by diagonalizing the Hamiltonian $\hat{H}_C$ or by applying classical optimizers.

Figure 3: WS-RQAOA for MAXCUT. At each iteration we run several WS-QAOAs that are initialized with different solutions from the GW randomized rounding. The resulting samples are aggregated to compute the combined correlation matrix $\mathcal{M}$ needed by RQAOA to eliminate a decision variable.

2.4 Further variants of warm-starting quantum optimization

In Sec. 2.2 and 2.3, we gave first examples of how to warm-start QAOA using a continuous relaxation and a randomized rounding. The key algorithm-design questions in warm-starting quantum optimization are: what to round, when to round it, and how to round it. For each of these questions, there are multiple options available, as suggested in the previous discussion and summarized in Tab. 1.

First, there are many options for what to round, outside of the (QP) relaxation and the (SDP) relaxation. For example, the (QP) relaxation can be seen as a second-order cone programming (SOCP) relaxation, and could be strengthened iteratively [74], until its objective-function value coincides with the objective-function value of the non-convex (QUBO), albeit at the cost of an exponential growth of the relaxation. Similarly, one could strengthen the (SDP) relaxation either by using an entropy-penalizing term [76] or by using the Moment/SOS hierarchy [72] and its sparse variant [73].

Second, there are two options for when to round: either in the classical pre-processing — within the initial state preparation which leads to the WS-RQAOA discussed in Sec. 2.3 on the example of the (SDP) relaxation — or within the quantum circuit. In its simplest form, the latter can be a quantum measurement, as discussed in Sec. 2.2 on the example of the (QP) relaxation.

Third, there are several options for the rounding procedure. Even the simplest rounding mechanisms often perform well: on MAXCUT, for example, disregarding the relaxation and coordinate-wise

assigning a value uniformly at random achieves a 0.5 approximation ratio [77] and can be derandomized [78, Chapter 6]. The random-hyperplane rounding of GW [53], as explained in Appendix B, improves the performance ratio on MAXCUT to $\alpha = \frac{2}{\pi} \min_{0 \leq \theta \leq \pi} \frac{\theta}{1 - \cos \theta} \approx 0.878$. The same ratio can also be obtained with an iterative procedure that rounds coordinates that are close to being integral to integers [79, 80] and removes them from further processing², as explained in Appendix D. Plausibly, the same ratio could also be achieved with a number of novel and very different iterative procedures, such as [81].

2.5 Discussion of warm-starting quantum optimization

On noisy quantum hardware it seems appealing to use the WS-RQAOA with the strongest available relaxations [72–74] in the classical pre-processing. However, higher-order relaxations within these hierarchies [72–74] require a run-time of the classical SDP solver which is super-polynomial in the number n of integral decision variables in (QUBO) and the order in the hierarchy [72–74]. Therefore, we limit ourselves to the use of WS-RQAOA with the basic (SDP) relaxation, whose value can be approximated classically to any fixed precision in polynomial time.

In contrast, one could extend the use of the continuous-valued solution c^* of the (QP) relaxation to either the solution Y^* of the basic (SDP) relaxation, or its strengthened variants [72, 73], when preparing the initial state. However, this may require more qubits than would be practical in the near-term. For example, a naïve approach to prepare the initial state would utilize $\Theta(n^2)$ and $\Omega(n^2)$ qubits to represent the optimum Y^* of the basic (SDP) relaxation and its strengthened variants, respectively³. At the same time, strong performance guarantees would be readily available for such variants of warm-started quantum optimization. For example, consider representing the matrix-valued solution of a (SDP) relaxation in a $O(n^2)$ -qubit initial state, applying a parametrized quantum circuit that allows for the identity in the unitary representation, at least for some choice of its parameters, and then, measuring the qubit register. This can be seen as a randomized-rounding algorithm, and one can hence analyze the quality of the measured output.

A recently-proposed [79] avenue for the analysis of such randomized-rounding algorithms utilizes the Sticky Brownian Motion [82], a well-known concept

²Notice that when the WS-QAOA leverages both a classical optimization routine with a classical conditional statement and a noisy quantum computer, such a rounding procedure can also be considered.

³While more elaborate representations of the matrix have been proposed [66, 67], it is not yet clear how to implement the related oracles in practice.

in Stochastic Analysis, possibly with a slowdown due to the use of a speed-function [79], as explained in Appendix D. In the case of a (SDP) warm start, one can obtain approximation guarantees for rounded solutions that match the best guarantees available classically in polynomial time, see Appendix D.

A particularly important question is whether any of these variants would strictly improve upon the guarantees of GW [53]. Under the Unique Games Conjecture [83, 84], it is strictly impossible to improve upon the guarantees of GW using either quantum or classical algorithms unless a quantum computer can solve NP-Hard problems in polynomial time, which is not believed to be the case [85], or if $P = NP$. However, a richer picture emerges if this conjecture were to be false and it may be possible to improve approximation ratios using both classical and quantum algorithms.

3 Simulations with Continuous-Valued Warm-start

As a first computational illustration of WS-QAOA, we solve combinatorial-optimization problems framed as a financial-portfolio optimization with a budget constraint [15]. An optimal portfolio minimizes risk and maximizes return by exploiting imperfect correlations in a covariance matrix Σ between n assets with expected returns μ [86]. Selecting B assets out of n with equal weights thus requires solving

$$\min_{x \in \{0,1\}^n} qx^T \Sigma x - \mu^T x \quad (4)$$

$$\text{such that } \mathbf{1}^T x = B, \quad (5)$$

where q controls the risk-return trade-off.

We create random instances of this problem with $n = 6$ assets by simulating the time-series of the asset prices and computing the covariance matrix and returns, see Appendix H. We enforce a budget constraint $B = 3$ with a large quadratic penalty term $\lambda(\mathbf{1}^T x - B)^2$ where we chose $\lambda = 3$ as it is much larger than Σ and μ . Each instance is mapped to an Ising Hamiltonian $\hat{H}_C$. To measure the performance of standard and warm-start QAOA we compute the energy of the optimized trial state $\langle \psi(\beta^*, \gamma^*) | \hat{H}_C | \psi(\beta^*, \gamma^*) \rangle$ labeled as E_{cold}^* and E_{warm}^* , respectively. We normalize E_{cold}^* and E_{warm}^* to the minimum energy E_0 found by diagonalizing $\hat{H}_C$. Since the state-vector simulator in Qiskit [87] evaluates the quantum circuits the only source of randomness is the initial guess for β and γ chosen uniformly from $\pm 2\pi$ given to the COBYLA optimizer we use to find β^* and γ^* . The optimal solution c^* of the continuous relaxation of the problem used to warm-start QAOA is found with IBM[®] ILOG[®] CPLEX[®] 12.10.0 (CPLEX). The probability of sampling the optimal binary solution d^* is more than 5 times higher

Figure 4: (a) Probability to sample the optimal state $|d^*\rangle$ from the optimized trial state $|\psi^*\rangle$ and (b) energy of $|\psi^*\rangle$ for warm-start and standard QAOA at different depths for $n = 6$ assets and $q = 2$. The optimal discrete and continuous solutions are $d^* = (0, 0, 1, 1, 1, 0)$ and $c^* \simeq (0.17, 0, 0.97, 0.73, 1.0, 0.14)$, respectively. QAOA is run ten times with different initial random guesses for (β, γ) chosen uniformly from $\pm 2\pi$. The thick lines show the median of the ten runs while the shaded areas indicate the 25% and 75% quantiles. The gray dashed line shows E_0 .

with WS-QAOA than standard QAOA for the simulated depths $1 \leq p \leq 5$, see Fig. 4(a). Furthermore, the quality of the solution found by WS-QAOA is better than standard QAOA since E_{warm}^* is closer to E_0 than E_{cold}^* , see Fig. 4(b). At depth $p \geq 4$ standard QAOA has enough free parameters to satisfy the budget constraint, as shown by the low energy in Fig. 4(b), but still fails to produce a trial state which contains the optimal solution with high probability.

We investigate the role of the warm-start mixer operator $\hat{H}_M^{(ws)}$ by replacing it with the standard mixer $-\sum_{i=0}^{n-1} \hat{X}_i$ while using the initial state given by the continuous solution c^* . Under these conditions the energy of the optimized state does not converge to the minimum energy, see blue triangles in Fig. 4(b). The probability of sampling the optimal discrete solution is between warm-start and standard QAOA but depends heavily on the initial point given to COBYLA, see Fig. 4(a). These results further justify the use of the modified mixer in WS-QAOA.

To further illustrate the advantage of a warm-start at low depth we solve 250 random portfolio instances with warm-start and standard QAOA, both at depth $p = 1$. Here, the standard QAOA produces variational states that poorly approximate the ground state, see the histogram of E_{cold}^* in Fig. 5(a). However, WS-QAOA produces optimized variational states that are much closer to the minimum energy of each problem Hamiltonian. Furthermore, we find that WS-QAOA tends to produce better solutions when the overlap $d^{*T} c^* / B$ between the optimal solutions

Figure 5: Improvement of depth-one WS-QAOA over standard QAOA for 250 random portfolio instances with $q = 2$. (a) Histogram of the energy of the optimized trial states $|\psi(\beta^*, \gamma^*)\rangle$ with (orange) and without (blue) warm-start normalized to E_0 . We found $|\psi(\beta^*, \gamma^*)\rangle$ with COBYLA seeded with random initial guesses for β and γ chosen uniformly from $\pm 2\pi$. The minimum energy E_0 is found by direct diagonalization. (b) Energy difference of WS-QAOA with the optimal solution, i.e. $\Delta_{\text{warm}} = E_{\text{warm}}^* - E_0$, normalized to the energy difference obtained with standard QAOA, i.e. $\Delta_{\text{cold}} = E_{\text{cold}}^* - E_0$, as a function of the overlap between the optimal solution of the problem with binary weights and continuous weights. $\Delta_{\text{warm}}/\Delta_{\text{cold}} < 1$ implies that WS-QAOA improved the energy of the trial state and $\Delta_{\text{warm}}/\Delta_{\text{cold}} = 0$ implies that WS-QAOA found the optimal portfolio.

Figure 6: Energy of the quantum state normalized to $[\min(\hat{H}_C), \max(\hat{H}_C)]$ after Trotterized annealing for the warm-start mixer (a) and the equal superposition mixer (b). The triangles in (a) and (b) indicate the points for which the time-resolved energy in (c) was obtained.

to the discrete and relaxed problems is closer to 1, see Fig. 5(b).

