

Fermilab

Iterative quantum optimization of spin glass problems with rapidly oscillating transverse fields

FERMILAB-PUB-24-0901-SQMS-V

arXiv:2408.06571

This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

Iterative quantum optimization of spin glass problems with rapidly oscillating transverse fields

Brandon Barton^{1,2,3†}, Jacob Sagal^{1†}, Sean Feeney¹, George Grattan^{1,5}, Pratik Patnaik⁴, Vadim Oganessian^{3,6,7}, Lincoln D Carr^{1,2,4} and Eliot Kapit^{1,4}

¹ Quantum Engineering Program, Colorado School of Mines, 1523 Illinois St, Golden, CO 80401, United States of America

² Department of Applied Mathematics and Statistics, Colorado School of Mines, 1500 Illinois St, Golden, CO 80401, United States of America

³ Center for Computational Quantum Physics, Flatiron Institute, 162 5th Avenue, New York, NY 10010, USA

⁴ Department of Physics, Colorado School of Mines, 1523 Illinois St, Golden, CO 80401, United States of America

⁵ Department of Computer Science, Colorado School of Mines, 1500 Illinois St, Golden CO 80401, United States of America

⁶ Department of Physics and Astronomy, College of Staten Island, City University of New York, Staten Island, NY 10314, United States of America

⁷ Physics program and Initiative for the Theoretical Sciences, The Graduate Center, City University of New York, New York, NY 10016, United States of America

E-mail: ekapit@mines.edu

July 2024

Abstract. In this work, we introduce a new iterative quantum algorithm, called Iterative Symphonic Tunneling for Satisfiability problems (IST-SAT), which solves quantum spin glass optimization problems using high-frequency oscillating transverse fields. IST-SAT operates as a sequence of iterations, in which bitstrings returned from one iteration are used to set spin-dependent phases in oscillating transverse fields in the next iteration. Over several iterations, the novel mechanism of the algorithm steers the system toward the problem ground state. We benchmark IST-SAT on sets of hard MAX-3-XORSAT problem instances with exact state vector simulation, and report polynomial speedups over trotterized adiabatic quantum computation (TAQC) and the best known semi-greedy classical algorithm. When IST-SAT is seeded with a sufficiently good initial approximation, the algorithm converges to exact solution(s) in a polynomial number of iterations. Our numerical results identify a critical Hamming radius(CHR), or quality of initial approximation, where the time-to-solution crosses from exponential to polynomial scaling in problem size. By combining IST-SAT with future classical or quantum approximation algorithms, larger gains may be achieved. The mechanism we present in this work thus presents a new path toward achieving quantum advantage in optimization.

Keywords: quantum optimization, heuristic quantum algorithms, combinatorial optimization, satisfiability problems, macroscopic quantum tunneling.

1. Introduction

Binary constraint satisfaction problems represent a promising opportunity to achieve practical quantum advantage in real-world problems found in optimization, artificial intelligence, and cryptography. In the worst case, and often, the typical case, these problems are exponentially hard to solve or even approximate for classical machines. Therefore, an efficient solution mechanism would have broad impacts and applicability.

So far however, demonstrating consistent quantum advantage with heuristic algorithms has remained elusive. A wide array of heuristic quantum algorithms such as analog quantum annealing [1, 2, 3, 4, 5, 6, 7], adiabatic quantum computing (AQC) [8, 9], quantum approximation optimization algorithms (QAOA) [10], and more exotic variations such as adaptive derivative assembled problem tailored (ADAPT)-QAOA [11], recursive QAOA [12], QAOA supplemented with amplitude amplification [13], and energy matching/population transfer algorithms [14, 15, 16, 17], have been proposed for spin glass optimization. Their empirical performance, however, has been decidedly mixed, and the frequently observed quadratic speedups from schedule optimization (a scheme dating back to the adiabatic formulation of Grover’s search [18]) are fragile and likely not feasible at large N [19]. Furthermore, if the problem Hamiltonian is stoquastic [20], quantum Monte Carlo [21, 22, 23, 24, 25] can simulate small systems sizes in the quadratic (incoherent) scaling limit.

Motivated by the recent search for applications of quantum algorithms to machine learning, variational quantum algorithms have also been applied to spin glass optimization problems [26]. Yet, large gains from these methods may be limited by the cost of computing gradients and may often get stuck in local minima and encounter barren plateaus, or loss of variance in the optimization landscape [27, 28]. Thus, the full capabilities of quantum computing in this space, and potential asymptotic limits for very general algorithms in this class, deserve further exploration. Given that beyond-quadratic speedups are critical to achieve useful quantum advantage [29], new mechanisms that expand the quantum optimization toolbox for reaching exact solutions [30], or even approximate solutions may have significant impacts.