We now illustrate the advantage of a warm-start in the context of quantum annealing. We simulate annealing by Trotterizing the time-evolution start-

ing from the ground state of the mixer Hamiltonian. At time step $k \in \{0, \dots, N\}$ we apply the operator $\exp(-i\beta_k \hat{H}_M) \exp(-i\gamma_k \hat{H}_C)$ where the annealing schedules are $\beta_k = 2\delta t(1 - i/N)$ and $\gamma_k = 2i\delta t/N$ with $\delta t = T/N$. We compute the energy as function of T , which controls the total annealing time, and N using the equal-superposition mixer $\sum_{i=0}^{n-1} \hat{X}_i$ and the warm-start mixer $\hat{H}_M^{(ws)}$. Warm start requires less annealing time T and a smaller number of steps N than the equal superposition mixer to minimize the energy, as can be seen by comparing (a) and (b) in Fig. 6. This is further confirmed by computing the energy at each time step for $T = 40$ and $N = 60$. When the warm-start mixer $\hat{H}_M^{(ws)}$ is used, the initial and final energies are lower than when the equal superposition mixer is used. These data emphasize, from an annealing point of view, that WS-QAOA converges faster than QAOA.

4 Simulations with Rounded Warm-Start

Next, we discuss warm-starting QAOA for MAXCUT. The maximum cut of an edge-weighted graph $G = (V, E)$ with nodes V , edges E , and weights $\omega_{ij}, \{i, j\} \in E$ is a partitioning of the set of nodes V in two such that the sum of the edge weights ω_{ij} where i and j are in different parts is maximized. This is cast as

$$\max \frac{1}{2} \sum_{(i,j) \in E} \omega_{ij} (1 - z_i z_j) \quad (6)$$

such that $z \in \{-1, 1\}^{|V|}$,

where the binary variable z_i indicates which side of the cut node i is on. In the case of positive edge weights ω_{ij} , for any ϵ , the problem cannot be approximated within the ratio of $16/17 - \epsilon$ classically in polynomial time [88], unless $P = NP$. In the case of the real-valued edge-weights ω_{ij} , the hardness factor is $11/13 - \epsilon$ [75]. In both cases, under the Unique Games Conjecture [55–57], the best guarantees obtainable classically in polynomial time are those of the random-hyperplane rounding [53, 54, 75], as we detail in Appendices B and C.

Since we cannot relax MAXCUT to a (QP) with a positive semi-definite matrix, without a projection onto the cone of positive-semidefinite matrices, we now warm-start QAOA with a *binary* solution obtained using the GW rounding. Here, the variational form can only produce states different from the initial GW cut when the regularization parameter $\epsilon > 0$. We study the effect of ϵ by minimizing the energy of depth-one WS-QAOAs applied to ten fully connected graphs with 30 nodes and edge weights uniformly chosen from $\{-10, -9, \dots, 0, \dots, 9, 10\}$. For each graph we generate ten GW cuts and study the

five best cuts with WS-QAOA. To find the optimal β_1 and γ_1 we seed COBYLA with an initial point obtained as follows. We perform a grid search in γ_1 with 25 equally spaced values from -2π to 2π . For each value of γ_1 we evaluate the energy for the two points $\beta_1 \in \{3\pi/4, 3\pi/8\}$ and fit the result to $a \sin(2\beta_1) + b \sin(4\beta_1)$ to find the β_1 that minimizes the energy without having to perform a two-dimensional grid scan. The (β_1, γ_1) with the lowest energy is given as initial point to COBYLA. At $\varepsilon = 0$ the median energy, normalized to the energy of the maximum cut, is 0.907 and corresponds to the energy of the GW cuts used to warm-start QAOA, see Fig. 7. As ε increases, the normalized energy decreases. However, around $\varepsilon = 0.15$ the median energy starts to increase, and for $\varepsilon = 0.25$ rises beyond the energy of the GW cut to 0.929, which suggests that warm-starting quantum optimization may lead to algorithms which can outperform the GW randomized rounding. This increase in ε is not observed when the mixer from Eq. (2) is used, see Appendix I.

Next, we illustrate the WS-RQAOA algorithm outlined in Sec. 2.3 at depth one by searching for the maximum cut of arbitrary graphs with $n = 20$, and $n = 30$ nodes. Two types of graphs are solved, one where each edge appears with a $p_E = 1/2$ probability and has a $1/2$ probability of having a positive or negative unit weight. The second type of graphs are fully connected with uniformly distributed edge weights sampled from $\{-10, -9, \dots, 0, \dots, 9, 10\}$. We expect that finding the maximum cut for the fully connected graphs will be harder than those with $p_E = 1/2$ [89] and that the resulting QAOA circuits will be deeper as they have more edges [90]. For each graph size and type we randomly generate 100 graphs. At each iteration a GW pre-solver generates $N = 10$ cuts of which we select the best $M = 5$ unique cuts to warm-start five QAOA solvers with a depth $p = 1$ and $\varepsilon = 0.25$, chosen based on the results from Fig. 7. We chose a low N to avoid systematically giving QAOA the maximum cut, see Appendix B. This only holds for the small graphs with which we illustrate WS-RQAOA. For larger graphs we would, however, choose a much larger N as GW cuts can be efficiently generated. The standard and warm-start depth-one RQAOA algorithms are efficiently simulated by computing the correlations $\langle \hat{Z}_i \hat{Z}_j \rangle$ at each iteration, see Appendix F. The parameters β_1 and γ_1 are optimized with COBYLA which is initialized with a good initial point obtained from a grid search, as discussed above, to avoid local-minima. When a graph reaches $n_{\text{stop}} = n/2$ nodes we diagonalize the corresponding Hamiltonian to find the maximum cut of this reduced problem. Together with the replacements from RQAOA we obtain an approximation of the maximum cut of the original graph with n nodes. We compare WS-RQAOA with standard RQAOA.

Our simulations indicate that WS-RQAOA outper-

Figure 7: Energy, normalized to the energy of the maximum cut, as a function of ε for ten graphs each solved five times with different GW cuts. The shaded area indicates the 25% to 75% quantiles and the line shows the median. Each small dot is the energy from one WS-QAOA. The dotted line shows the median normalized energy at $\varepsilon = 0.25$.

forms standard RQAOA, see Fig. 8, and that the number of maximum cuts found decreases with graph size. Fully connected graphs with 30 nodes are the hardest to solve among the graphs we consider. Still, for the graphs in Fig. 8, we observe that the optimal β_1^* and γ_1^* are systematically found. This indicates that when warm-start and standard RQAOA fail, it is because the depth-one variational form is not versatile enough to capture the correlations in the maximum cut. At $\varepsilon = 0.25$ we often observed that $\beta_1^* = \pi/2$ and $\gamma_1^* = 0$, see Fig. 9. We therefore benchmark WS-RQAOA against a classical recursive optimization procedure, where the average correlation matrix of the five best GW cuts is used to eliminate decision variables in each iteration, similarly to RQAOA, see black bars in Fig. 8. This classical algorithm performs better than standard RQAOA, but slightly worse than WS-RQAOA.

We now investigate WS-QAOA for $p > 1$ in a non-recursive setting. Since the efficient algorithm outlined in Appendix F is not valid for $p > 1$ we solve a small, fully connected graph with six nodes and edge weights uniformly distributed in $\{-10, -9, \dots, 0, \dots, 10\}$, see Fig. 9(a). The maximum cut of this graph has size 27. By comparing the energy landscape $E(\beta_1, \gamma_1)$ of a depth-one and depth-three WS-QAOA initialized with the cut 001111, of size 23, we observe that the optimal trial state of deeper variational forms is no-longer the initial GW cut, see Fig. 9(b-d). We study the probability of sampling the maximum cut as a function of p by running 30 WS-QAOAs each with a random initial β and γ chosen uniformly from $[0, 2\pi]$ and $\pm 2\pi$, respectively, for depths $p = 1, \dots, 6$. The probability to sample the maximum cut for this graph increases with p while the energy of the optimized trial state decreases, see Fig. 10. This matches our expectations from the-

Figure 8: Histograms of cut sizes, relative to the maximum cut found by CPLEX, for the best out of ten cuts generated by GW on the initial graph (blue), standard RQAOA (orange), WS-RQAOA (red), and the recursive classical solver based on GW (black). (a) and (b) correspond to random graphs with 20 and 30 nodes, respectively, $p_E = 1/2$ and edge weights in $\{-1, 1\}$. (c) and (d) correspond to fully-connected graphs with 20 and 30 nodes, respectively, with edge weights uniformly distributed in $\{-10, -9, \dots, 0, \dots, 10\}$. The number of maximum cuts found is shown in the gray shaded sub-plots and not the main plot. The dashed line shows the hardness factor $11/13$.

Figure 9: (a) The graph used to study the energy-landscape as a function of β_1 and γ_1 for depth-one (b) and depth-three (c) WS-QAOA. The initial cut is 001111 and $\varepsilon = 0.25$. Edge and node colors indicate the edge weight and maximum cut with value 27, respectively. (d) Ten highest probability cuts in the optimized depth-three trial state. The numbers indicate the cut-size. In (c) the values of β_i and γ_i for $i = 2, 3$ are given by COBYLA after minimizing the energy $E(\beta, \gamma)$ and correspond to the best point in Fig. 10. The inset in (d) is a zoom of (c) around the optimal point (orange dot) found by COBYLA.

ory since the circuit at depth $p + 1$ can reproduce all states of the depth p variational form while being more flexible. Since the energy landscape is non-convex and contains many local minima it is challenging to find globally optimal parameters starting from random guesses of β and γ [32, 91–93]. Even at depth-one with random initial guesses for β_1 and γ_1 COBYLA does not always find the optimal $\beta_1 = \pi/2$ and $\gamma_1 = 0$, see Fig. 9 and Fig. 10(b). The complexity

of the energy landscape, even for this six-node graph, may therefore explain why the energy of the optimized trial state decreases slowly with p .

Figure 10: (a) Probability of sampling the maximum cut of the six-node graph shown in Fig. 9(a) from the optimized trial state and (b) its energy as a function of QAOA depth p . The shaded areas indicate the 25% and 75% quantiles of 30 runs represented as small dots. Their medians are the large dots. The green triangle is the energy of the depth-one trial state with $\beta_1 = \pi/2$ and $\gamma_1 = 0$.

5 Discussion and Conclusion

We hope to have contributed towards a framework for the design of quantum optimization algorithms with a warm start, and towards reasoning about their properties. Currently, these algorithms can achieve the same guarantees as the classical relaxations upon which they are based. If the Unique Games Conjecture is true, these guarantees cannot be improved upon by classical or quantum algorithms running in

polynomial time, unless we can solve NP-Hard problems in polynomial time. However, if this conjecture is false then both quantum and classical algorithms may be able to improve the existing performance guarantees.

An implementation of WS-QAOA is available in Qiskit [94], the open-source software development kit for working with quantum computers. Our simulations show that warm-starting quantum heuristics provides an advantage at low depth. This is particularly important for dense optimization problems intended to be solved on noisy quantum hardware that struggles to implement deep quantum circuits. The portfolio optimization simulations indicate that WS-QAOA finds better solutions than standard QAOA. Here, future work could investigate tying budget constraints into the quantum circuit of WS-QAOA [95].

We have also demonstrated how to continuously transform WS-QAOA to conventional QAOA using the regularization parameter ε . We also showed how to improve QAOA for MAXCUT using the GW algorithm to warm-start RQAOA, albeit by introducing an inconsistency between the mixer and the initial state. By using a grid scan at depth one, we mitigated the effect of local optima. We applied WS-RQAOA at depths $p > 1$ on a single graph. The results suggest that the algorithm provides better solutions as p increases. Future work may extend these simulations to more graphs with different sizes. Further work could also exploit other possible warm-starts, e.g., based on polynomially-solvable special cases [62, 63], where one could for example consider low-rank approximations of Σ or the SDP [96] as well as analysis of the convergence properties when using a modified mixer that does not have the initial state as eigenstate. One may also investigate RQAOA in the context of the continuous warm-start discussed in Sec. 3. Furthermore, we may also warm-start quantum optimization algorithms from candidate solutions obtained with classical solvers such as CPLEX or GUROBI with a time-limit termination criterion.

We expect warm-start to be applicable to other problems within Combinatorial Optimization and Integer Programming, for which a good solution can be found through randomized rounding [52], possibly following an encoding into a QUBO [70, 71, 97], a mixed-integer linear optimization problem [60], or a polynomial unconstrained binary optimization problem [90]. Indeed, both the recipe to obtain SDP relaxations [65] and the analytical tools of Appendix D are applicable to linearly constrained problems equally well. For example, the particle-hole representation for VQE can be seen as a form of warm-start [98]. We anticipate that WS-QAOA is also applicable to other binary optimization problems for which an approximate solution can easily be found using relaxed versions of the problem, without the use of randomized rounding, albeit more research needs to be done in this direction.

Acknowledgements

The authors acknowledge useful discussions with Libor Caha, Donny Greenberg, Sergey Bravyi and Giacomo Nannicini. IBM, the IBM logo, and ibm.com are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. The current list of IBM trademarks is available at <https://www.ibm.com/legal/copytrade>. Jakub Mareček's research has been supported by the OP RDE project CZ.02.1.01/0.0/0.0/16_019/0000765 Research Center for Informatics. One month after releasing our work Ref. [96] appeared on the arXiv.

A The Unique Games Conjecture

We now summarize the Unique Games Conjecture (UGC) without presenting any original material. Inapproximability results suggest that finding an approximate solution of a certain problem-dependent approximation ratio is no easier than finding the optimal solution. The PCP theorem [99–101] shows that in a constraint satisfaction problem, the fraction of satisfiable constraints is NP-Hard to approximate within *some* constant factor. In particular, for a constraint satisfaction problem with at most k variables per constraint, it suggests there is a constant $0 < \alpha < 1$, such that it is NP-Hard to decide whether either all constraints are simultaneously satisfiable or whether every assignment satisfies fewer than an α fraction of the constraints. We refer to [102, Chapters 18 and 19] for an excellent overview.

By building on the PCP theorem Khot suggested the Unique Games Conjecture [103] which can be formulated with two-prover one-round games as well as unique label cover problems. Here, we state the UGC with the unique label cover problem as it relates to MAXCUT [55]. In a unique label cover problem there is a bipartite graph $G = (V, W, E)$ with partition V, W and edges E , an alphabet M , and a bijection $\pi_e: M \rightarrow M$ for every edge $e \in E$. We denote this by $\mathcal{L}(V, W, E, M, \{\pi_e\}_E)$. Given the assignments $\mathcal{A}_V: V \rightarrow M$ and $\mathcal{A}_W: W \rightarrow M$, an edge $e = (v, w) \in E$ is satisfied if and only if $\pi_e(\mathcal{A}_W(w)) = \mathcal{A}_V(v)$. Now let $\text{opt}(\mathcal{L})$ be the maximum weight of edges satisfied by any assignment. The Unique Label Cover problem with parameter δ , denoted $\text{ULC}(\delta)$, is the problem of deciding, given an instance $\mathcal{L}(V, W, E, M, \{\pi_e\}_E)$, whether $\text{opt}(\mathcal{L}) \geq 1 - \delta$ or $\text{opt}(\mathcal{L}) \leq \delta$, see examples in Fig. 11. Then:

Conjecture 1 (Unique Games Conjecture [103]). *For arbitrarily small constants $\zeta, \delta > 0$, there exists a constant $k = k(\zeta, \delta)$ such that it is NP-Hard to determine whether a unique label cover instance with the label set size k , i.e. $k = |M|$, has optimum at least $1 - \zeta$ or at most δ .*

Figure 11: Unique label cover. (a) and (c) are examples of unique label cover problems with a graph with four nodes and four edges denoted $\mathcal{L}_a$ and $\mathcal{L}_c$, respectively. The alphabet has two colors: orange and blue. (b) An assignment of node colors for $\mathcal{L}_a$ where all edges are satisfied, i.e. $\text{opt}(\mathcal{L}_a) = 1$. (d) An assignment of node colors for $\mathcal{L}_c$ where one edge is not satisfied. We have $\text{opt}(\mathcal{L}_c) = 3/4$.

Under the UGC a canonical semidefinite programming relaxation of a constraint satisfaction problem provides the best possible approximation ratio [55, 104, 105]. Nevertheless, the UGC itself has been neither proven nor disproven, despite much effort. The closest to proving the UGC, namely Khot *et al.* [106], essentially proved the so-called 2-to-2-Games Conjecture, which is a specialisation of the UGC to finite fields on two elements. Independently, [107, 108] have shown that there is a subexponential time approximation for the original Unique Games problem, which utilizes a hierarchy of LP [108] or SDP [107] relaxations. This seems to point to the full UGC being true, but does not prove it yet.

Research in quantum complexity is still ongoing. The EPR paradox [109] and the Tsirelson problem [110] can be seen as two-prover one-round games in which the provers may share entanglement while the verifier and all communication are classical. In this spirit, Kempe *et al.* [84] showed that the “unique games with entangled provers” is false by rounding an SDP relaxation of the value of a unique game with entangled provers and with more than two possible answers, in polynomial time. The “games quantum PCP conjecture”, where one wishes to distinguish between the cases when the provers of a game with a classical verifier have a strategy using entanglement that succeeds with probability 1, or when no such strategy succeeds with probability larger than $1/2$, has been shown to be true under randomized reductions [111]. By contrast, the “constraint satisfaction” variant of the quantum PCP conjecture [112], which considers constant-factor approximations to the minimal energy of a local Hamiltonian normalized to 1, remains open.

B Goemans-Williamson Algorithm

The GW algorithm [53] first solves the continuous relaxation of MAXCUT

$$\max \frac{1}{2} \sum_{i < j} \omega_{ij} (1 - v_i^T v_j) \quad (7)$$

with positive edge weights ω_{ij} , where the decision variables v_i are n -dimensional vectors with unit Euclidean norm instead of binary variables $z_i \in \{-1, 1\}$. We denote this $v_i \in S^n$ with S in plain font, in contrast to $\mathbb{S}^n$ for symmetric matrices. The relaxation (7) is efficiently solvable as a semidefinite programming problem [48] to get the optimal vectors v_i^* .

Next, the GW algorithm generates a cut by selecting a vector r uniformly at random on the unit sphere and assigning $z_i = \text{sign}(r^T v_i^*)$ for each node, where the sign function returns 1 for non-negative inputs and -1 elsewhere. That is, the rounding depends on which side of the hyperplane (defined by r) passing through the origin the node lies.

Informally speaking, cuts generated in this way are guaranteed to be on average 87.9% of the size of the maximum cut [53], when averaging over the choice of the random hyperplane in the case of the positive edge weights. Formally,

Proposition 2 (Based on Theorem 3.1 in [53]). *The expected value, with respect to the random hyper-plane defined by the vector r , of the cut size W generated by rounding of the MAXCUT SDP relaxation (7) is:*

$$\begin{aligned} \mathbb{E}[W] &= \sum_{1 \leq i < j \leq n} \omega_{ij} \text{Prob}[\text{sign}(r^T v_i) \neq \text{sign}(r^T v_j)] \\ &= \frac{1}{\pi} \sum_{1 \leq i < j \leq n} \omega_{ij} \arccos(v_i^T v_j) \\ &\geq \alpha W^*, \end{aligned} \quad (8)$$

where W^* denotes the value of the maximum cut and the hardness factor is

$$\alpha = \frac{2}{\pi} \min_{0 \leq \theta \leq \pi} \frac{\theta}{1 - \cos \theta} \approx 0.878. \quad (9)$$

Further, conditional on the Unique Games Conjecture [55–57], this is the best possible guarantee that can be obtained by any classical algorithm in polynomial time.

C Extensions towards QUBO

MAXCUT is a special case of (QUBO)⁴. The GW performance ratio is valid only for MAXCUT, as the

⁴Consider the matrix Σ , where for each edge weight ω_{ij} there is an entry $\Sigma_{ij} = \Sigma_{ji} = -\omega_{ij}$, and for all other values there are zeros.