In this work, we propose a non-classical steering mechanism that guides quantum optimization algorithms towards the true ground state in spin glass problems. We demonstrate this mechanism by introducing a new heuristic quantum algorithm which we call *Iterative Symphonic Tunneling for Satisfiability problems* (IST-SAT). The IST-SAT algorithm modifies quasi-continuous time Trotterized AQC (TAQC), with total evolution time increasing linearly with N , by adding a monochromatic fast oscillating field along Y to all spins. The addition of the oscillating field is inspired by a recently discovered mechanism termed *symphonic tunneling* which has exponentially accelerated macroscopic quantum tunneling rates between ferromagnetic ground states [31, 32]. IST-

SAT builds on this work by expanding the use of high frequency oscillating transverse fields to quantum optimization. We consider disordered spin glass instances drawn from the MAX-3-XORSAT problem class [30], where the problem Hamiltonian H_P is diagonal in Z and the DC transverse field is along X . In contrast to accelerating collective tunneling in the transverse-field Ising model (TFIM) where the classical ground states are known a priori and there is no frustration [32], applying a uniform oscillating field with identical frequencies and phases to all sites produces no meaningful benefits for finding low energy states of MAX-3-XORSAT problems.

To formulate the IST-SAT algorithm, we propose an iterative strategy, where bitstrings returned from TAQC are used to set spin-dependent phases for the oscillating drive terms in subsequent iterations of the algorithm. We define the list of phases $P = [\varphi_1, \varphi_2, \dots, \varphi_N]$, with $\varphi_j \in \{0, \pi\}$ assigned to each spin. We test the performance of IST-SAT on the planted partial solution problem (PPSP) class defined in [33], and with extensive exact state-vector simulations. Our results show that when *the phase pattern* more closely matches the pattern of bits in the planted classical ground state G , the probabilities of reaching G and states nearby to the global optima increase monotonically. Our simulations identify a *critical Hamming radius* (CHR) r_c where the time to find G in iterative optimization scales polynomially, starting from any bitstring with Hamming distance $D_H \leq r_c N$ from G . IST-SAT achieves a polynomial speedup over both trotterized AQC (TAQC) and the best-known quasi-greedy classical algorithm for this problem. We also demonstrate that using TAQC as a seed algorithm for classical methods yields meaningful polynomial speedups over both algorithms on their own. Taken together, these results suggest new and promising routes to gradient-free quantum optimization, through a new non-classical steering mechanism to find ground-states with quantum algorithms.

2. Problem and algorithm definitions

In contrast to many problems where asymptotic exponential scaling is not reached until prohibitively large system sizes for classical computers and NISQ era devices [34], 3-XORSAT problems have exponential scaling that is typically observed at small N . We test IST-SAT on the MAX-3-XORSAT [30] problem, defined on a random 3-uniform hypergraph consisting of N_C three-body constraint terms. The problem Hamiltonian is given by

$$H_P = - \sum_{ijk}^{N_C} V_{ijk} Z_i Z_j Z_k, \quad V_{ijk} = \pm 1, \quad (1)$$

where for a given bitstring a constraint is satisfied if $\langle V_{ijk} Z_i Z_j Z_k \rangle = +1$, and unsatisfied otherwise. Since the problem is linear, Gaussian elimination can be used to check if the problem is satisfiable in polynomial time. However, when not all the constraints can be satisfied, finding the lowest energy state(s) is known to be NP-hard [35].

This work utilizes a family of instances called planted partial solution problems

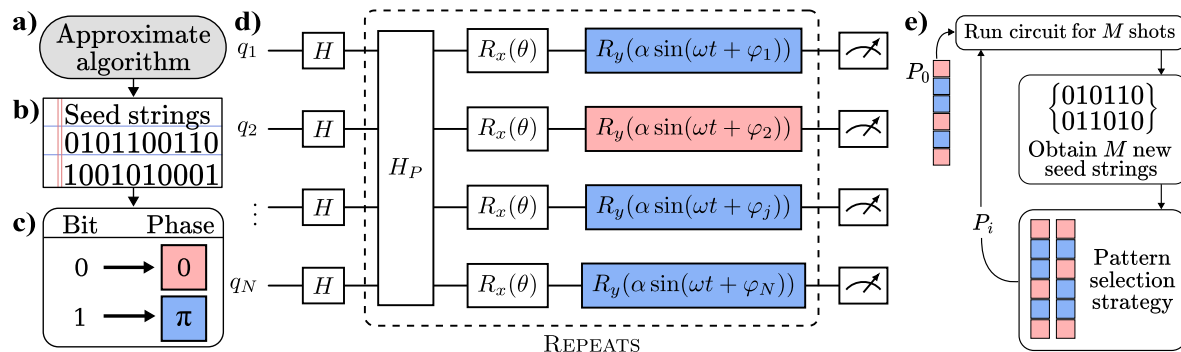


Figure 1. Schematic of the IST-SAT algorithm. An approximate algorithm—which can include random guessing and high-depth TAQC—in (a) is used to find “seed strings” (b) for the phase pattern in the H_{ST} drive Hamiltonian. The local phases $\{\varphi_1, \varphi_2, \dots, \varphi_N\}$ in H_{ST} are set from bits of the seed states using the key (c), which we denote with colored R_y rotations in (d). After providing the algorithm an initial phase pattern P_0 , the iterative process described in (e) continues, where the circuit is run for M shots, generating new seed strings which are converted to phase patterns. A new phase pattern P_i is then selected to be used in the next iteration. The total number of “REPEATS” is set by the ratio of total runtime and Trotter step, t/dt (see text).