Figure 12: (a) and (b) show the size of GW cuts relative to the maximum cut for the graphs in Fig. 8(b) and (d), respectively. For each graph, N GW cuts are generated and the averages, minima, and maxima of the per-graph objective-function values are computed. These averages, minima, and maximums are then averaged over the 100 graphs. The dotted-dashed green line shows the fraction of graphs for which the maximum cut was found. It is harder to find the maximum cut on fully connected uniform random graphs than random graphs with $p_E = 1/2$.

special case of (QUBO) [75], and likewise the constants in the inapproximability results.

One can, however, encode most problems in combinatorial optimization into a so-called constraint satisfaction problem [75, 104, 113], for which there is a well-known SDP relaxation and a subsequent rounding procedure [75, 104, 113]. Likewise, one can derive optimal inapproximability results conditional on the Unique Games Conjecture [55–57]. See, for example, Figure 2 of [56].

For example, for MAXCUT with real-valued edge weights [75], which actually generalises the QUBO we have presented, as it does not assume Σ is symmetric, we have:

Proposition 3 (Based on Lemma 6 in [75]). *Let w stand for the total weight of edges in a MAXCUT instance, where it is NP-Hard to decide whether the optimal cut is larger or equal than k or less than αk , where α is the hardness factor (9) for MAXCUT. Then for every $\epsilon > 0$ it is NP-Hard to distinguish instances of QUBO with optimum greater or equal to $2k - w$ from instances of QUBO whose optimum is at most $2\alpha k - w$. The ratio of these two bounds on the optimum is*

$$\frac{2\alpha k - w}{2k - w} = \alpha + w \frac{\alpha - 1}{2k - w}. \quad (10)$$

Moreover, the optimum hardness factor is achieved by

the randomized rounding of an SDP relaxation [75].

This can be used to prove the inapproximability results for MAXCUT with real weights [75], both conditional and independent of the Unique Games Conjecture.

We illustrate the performance of GW on the random graphs with 30 nodes used in Sec. 4. For each graph we generate N cuts with GW and normalize them to the maximum cut which is found with CPLEX. We next calculate the minimum, maximum, and average size of these N cuts for each graph. Finally, we average the minimum, maximum, and average of the 100 graphs, see Fig. 12. The average is stable at 85.2% and 83.7% for the random graphs with $p_E = 1/2$ and the fully connected graphs, respectively, see Fig. 12(a) and (b). These averages are slightly below the GW approximation ratio⁵. When $N > 100$ the maximum cut for more than 80 of the 100 graphs in Fig. 12(a) is found which is why we chose $N = 10$ in Sec. 4. When the graphs are fully connected the GW algorithm does not find as many maximum cuts. For instance, 61 maximum cuts are found at $N = 10^5$ for fully connected graphs, see Fig. 12(b).

D A Stochastic-Analysis Viewpoint

Many randomized rounding procedures can be seen from the viewpoint of stochastic analysis: one obtains random unit vectors $u_1, \dots, u_n \in S^n$ and produces signs $\sigma_1, \dots, \sigma_n \in \{-1, 1\}$. In a natural view of [79], the sign is extracted when an associated stochastic process $\{u_i^T B(t)\}_{t \geq 0}$ first reaches $\{-1, 1\}$, where $\{B(t)\}_{t \geq 0}$ is a Brownian motion in $\mathbb{R}^n$ adapted to the filtration $\{\mathcal{F}_t\}_{t \geq 0}$. The corresponding Sticky Brownian Motion $\forall i \in \{1, \dots, n\}$ is

$$\sigma_i := u_i^T B(T_i), \quad (11)$$

where

$$T_i := \min\{t \geq 0 : |u_i^T B(t)| = 1\}. \quad (12)$$

This can be extended to “Slowed-down” Sticky Brownian Motion [79, 80]. Considering first a speed function $\varphi : [-1, 1] \rightarrow [0, \infty)$ that satisfies

$$\lim_{s \rightarrow 1^-} \varphi(s) = \lim_{s \rightarrow -1^+} \varphi(s) = 0 \quad (13)$$

and

$$\forall s \in (-1, 1), \quad \varphi(s) > 0 \quad (14)$$

and second a stochastic process $\{W_u^\varphi(t)\}_{(u,t) \in \mathbb{R}^n \times [0, \infty)}$ that satisfies:

$$\forall (u, t) \in \mathbb{R}^n \times [0, \infty), \quad W_u^\varphi(0) = 0, \quad (15)$$

⁵This could be seen as an instance of a phase transition [114], beyond which the problem becomes computationally difficult for classical algorithms running in polynomial time. Notice that the MAX-2-SAT of [114] is closely related to MAXCUT [115].

and

$$dW_u^\varphi(t) = \varphi(W_u^\varphi(t))u^T dB(t), \quad (16)$$

one obtains, under mild assumptions [79, 80],

$$\sigma_u^\varphi := \lim_{t \rightarrow \infty} W_u^\varphi(t) \in \{-1, 1\} \text{ a.s.} \quad (17)$$

For the “Slowed-down” Sticky Brownian Motion [79, 80], one can show:

Proposition 4 (Based on [79]). *For $u \in \mathbb{R}^n$ and $t \geq 0$ write $W_u^\xi(t) = W_u(t)$ and $\sigma_u^\xi = \sigma_u \in \{-1, 1\}$, where*

$$\forall s \in [-1, 1], \quad \xi(s) = (1 - s^2)^\alpha, \quad (18)$$

for $\alpha > 0$. Then,

$$\forall u, v \in S^{d-1}, \quad \mathbb{E}[\sigma_u \sigma_v] \approx 0.878. \quad (19)$$

We note that the constant in (19) is not exactly the constant of the Goemans-Williamson [53] work. (See also Appendix B.) However, one can consider a different speed-function ξ to obtain the GW constant. In particular, by seeing the processes as Krivine diffusions, one can obtain:

Proposition 5 (Based on [80]). *For $u \in \mathbb{R}^n$ and $t \geq 0$ write $W_u^\xi(t) = W_u(t)$ and $\sigma_u^\xi = \sigma_u \in \{-1, 1\}$, where*

$$\forall s \in [-1, 1], \quad \xi(s) = \frac{\sqrt{2}}{\sqrt{\pi}} e^{-\frac{1}{2}\Phi^{-1}(\frac{1-s}{2})^2}, \quad (20)$$

where $\Phi: \mathbb{R} \rightarrow \mathbb{R}$ is the standard Gaussian cumulative distribution function, i.e.,

$$\forall x \in \mathbb{R}, \quad \Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-\frac{s^2}{2}} ds. \quad (21)$$

Then,

$$\forall u, v \in S^{d-1}, \quad \mathbb{E}[\sigma_u \sigma_v] = \frac{2}{\pi} \arcsin(u^T v). \quad (22)$$

Compare this to the statement of Proposition 2, noting that $\arccos(t) + \arcsin(t) = \pi/2$ for $-1 \leq t \leq 1$. The proof relies in seeing the process as discrete-time Krivine diffusions [80] and applying Theorem 3 of [80].

E Recursive QAOA

RQAOA [24] is a recursive algorithm to find the ground state of an Ising Hamiltonian $\hat{H}_n = \sum_{i,j} J_{i,j} \hat{Z}_i \hat{Z}_j + \sum_k J_k \hat{Z}_k$ with $J_{i,j}, J_k$ as arbitrary real coefficients and n decision variables. At each step of the recursion a standard QAOA is run to find the state $|\psi^*\rangle = \hat{U}(\beta^*, \gamma^*)|+\rangle^{\otimes n}$ that minimizes the energy $\langle \psi^* | \hat{H}_n | \psi^* \rangle$. For each edge $(i, j) \in E$ the correlator $\mathcal{M}_{i,j} = \langle \psi^* | \hat{Z}_i \hat{Z}_j | \psi^* \rangle$ is computed. Next, the decision variable z_i for which $|\mathcal{M}_{i,j}|$ is largest is replaced

with $\text{sign}(\mathcal{M}_{i,j})z_j$ to generate a new Ising Hamiltonian $\hat{H}_{n-1}$ with $n-1$ decision variables. The recursion stops once the number of variables is below a threshold n_{stop} . The remaining problem is solved with a classical solver. We refer to Appendix C of [24] for the pseudocode and detailed discussion.

F Depth-one RQAOA

Depth-one RQAOA can efficiently be simulated classically [91]. Here, we show the algorithm we used to efficiently simulate depth-one WS-RQAOA. To evaluate the correlator $\langle \hat{Z}_i \hat{Z}_j \rangle = \text{Tr}\{\rho_{i,j} \hat{Z}_i \hat{Z}_j\}$ we only need the density matrix $\rho_{i,j}$ of qubits i and j , see the circuit in Fig. 13. Qubits i and j are first prepared in the state $(\sqrt{1-c_i^*}|0_i\rangle + \sqrt{c_i^*}|1_i\rangle) \otimes (\sqrt{1-c_j^*}|0_j\rangle + \sqrt{c_j^*}|1_j\rangle)$. For each qubit $k \neq i, j$, the cost Hamiltonian applies the gate $\hat{U}_{i,k} \otimes \hat{U}_{j,k}$, where $\hat{U}_{i,k}$ is

$$\hat{U}_1(\gamma\omega_{i,k}) \otimes \hat{U}_1(\gamma\omega_{j,k}) \cdot \text{C-Phase}(-2\gamma\omega_{i,k}). \quad (23)$$

Here, the single-qubit gate $\hat{U}_1(\phi)$ is $\text{diag}(1, e^{i\phi})$. Since the controlled-phase gate $\text{C-Phase}(\phi) = \text{diag}(1, 1, 1, e^{i\phi})$ commutes with the $\hat{U}_1$ gate we move all $\hat{U}_1$ gates to the front of the circuit and apply the phases $\hat{U}_1(\gamma \sum_{k \neq i,j} \omega_{i,k})$ and $\hat{U}_1(\gamma \sum_{k \neq i,j} \omega_{j,k})$ to qubits i and j , respectively. Next, we include the effect of the controlled-phase gate of each qubit $k \neq i, j$, initially in the state $\sqrt{1-c_k^*}|0\rangle + \sqrt{c_k^*}|1\rangle$, on the density matrix $\rho_{i,j}$ by computing $(1-c_k^*)\rho_{i,j} + c_k^* \hat{U}_{ijk} \rho_{i,j} \hat{U}_{ijk}^\dagger$ where $\hat{U}_{ijk} = \text{diag}(1, e^{-2i\gamma\omega_{i,k}}) \otimes \text{diag}(1, e^{-2i\gamma\omega_{j,k}})$. Then we apply the two-qubit operation from the $\omega_{i,j} \hat{Z}_i \hat{Z}_j$ term, i.e. $\hat{U}_{i,j} = \text{diag}(1, e^{i\gamma\omega_{i,j}}, e^{i\gamma\omega_{i,j}}, 1)$, and finally apply the mixer operator before measuring. This is summarized in Alg. 1.

Algorithm 1: Depth-one RQAOA

Initialization: qubit i and j in state $|0\rangle$.

Output: Correlator $\langle \hat{Z}_i \hat{Z}_j \rangle$

Apply $\hat{R}_Y(\theta_i)$ and $\hat{R}_Y(\theta_j)$ to qubit i and j .

Apply $\hat{U}_1(\gamma \sum_{k \neq i,j} \omega_{i,k})$ to qubit i .

Apply $\hat{U}_1(\gamma \sum_{k \neq i,j} \omega_{j,k})$ to qubit j .

for $k \neq i, j$ **do**

$\rho_{i,j} \leftarrow (1 - c_k^*)\rho_{i,j} + c_k^* \hat{U}_{ijk} \rho_{i,j} \hat{U}_{ijk}^\dagger$

end

Apply $\rho_{i,j} \leftarrow \hat{U}_{i,j} \rho_{i,j} \hat{U}_{i,j}^\dagger$

Apply mixer

$$e^{-i\beta(\hat{H}_{M,i}^{(ws)} \otimes \hat{H}_{M,j}^{(ws)})} \rho_{i,j} e^{i\beta(\hat{H}_{M,i}^{(ws)} \otimes \hat{H}_{M,j}^{(ws)})}$$

Measure correlator $\langle \hat{Z}_i \hat{Z}_j \rangle = \text{Tr}\{\rho_{i,j} \hat{Z}_i \hat{Z}_j\}$

Figure 13: Quantum circuit used to compute the correlator $\langle \hat{Z}_i \hat{Z}_j \rangle$. The gates highlighted in red do not need to be taken into account.

G MAXCUT reduction

Here we show that replacing $\hat{Z}_i$ by $\pm\hat{Z}_j$ in a MAXCUT problem results in a new MAXCUT problem with one node less. Without loss of generality we label the nodes from 1 to n such that the spin operator $\hat{Z}_n$ of node n will be replaced by $\hat{Z}_n = \alpha\hat{Z}_k$ with $\alpha = \pm 1$ and $k < n$. The MAXCUT Hamiltonian of the weighted graph is

$$\begin{aligned} \hat{H} &= \frac{1}{4} \sum_{i,j=1}^n \omega_{i,j} (1 - \hat{Z}_i \hat{Z}_j) \\ &= \frac{1}{4} \sum_{i,j=1}^{n-1} \omega_{i,j} (1 - \hat{Z}_i \hat{Z}_j) + \frac{1}{2} \sum_{i=1}^{n-1} \omega_{i,n} (1 - \hat{Z}_i \hat{Z}_n) \end{aligned} \quad (24)$$

We now replace $\hat{Z}_n = \alpha\hat{Z}_k$ in the last term and, since $\omega_{i,n} (1 - \alpha\hat{Z}_i \hat{Z}_k) = \alpha\omega_{i,n} (1 - \hat{Z}_i \hat{Z}_k) + \omega_{i,n}(1 - \alpha)$, we may write

$$\begin{aligned} \sum_{i=1}^{n-1} \omega_{i,n} (1 - \hat{Z}_i \hat{Z}_n) &= \sum_{i=1}^{n-1} \omega_{i,n} (1 - \alpha) \\ &\quad + \sum_{i=1}^{n-1} \alpha\omega_{i,n} (1 - \hat{Z}_i \hat{Z}_k). \end{aligned} \quad (25)$$

We neglect the first sum since it is an energy offset that does not affect the optimization. The Hamiltonian of the reduced problem is therefore

$$\begin{aligned} \hat{H}_{n-1} &= \frac{1}{4} \sum_{i,j=1}^{n-1} \omega_{i,j} (1 - \hat{Z}_i \hat{Z}_j) \\ &\quad + \frac{1}{2} \sum_{i=1}^{n-1} \alpha\omega_{i,n} (1 - \hat{Z}_i \hat{Z}_k) \end{aligned} \quad (26)$$

This Hamiltonian corresponds to a new graph E' in which the weights $\omega'_{i,j}$ with $i, j = 1, \dots, n-1$ have been updated according to

$$\omega'_{ij} = \begin{cases} \omega_{i,j} & \text{if } j \neq k, \\ \omega_{i,j} + \alpha\omega_{i,n} & \text{if } j = k. \end{cases} \quad (27)$$

Figure 14: Energy, normalized to the energy of the maximum cut, as a function of ε for a depth one WS-QAOA. The graphs are the same as those in Fig. 7 but we used the mixer from Eq. (2) which cannot reproduce the GW at $\varepsilon = 0.25$ and depth one.

H Portfolio data

The return vectors and covariance matrices used in Sec. 3 are obtained by simulating the price of each asset following a Geometric Brownian motion for $N = 250$ days. The price of asset i on the k^{th} day is

$$S_{i,k} = S_{i,0} \exp \left[(\mu_i - \sigma_i^2/2)k/N + \sigma_i W_k \right]. \quad (28)$$

Without loss of generality we set the initial price $S_{i,0} = 1$. We randomly chose each mean μ_i and standard deviation σ_i uniformly from $[-5\%, 5\%]$ and $[-20\%, 20\%]$, respectively. The Brownian motion is given by $W_k = \sum_{l=0}^k z_l / \sqrt{N}$ where z_l is drawn from the normal distribution. The return of asset i on the k^{th} day is $r_{i,k} = S_{i,k}/S_{i,k-1} - 1$. The average of $r_{i,k}$ gives the mean return of asset i and the covariance of asset i and j is obtained from $r_{i,k}$ and $r_{j,k}$ where $k = 1, \dots, N$.

I WS-QAOA for MAXCUT with the warm-start mixer

In the main text we modified the WS mixer to retain the GW cut. Here, we explore WS-QAOA for MAX-

Figure 15: Histograms of cut sizes, relative to the maximum cut found by CPLEX, for WS-RQAOA with the modified mixer (red, same data as in Fig. 8), and the WS-QAOA with the mixer from Eq. (2). The dashed line shows the hardness factor 11/13.

CUT with the mixer of Eq. (2). By repeating the analysis done in Fig. 7 we observe that the energy, normalized to the maximum, decreases as a function of ε and does not recover at $\varepsilon = 0.25$, see Fig. 14.

In addition, we repeat the analysis done in Fig. 8 but with the mixer of Eq. (2) and compare it to WS-RQAOA with the modified mixer. Both algorithms have a similar performance, see Fig. 15, which suggests that the amount of correlation between the variables at each iteration, e.g. as in Fig. 14, is still sufficient for WS-RQAOA to produce good results even though the GW cut cannot be sampled with certainty at each iteration. As the size and complexity of the graphs is increased the performance of WS-RQAOA with the mixer from Eq. (2) decreases compared to WS-RQAOA with the modified mixer and indicates the importance of being able to retain the GW cut.

References

- [1] Nikolaï Moll, Panagiotis Barkoutsos, Lev S. Bishop, Jerry M. Chow, Andrew Cross, Daniel J. Egger, Stefan Filipp, Andreas Fuhrer, Jay M. Gambetta, Marc Ganzhorn, and *et al.* Quantum optimization using variational algorithms on near-term quantum devices. *Quantum Sci. Technol.*, 3(3):030503, 2018. DOI: [10.1088/2058-9565/aab822](https://doi.org/10.1088/2058-9565/aab822).
- [2] Abhinav Kandala, Kristan Temme, Antonio D. Corcoles, Antonio Mezzacapo, Jerry M. Chow, and Jay M. Gambetta. Error mitigation extends the computational reach of a noisy quantum processor. *Nature*, 567:491–495, 2018. DOI: [10.1038/s41586-019-1040-7](https://doi.org/10.1038/s41586-019-1040-7).
- [3] Marc Ganzhorn, Daniel J. Egger, Panagiotis Kl. Barkoutsos, Pauline Ollitrault, Gian Salis, Nikolaï Moll, Andreas Fuhrer, Peter Mueller, Stefan Woerner, Ivano Tavernelli, and Stefan Filipp. Gate-efficient simulation of molecular eigenstates on a quantum computer. *Phys. Rev. Applied*, 11:044092, Apr 2019. DOI: [10.1103/PhysRevApplied.11.044092](https://doi.org/10.1103/PhysRevApplied.11.044092).
- [4] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. Quantum machine learning. *Nature*, 549(7671):195–202, 2017. DOI: [10.1038/nature23474](https://doi.org/10.1038/nature23474).
- [5] Vojtech Havlicek, Antonio D. Corcoles, Kristan Temme, Aram W. Harrow, Abhinav Kandala, Jerry M. Chow, and Jay M. Gambetta. Supervised learning with quantum-enhanced feature spaces. *Nature*, 567:209 – 212, 2019. DOI: [10.1038/s41586-019-0980-2](https://doi.org/10.1038/s41586-019-0980-2).
- [6] Daniel J. Egger, Claudio Gambella, Jakub Mareček, Scott McFaddin, Martin Mevissen, Rudy Raymond, Aandrea Simonetto, Sefan Woerner, and Elena Yndurain. Quantum computing for finance: State-of-the-art and future prospects. *IEEE Trans. on Quantum Eng.*, 1: 1–24, 2020. DOI: [10.1109/TQE.2020.3030314](https://doi.org/10.1109/TQE.2020.3030314).
- [7] Stefan Woerner and Daniel J. Egger. Quantum risk analysis. *npj Quantum Inf.*, 5:15, 2019. DOI: [10.1038/s41534-019-0130-6](https://doi.org/10.1038/s41534-019-0130-6).
- [8] Patrick Rebentrost, Brajesh Gupt, and Thomas R. Bromley. Quantum computational finance: Monte Carlo pricing of financial derivatives. *Phys. Rev. A*, 98:022321, Aug 2018. DOI: [10.1103/PhysRevA.98.022321](https://doi.org/10.1103/PhysRevA.98.022321).
- [9] Nikitas Stamatopoulos, Daniel J. Egger, Yue Sun, Christa Zoufal, Raban Iten, Ning Shen, and Stefan Woerner. Option pricing using quantum computers. *Quantum*, 4:291, 2020. DOI: [10.22331/q-2020-07-06-291](https://doi.org/10.22331/q-2020-07-06-291).
- [10] Ana Martin, Bruno Candelas, Ángel Rodríguez-Rozas, José D. Martín-Guerrero, Xi Chen, Lucas Lamata, Román Orús, Enrique Solano, and Mikel Sanz. Toward pricing financial derivatives

- with an ibm quantum computer. *Phys. Rev. Research*, 3:013167, Feb 2021. DOI: [10.1103/PhysRevResearch.3.013167](https://doi.org/10.1103/PhysRevResearch.3.013167).
- [11] Roman Orus, Samuel Mugel, and Enrique Lizaso. Quantum computing for finance: Overview and prospects. *Rev. Phys.*, 4:100028, 2019. DOI: [10.1016/j.revip.2019.100028](https://doi.org/10.1016/j.revip.2019.100028).
 - [12] Daniel J. Egger, Ricardo G. Gutierrez, Jordi Cahue Mestre, and Stefan Woerner. Credit risk analysis using quantum computers. *IEEE Trans. Comput.*, 1:1–1, Nov 2020. DOI: [10.1109/TC.2020.3038063](https://doi.org/10.1109/TC.2020.3038063).
 - [13] Almudena Carrera Vazquez and Stefan Woerner. Efficient state preparation for quantum amplitude estimation. *Phys. Rev. Applied*, 15:034027, Mar 2021. DOI: [10.1103/PhysRevApplied.15.034027](https://doi.org/10.1103/PhysRevApplied.15.034027).
 - [14] Lee Braine, Daniel J. Egger, Jennifer Glick, and Stefan Woerner. Quantum algorithms for mixed binary optimization applied to transaction settlement. *IEEE Trans. on Quantum Eng.*, 2:1–8, 2021. DOI: [10.1109/TQE.2021.3063635](https://doi.org/10.1109/TQE.2021.3063635).
 - [15] Panagiotis Kl. Barkoutsos, Giacomo Nannicini, Anton Robert, Ivano Tavernelli, and Stefan Woerner. Improving variational quantum optimization using cvar. *Quantum*, 4:256, Apr 2020. DOI: [10.22331/q-2020-04-20-256](https://doi.org/10.22331/q-2020-04-20-256).
 - [16] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm, 2014. URL <https://arxiv.org/abs/1411.4028>.
 - [17] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm applied to a bounded occurrence constraint problem, 2014. URL <https://arxiv.org/abs/1412.6062>.
 - [18] Zhi-Cheng Yang, Armin Rahmani, Alireza Shabani, Hartmut Neven, and Claudio Chamon. Optimizing variational quantum algorithms using pontryagin’s minimum principle. *Phys. Rev. X*, 7:021027, May 2017. DOI: [10.1103/PhysRevX.7.021027](https://doi.org/10.1103/PhysRevX.7.021027).
 - [19] Mark W. Johnson, Mohammad HS Amin, Suzanne Gildert, Trevor Lanting, Firas Hamze, Neil Dickson, Richard Harris, Andrew J. Berkley, Jan Johansson, Paul Bunyk, and *et al.* Quantum annealing with manufactured spins. *Nature*, 473(7346):194–198, May 2011. DOI: [10.1038/nature10012](https://doi.org/10.1038/nature10012).
 - [20] Glen Bigan Mbeng, Rosario Fazio, and Giuseppe Santoro. Quantum annealing: a journey through digitalization, control, and hybrid quantum variational schemes, 2019. URL <https://arxiv.org/abs/1906.08948>.
 - [21] Michael Juenger, Elisabeth Lobe, Petra Mutzel, Gerhard Reinelt, Franz Rendl, Giovanni Rinaldi, and Tobias Stollenwerk. Performance of a quantum annealer for Ising ground state computations on chimera graphs, 2019. URL <https://arxiv.org/abs/1904.11965>.
 - [22] Rami Barends, Alireza Shabani, Lucas Lamata, Julian Kelly, Antonio Mezzacapo, Urtzi Las Heras, Ryan Babbush, Austin G. Fowler, Brooks Campbell, Yu Chen, and *et al.* Digitized adiabatic quantum computing with a superconducting circuit. *Nature*, 534(7606):222–226, Jun 2016. DOI: [10.1038/nature17658](https://doi.org/10.1038/nature17658).
 - [23] Madita Willsch, Dennis Willsch, Fengping Jin, Hans De Raedt, and Kristel Michielsen. Benchmarking the quantum approximate optimization algorithm. *Quantum Inf. Process.*, 19(7):197, Jun 2020. DOI: [10.1007/s11128-020-02692-8](https://doi.org/10.1007/s11128-020-02692-8).
 - [24] Sergey Bravyi, Alexander Kliesch, Robert Koenig, and Eugene Tang. Obstacles to variational quantum optimization from symmetry protection. *Phys. Rev. Lett.*, 125:260505, Dec 2020. DOI: [10.1103/PhysRevLett.125.260505](https://doi.org/10.1103/PhysRevLett.125.260505).
 - [25] Gavin E. Crooks. Performance of the quantum approximate optimization algorithm on the maximum cut problem, 2018. URL <https://arxiv.org/abs/1811.08419>.
 - [26] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Hartmut Neven. Quantum algorithms for fixed qubit architectures, 2017. URL <https://arxiv.org/abs/1703.06199>.
 - [27] Stuart Hadfield, Zhihui Wang, Bryan O’Gorman, Eleanor Rieffel, Davide Venturelli, and Rupak Biswas. From the quantum approximate optimization algorithm to a quantum alternating operator ansatz. *Algorithms*, 12(2):34, Feb 2019. DOI: [10.3390/a12020034](https://doi.org/10.3390/a12020034).
 - [28] Linghua Zhu, Ho Lun Tang, George S. Barron, Nicholas J. Mayhall, Edwin Barnes, and Sophia E. Economou. An adaptive quantum approximate optimization algorithm for solving combinatorial problems on a quantum computer, 2020. URL <https://arxiv.org/abs/2005.10258>.
 - [29] Zhihui Wang, Nicholas C. Rubin, Jason M. Dominy, and Eleanor G. Rieffel. XY mixers: Analytical and numerical results for the quantum alternating operator ansatz. *Phys. Rev. A*, 101:012320, Jan 2020. DOI: [10.1103/PhysRevA.101.012320](https://doi.org/10.1103/PhysRevA.101.012320).
 - [30] Sami Khairy, Ruslan Shaydulin, Lukasz Cincio, Yuri Alexeev, and Prasanna Balaprakash. Learning to optimize variational quantum circuits to solve combinatorial problems. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(03):2367–2375, Apr 2020. DOI: [10.1609/aaai.v34i03.5616](https://doi.org/10.1609/aaai.v34i03.5616).
 - [31] Matteo M. Wauters, Emanuele Panizon, Glen B. Mbeng, and Giuseppe E. Santoro. Reinforcement-learning-assisted quantum optimization. *Phys. Rev. Research*, 2:

033446, Sep 2020. DOI: [10.1103/PhysRevResearch.2.033446](https://doi.org/10.1103/PhysRevResearch.2.033446).

- [32] Ruslan Shaydulin, Ilya Safro, and Jeffrey Larson. Multistart methods for quantum approximate optimization. *2019 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–8, Sep 2019. DOI: [10.1109/HPEC.2019.8916288](https://doi.org/10.1109/HPEC.2019.8916288).
- [33] Ruslan Shaydulin and Yuri Alexeev. Evaluating quantum approximate optimization algorithm: A case study. In *2019 Tenth International Green and Sustainable Computing Conference (IGSC)*, pages 1–6, 2019. DOI: [10.1109/IGSC48788.2019.8957201](https://doi.org/10.1109/IGSC48788.2019.8957201).
- [34] Roeland Wiersema, Cunlu Zhou, Yvette de Sereville, Juan Felipe Carrasquilla, Yong Baek Kim, and Henry Yuen. Exploring entanglement and optimization within the hamiltonian variational ansatz. *PRX Quantum*, 1:020319, Dec 2020. DOI: [10.1103/PRXQuantum.1.020319](https://doi.org/10.1103/PRXQuantum.1.020319).
- [35] Fernando G. S. L. Brandao, Michael Broughton, Edward Farhi, Sam Gutmann, and Hartmut Neven. For fixed control parameters the quantum approximate optimization algorithm’s objective function value concentrates for typical instances, 2018. URL <https://arxiv.org/abs/1812.04170>.
- [36] Matthew B. Hastings. Classical and quantum bounded depth approximation algorithms, 2019. URL <https://arxiv.org/abs/1905.07047>.
- [37] Sergey Bravyi, Alexander Kliesch, Robert Koenig, and Eugene Tang. Hybrid quantum-classical algorithms for approximate graph coloring, 2020. URL <https://arxiv.org/abs/2011.13420>.
- [38] Mahabubul Alam, Abdullah Ash-Saki, and Swaroop Ghosh. Analysis of quantum approximate optimization algorithm under realistic noise in superconducting qubits, 2019. URL <https://arxiv.org/abs/1907.09631>.
- [39] Vishwanathan Akshay, Hariphan Philathong, Mauro E. S. Morales, and Jacob D. Biamonte. Reachability deficits in quantum approximate optimization. *Phys. Rev. Lett.*, 124:090504, Mar 2020. DOI: [10.1103/PhysRevLett.124.090504](https://doi.org/10.1103/PhysRevLett.124.090504).
- [40] Matthew P. Harrigan, Kevin J. Sung, Matthew Neeley, Kevin J. Satzinger, Frank Arute, Kunal Arya, Juan Atalaya, Joseph C. Bardin, Rami Barends, Sergio Boixo, and *et al.* Quantum approximate optimization of non-planar graph problems on a planar superconducting processor. *Nat. Phys.*, 17(3):332–336, Mar 2021. DOI: [10.1038/s41567-020-01105-y](https://doi.org/10.1038/s41567-020-01105-y).
- [41] Yulong Dong, Xiang Meng, Lin Lin, Robert Kosut, and K. Birgitta Whaley. Robust control optimization for quantum approximate optimization algorithms. *IFAC-PapersOnLine*, 53(2):242–249, 2020. DOI: [10.1016/j.ifacol.2020.12.130](https://doi.org/10.1016/j.ifacol.2020.12.130). 21th IFAC World Congress.
- [42] Nathan Lacroix, Christoph Hellings, Christian Kraglund Andersen, Agustin Di Paolo, Ants Remm, Stefania Lazar, Sebastian Krinner, Graham J. Norris, Mihai Gabureac, Johannes Heinsoo, and *et al.* Improving the performance of deep quantum optimization algorithms with continuous gate sets. *PRX Quantum*, 1:110304, Oct 2020. DOI: [10.1103/PRXQuantum.1.020304](https://doi.org/10.1103/PRXQuantum.1.020304).
- [43] Nathan Earnest, Caroline Tornow, and Daniel J. Egger. Pulse-efficient circuit transpilation for quantum applications on cross-resonance-based hardware, 2021. URL <https://arxiv.org/abs/2105.01063>.
- [44] Pranav Gokhale, Ali Javadi-Abhari, Nathan Earnest, Yunong Shi, and Frederic T. Chong. Optimized quantum compilation for near-term algorithms with openpulse, 2020. URL <https://www.microarch.org/micro53/papers/738300a186.pdf>.
- [45] David C. McKay, Thomas Alexander, Luciano Bello, Michael J. Biercuk, Lev Bishop, Jiayin Chen, Jerry M. Chow, Antonio D. Córcoles, Daniel J. Egger, Stefan Filipp, and *et al.* Qiskit backend specifications for OpenQASM and OpenPulse experiments, 2018. URL <https://arxiv.org/abs/1809.03452>.
- [46] Thomas Alexander, Naoki Kanazawa, Daniel J. Egger, Lauren Capelluto, Christopher J. Wood, Ali Javadi-Abhari, and David C. McKay. Qiskit pulse: programming quantum computers through the cloud with pulses. *Quantum Sci. Technol.*, 5(4):044006, Aug 2020. DOI: [10.1088/2058-9565/aba404](https://doi.org/10.1088/2058-9565/aba404).
- [47] Anirudha Majumdar, Georgina Hall, and Amir Ali Ahmadi. Recent scalability improvements for semidefinite programming with applications in machine learning, control, and robotics. *Annu. Rev. Control Robot. Auton. Syst.*, 3:331–360, 2020. DOI: [10.1146/annurev-control-091819-074326](https://doi.org/10.1146/annurev-control-091819-074326).
- [48] Miguel F. Anjos and Jean B. Lasserre. *Handbook on semidefinite, conic and polynomial optimization*, volume 166. Springer Science & Business Media, 2011. DOI: [10.1007/978-1-4614-0769-0](https://doi.org/10.1007/978-1-4614-0769-0).
- [49] Lenore Blum, Felipe Cucker, Michael Shub, and Steve Smale. *Complexity and real computation*. Springer Science & Business Media, 2012. DOI: [10.1007/978-1-4612-0701-6](https://doi.org/10.1007/978-1-4612-0701-6).
- [50] Lorant Porkolab and Leonid Khachiyan. On the complexity of semidefinite programs. *J. Glob. Optim.*, 10(4):351–365, 1997. DOI: [10.1023/A:1008203903341](https://doi.org/10.1023/A:1008203903341).

- [51] Alp Yurtsever, Joel A. Tropp, Olivier Fercoq, Madeleine Udell, and Volkan Cevher. Scalable semidefinite programming. *SIAM J. Math. Data Sci.*, 3(1):171–200, 2021. DOI: [10.1137/19M1305045](https://doi.org/10.1137/19M1305045).
- [52] Prabhakar Raghavan and Clark D. Thompson. Randomized rounding: A technique for provably good algorithms and algorithmic proofs. *Combinatorica*, 7(4):365–374, Dec 1987. DOI: [10.1007/BF02579324](https://doi.org/10.1007/BF02579324).
- [53] Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42(6):1115–1145, nov 1995. DOI: [10.1145/227683.227684](https://doi.org/10.1145/227683.227684).
- [54] Howard Karloff. How good is the Goemans–Williamson MAX CUT algorithm? *SIAM J. Comput.*, 29(1):336–350, 1999. DOI: [10.1137/S0097539797321481](https://doi.org/10.1137/S0097539797321481).
- [55] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O’Donnell. Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? *SIAM J. Comput.*, 37(1):319–357, 2007. DOI: [10.1137/S0097539705447372](https://doi.org/10.1137/S0097539705447372).
- [56] Subhas Khot. On the unique games conjecture (invited survey). In *2012 IEEE 27th Conference on Computational Complexity*, pages 99–121, Los Alamitos, CA, USA, jun 2010. IEEE Computer Society. DOI: [10.1109/CCC.2010.19](https://doi.org/10.1109/CCC.2010.19).
- [57] Subhash A. Khot and Nisheeth K. Vishnoi. The unique games conjecture, integrality gap for cut problems and embeddability of negative-type metrics into ℓ_1 . *J. ACM*, 62(1):1–39, 2015. DOI: [10.1145/2629614](https://doi.org/10.1145/2629614).
- [58] Kunal Marwaha. Local classical MAX-CUT algorithm outperforms $p = 2$ QAOA on high-girth regular graphs. *Quantum*, 5:437, April 2021. DOI: [10.22331/q-2021-04-20-437](https://doi.org/10.22331/q-2021-04-20-437).
- [59] Peter L. Hammer and Sergiu Rudeanu. *Boolean methods in operations research and related areas*. Springer Science & Business Media, 1968. DOI: [10.1007/978-3-642-85823-9](https://doi.org/10.1007/978-3-642-85823-9).
- [60] Jean B. Lasserre. A MAX-CUT formulation of 0/1 programs. *Oper. Res. Lett.*, 44(2):158 – 164, 2016. DOI: [10.1016/j.orl.2015.12.014](https://doi.org/10.1016/j.orl.2015.12.014).
- [61] Panos M. Pardalos and Georg Schnitger. Checking local optimality in constrained quadratic programming is np-hard. *Oper. Res. Lett.*, 7(1):33–35, 1988. DOI: [10.1016/0167-6377\(88\)90049-1](https://doi.org/10.1016/0167-6377(88)90049-1).
- [62] Kim Allemand, Komei Fukuda, Thomas M. Lieblich, and Erich Steiner. A polynomial case of unconstrained zero-one quadratic optimization. *Math. Program.*, 91(1):49–52, 2001. DOI: [10.1007/s101070100233](https://doi.org/10.1007/s101070100233).
- [63] Milan Hladík, Michal Černý, and Miroslav Rada. A new polynomially solvable class of quadratic optimization problems with box constraints. *arXiv:1911.10877*, 2019. URL <https://arxiv.org/abs/1911.10877>.
- [64] Jacek Gondzio and Andreas Grothey. Solving nonlinear financial planning problems with 10^9 decision variables on massively parallel architectures. *WIT Trans Modelling Simul.*, 43:11, 2006. DOI: [10.2495/CF060101](https://doi.org/10.2495/CF060101).
- [65] Svatopluk Poljak, Franz Rendl, and Henry Wolkowicz. A recipe for semidefinite relaxation for (0, 1)-quadratic programming. *J. Glob. Optim.*, 7(1):51–73, 1995. DOI: [10.1007/BF01100205](https://doi.org/10.1007/BF01100205).
- [66] Joran Van Apeldoorn, András Gilyén, Sander Gribbling, and Ronald de Wolf. Quantum SDP solvers: Better upper and lower bounds. *Quantum*, 4:230, 2020. DOI: [10.22331/q-2020-02-14-230](https://doi.org/10.22331/q-2020-02-14-230).
- [67] Fernando G. S. L. Brandão, Amir Kalev, Tongyang Li, Cedric Yen-Yu Lin, Krysta M. Svore, and Xiaodi Wu. Quantum SDP Solvers: Large Speed-Ups, Optimality, and Applications to Quantum Learning. In *46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*, volume 132 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 27:1–27:14, Dagstuhl, Germany, 2019. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. ISBN 978-3-95977-109-2. DOI: [10.4230/LIPIcs.ICALP.2019.27](https://doi.org/10.4230/LIPIcs.ICALP.2019.27).
- [68] Nai-Hui Chia, Tongyang Li, Han-Hsuan Lin, and Chunhao Wang. Quantum-Inspired Sublinear Algorithm for Solving Low-Rank Semidefinite Programming. In *45th International Symposium on Mathematical Foundations of Computer Science (MFCS 2020)*, volume 170 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 23:1–23:15, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. ISBN 978-3-95977-159-7. DOI: [10.4230/LIPIcs.MFCS.2020.23](https://doi.org/10.4230/LIPIcs.MFCS.2020.23).
- [69] Jacek Gondzio. Warm start of the primal-dual method applied in the cutting-plane scheme. *Math. Program.*, 83(1-3):125–143, 1998. DOI: [10.1007/BF02680554](https://doi.org/10.1007/BF02680554).
- [70] Andrew Lucas. Ising formulations of many NP problems. *Front. Phys.*, 2:5, 2014. DOI: [10.3389/fphy.2014.00005](https://doi.org/10.3389/fphy.2014.00005).
- [71] Bas Lodewijks. Mapping np-hard and NP-complete optimisation problems to quadratic unconstrained binary optimisation problems. *arXiv:1911.08043*, 2019. URL <https://arxiv.org/abs/1911.08043>.
- [72] Jean B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM J. Optim.*, 11(3):796–817, 2001. DOI: [10.1137/S1052623400366802](https://doi.org/10.1137/S1052623400366802).
- [73] Jean B. Lasserre. Convergent SDP-relaxations

- in polynomial optimization with sparsity. *SIAM J. Optim.*, 17(3):822–843, 2006. DOI: [10.1137/05064504X](https://doi.org/10.1137/05064504X).
- [74] Bissan Ghaddar, Juan C. Vera, and Miguel F. Anjos. Second-order cone relaxations for binary quadratic polynomial programs. *SIAM J. Optim.*, 21(1):391–414, 2011. DOI: [10.1137/100802190](https://doi.org/10.1137/100802190).
- [75] Moses Charikar and Anthony Wirth. Maximizing quadratic programs: Extending Grothendieck’s inequality. In *45th Annual IEEE Symposium on Foundations of Computer Science*, pages 54–60. IEEE, 2004. DOI: [10.1109/FOCS.2004.39](https://doi.org/10.1109/FOCS.2004.39).
- [76] Mikhail Krechetov, Jakub Mareček, Yury Maximov, and Martin Takáč. Entropy-penalized semidefinite programming. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, 2019. DOI: [10.24963/ijcai.2019/157](https://doi.org/10.24963/ijcai.2019/157).
- [77] Sartaj Sahni and Teofilo Gonzalez. P-complete approximation problems. *J. ACM*, 23(3):555–565, 1976. DOI: [10.1145/321958.321975](https://doi.org/10.1145/321958.321975). See Lemma A2.
- [78] Michael Mitzenmacher and Eli Upfal. *Probability and computing: Randomization and probabilistic techniques in algorithms and data analysis*. Cambridge university press, 2017.
- [79] Sepehr Abbasi-Zadeh, Nikhil Bansal, Guru Guruganesh, Aleksandar Nikolov, Roy Schwartz, and Mohit Singh. Sticky brownian rounding and its applications to constraint satisfaction problems. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 854–873. SIAM, 2020. DOI: [10.1137/1.9781611975994.52](https://doi.org/10.1137/1.9781611975994.52).
- [80] Ronen Eldan and Assaf Naor. Krivine diffusions attain the goemans–williamson approximation ratio. *arXiv:1906.10615*, 2019. URL <https://arxiv.org/abs/1906.10615>.
- [81] Jamie Morgenstern, Samira Samadi, Mohit Singh, Uthaiapon Tantipongpipat, and Santosh Vempala. Fair dimensionality reduction and iterative rounding for SDPs. *arXiv:1902.11281*, 2019. URL <https://arxiv.org/abs/1902.11281v1>.
- [82] Samuel Karlin and Howard E. Taylor. *A second course in stochastic processes*. Elsevier, 1981. p. 257 and the following.
- [83] Julia Kempe, Oded Regev, and Ben Toner. The unique games conjecture with entangled provers is false. In *Algebraic Methods in Computational Complexity*, 2007.
- [84] Julia Kempe, Oded Regev, and Ben Toner. Unique games with entangled provers are easy. *SIAM J. Comput.*, 39(7):3207–3229, 2010. DOI: [10.1137/090772885](https://doi.org/10.1137/090772885).
- [85] Charles H. Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. Strengths and weaknesses of quantum computing. *SIAM J. Comput.*, 26(5):1510–1523, 1997. DOI: [10.1137/S0097539796300933](https://doi.org/10.1137/S0097539796300933).
- [86] Harry Markowitz. Portfolio selection. *J. Finance*, 7(1):77–91, 1952. DOI: [10.2307/2975974](https://doi.org/10.2307/2975974).
- [87] H. Abraham et al. Qiskit: An open-source framework for quantum computing, 2019. URL <https://doi.org/10.5281/zenodo.2562111>.
- [88] Johan Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001. DOI: [10.1145/502090.502098](https://doi.org/10.1145/502090.502098).
- [89] Vishwanathan Akshay, Hariphan Philathong, Igor Zacharov, and Jacob D. Biamonte. Reachability deficits implicit in google’s quantum approximate optimization of graph problems, 2020. URL <https://arxiv.org/abs/2007.09148>.
- [90] Rebekah Herrman, James Ostrowski, Travis S. Humble, and George Siopsis. Lower bounds on circuit depth of the quantum approximate optimization algorithm. *Quantum Inf. Process.*, 20(2):59, Feb 2021. DOI: [10.1007/s11128-021-03001-7](https://doi.org/10.1007/s11128-021-03001-7).
- [91] Zhihui Wang, Stuart Hadfield, Zhang Jiang, and Eleanor G. Rieffel. Quantum approximate optimization algorithm for maxcut: A fermionic view. *Phys. Rev. A*, 97:022304, Feb 2018. DOI: [10.1103/PhysRevA.97.022304](https://doi.org/10.1103/PhysRevA.97.022304).
- [92] Leo Zhou, Sheng-Tao Wang, Soonwon Choi, Hannes Pichler, and Mikhail D. Lukin. Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices. *Phys. Rev. X*, 10:021067, Jun 2020. DOI: [10.1103/PhysRevX.10.021067](https://doi.org/10.1103/PhysRevX.10.021067).
- [93] Jason Larkin, Matías Jonsson, Daniel Justice, and Gian Giacomo Guerreschi. Evaluation of quantum approximate optimization algorithm based on the approximation ratio of single samples, 2020. URL <https://arxiv.org/abs/2006.04831>.
- [94] qiskit-optimization. <https://github.com/Qiskit/qiskit-optimization>. Accessed: 25. 04. 2021.
- [95] Andreas Bärttschi and Stephan Eidenbenz. Grover mixers for qaoa: Shifting complexity from mixer design to state preparation. In *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)*, pages 72–82, 2020. DOI: [10.1109/QCE49297.2020.00020](https://doi.org/10.1109/QCE49297.2020.00020).
- [96] Reuben Tate, Majid Farhadi, Creston Herold, Greg Mohler, and Swati Gupta. Bridging classical and quantum with SDP initialized warmstarts for QAOA, 2020. URL <https://arxiv.org/abs/2010.14021>.
- [97] Iain Dunning, Swati Gupta, and John Silberholz. What works best when? a systematic evaluation of heuristics for Max-Cut and QUBO.

- INFORMS J. Comput.*, 30(3):608–624, 2018. DOI: [10.1287/ijoc.2017.0798](https://doi.org/10.1287/ijoc.2017.0798).
- [98] Panagiotis Kl. Barkoutsos, Jerome F. Gonthier, Igor Sokolov, Nikolaj Moll, Gian Salis, Andreas Fuhrer, Marc Ganzhorn, Daniel J. Egger, Matthias Troyer, Antonio Mezzacapo, and *et al.* Quantum algorithms for electronic structure calculations: Particle-hole hamiltonian and optimized wave-function expansions. *Phys. Rev. A*, 98:022322, Aug 2018. DOI: [10.1103/PhysRevA.98.022322](https://doi.org/10.1103/PhysRevA.98.022322).
 - [99] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of np. *J. ACM*, 45(1):70–122, 1998. DOI: [10.1145/273865.273901](https://doi.org/10.1145/273865.273901).
 - [100] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998. DOI: [10.1145/278298.278306](https://doi.org/10.1145/278298.278306).
 - [101] Irit Dinur. The PCP theorem by gap amplification. *J. ACM*, 54(3):12–es, jun 2007. DOI: [10.1145/1236457.1236459](https://doi.org/10.1145/1236457.1236459).
 - [102] Sanjeev Arora and Boaz Barak. *Computational complexity: a modern approach*. Cambridge University Press, 2009.
 - [103] Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing*, STOC '02, pages 767–775, New York, NY, USA, 2002. Association for Computing Machinery. DOI: [10.1145/509907.510017](https://doi.org/10.1145/509907.510017).
 - [104] Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 245–254, 2008. DOI: [10.1145/1374376.1374414](https://doi.org/10.1145/1374376.1374414).
 - [105] Prasad Raghavendra and David Steurer. How to round any CSP. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 586–594, 2009. DOI: [10.1109/FOCS.2009.74](https://doi.org/10.1109/FOCS.2009.74).
 - [106] Subhash Khot, Dor Minzer, and Muli Safra. Pseudorandom sets in grassmann graph have near-perfect expansion. In *Proceedings of the fifty-ninth Annual Symposium on Foundations of Computer Science (FOCS)*, pages 592–601, 2018. DOI: [10.1109/FOCS.2018.00062](https://doi.org/10.1109/FOCS.2018.00062).
 - [107] Boaz Barak, Prasad Raghavendra, and David Steurer. Rounding semidefinite programming hierarchies via global correlation. In *Proceedings of the fiftysecond annual symposium on foundations of computer science*, pages 472–481. IEEE, 2011. DOI: [10.1109/FOCS.2011.95](https://doi.org/10.1109/FOCS.2011.95).
 - [108] Samuel B. Hopkins, Tselil Schramm, and Luca Trevisan. Subexponential LPs approximate max-cut. In *Proceedings of the sixtyfirst Annual Symposium on Foundations of Computer Science (FOCS)*, pages 943–953. IEEE, 2020. DOI: [10.1109/FOCS46700.2020.00092](https://doi.org/10.1109/FOCS46700.2020.00092).
 - [109] Albert Einstein, Boris Podolsky, and Nathan Rosen. Can quantum-mechanical description of physical reality be considered complete? *Phys. Rev.*, 47(10):777, 1935. DOI: [10.1103/PhysRev.47.777](https://doi.org/10.1103/PhysRev.47.777).
 - [110] Boris S. Cirel'son. Quantum generalizations of Bell's inequality. *Lett. Math. Phys.*, 4(2):93–100, 1980. DOI: [10.1007/BF00417500](https://doi.org/10.1007/BF00417500).
 - [111] A. Natarajan and T. Vidick. Low-degree testing for quantum states, and a quantum entangled games PCP for QMA. In *Proceedings of the fifty-ninth Annual Symposium on Foundations of Computer Science (FOCS)*, pages 731–742, 2018. DOI: [10.1109/FOCS.2018.00075](https://doi.org/10.1109/FOCS.2018.00075).
 - [112] Dorit Aharonov, Itai Arad, Zeph Landau, and Umesh Vazirani. The detectability lemma and quantum gap amplification. In *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, STOC '09, pages 417–426, New York, NY, USA, 2009. Association for Computing Machinery. ISBN 9781605585062. DOI: [10.1145/1536414.1536472](https://doi.org/10.1145/1536414.1536472).
 - [113] Moses Charikar, Konstantin Makarychev, and Yury Makarychev. Near-optimal algorithms for unique games. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 205–214, 2006. DOI: [10.1145/1132516.1132547](https://doi.org/10.1145/1132516.1132547).
 - [114] Dimitris Achlioptas, Assaf Naor, and Yuval Peres. Rigorous location of phase transitions in hard optimization problems. *Nature*, 435(7043):759–764, 2005. DOI: [10.1038/nature03602](https://doi.org/10.1038/nature03602).
 - [115] Don Coppersmith, David Gamarnik, Mohammad T. Hajiaghayi, and Gregory B. Sorkin. Random MAX SAT, random MAX CUT, and their phase transitions. *Random Struct. Algorithms*, 24(4):502–545, 2004. DOI: [10.1002/rsa.20015](https://doi.org/10.1002/rsa.20015).