(PPSPs), used in recent quantum approximation algorithms for MAX-3-XORSAT [33]. To construct a PPSP we first pick a random hypergraph with N variables that participate in N_C constraints with $N_C \gg N$, a fixed small fraction ϵ (we use $\epsilon = 0.1$ in this work), and a random bitstring G to be the planted ground state. We randomly select $(1 - \epsilon)N_C$ constraints to be satisfied in G (by choosing the signs of the V_{ijk}) with the remaining left unsatisfied. For small ϵ and $N_C \gg N$, this construction makes G a unique ground state with high probability (the SAT/UNSAT transition here is at $N_C/N \sim 0.92$ [36]). We note that the PPSP instances in this work can be more difficult than 3-regular hypergraphs studied in previous work [37, 38], which are easy to approximate and can be solved efficiently if satisfiable. Throughout this work, we used the set of constraint densities $N_C/N \in \{1.5, 2, 4\}$ to test the performance of IST-SAT over several densities above the SAT/UNSAT transition.

IST-SAT starts from the “standard” quasi-continuous time AQC method of interpolating between a transverse field “driver” Hamiltonian and the problem Hamiltonian with Trotterized evolution from $t = 0$ to $t = t_f$:

$$H(t) = f(t)H_D + g(t)H_P, \quad H_D = -\sum_j X_j, \quad (2)$$

$$f(0) = g(t_f) = 1, \quad f(t_f) = g(0) = 0.$$

We interpolate between the problem and driver Hamiltonians using the following functions

$$f(t) = \sqrt{1 - t/t_f}, \quad g(t) = \sqrt{t/t_f}, \quad (3)$$

The choices of $f(t)$ and $g(t)$ empirically outperform linear interpolation, producing a better prefactor and modestly better scaling with N , though we expect the asymptotic

scaling of the two schedules may converge for this problem. We let the total evolution time t_f grow linearly with N .

IST-SAT modifies the base Hamiltonian in equation 2 by adding a high-frequency monochromatic oscillating field

$$H_{\text{ST}}(t) = \alpha \sum_j Y_j \sin(\omega t + \varphi_j). \quad (4)$$

Within $H_{\text{ST}}(t)$, the parameters $\varphi_j = \{0, \pi\}$ are single-qubit phases, and $\alpha = \alpha_s \ln N$ is the drive strength. The total Hamiltonian is then given by

$$H(t) = f(t)H_D + g(t)H_P + h(t)H_{\text{ST}}(t), \quad (5)$$

where $h(t)$ is a smooth function with initial and final conditions $h(t=0), h(t_f) = 0$. In this work, we used $h(t) = 4\sqrt{(1-t/t_f)(t/t_f)}$, $\alpha_s = 0.6$ and $\omega = 2\pi \times 6 \ln N$. The choice of α and ω both increasing logarithmically with N minimizes heating [39, 40] while ensuring that the novel low-frequency terms generated by the high frequency drive have constant magnitude. Time evolution is implemented via Trotterization and gate-based exact state-vector simulation using the qulacs python package [41].

The schematic work-flow of the IST-SAT algorithm is shown in Fig. 1, which begins with an approximate algorithm. The initial approximation algorithm can be a greedy classical approach, simulated annealing, or a quantum algorithm such as AQC/TAQC or QAOA. The initial approximate algorithm is used to obtain *seed strings* which are used to set phases in the high frequency AC drive $H_{\text{ST}}(t)$. We note the method of phase selection in H_{ST} is distinct from *warm start* methods [42], where preparing a good initial state with high overlap to the ground state is costly, and still expected to be hard. Instead, IST-SAT starts from the initial superposition state $|+\rangle^{\otimes N}$, and sets parameters in the time-dependent Hamiltonian $H_{\text{ST}}(t)$. In one iteration of IST-SAT, the binary configuration of a selected seed string is used to form the phase pattern P , which sets parametrizes the quantum circuit in figure 1. After evolving the circuit, new measurements in the z -basis generate M new phase patterns that may be used in further iterations.

For the time evolution of equation 5, we used a mean evolution time of $t_f = N/32$, with $dt = 0.4/\omega$ to ensure high frequencies are appropriately sampled. The total circuit depth thus scales as $N \ln N$ problem Hamiltonian applications. All data is averaged over 1000 random problem instances for each N . The total run-time is total averaged between $T = 2t_f/3$ and $T = 4t_f/3$, for all simulations of IST-SAT and TAQC. This averaging procedure has shown to effectively smooth out unpredictable diabatic effects that can make reliably estimating scaling difficult [31], particularly when success probabilities decay exponentially with N . The parameters used in this work are the result of trial and error on small system sizes. Therefore further parameter optimization techniques are expected to produce further benefits.

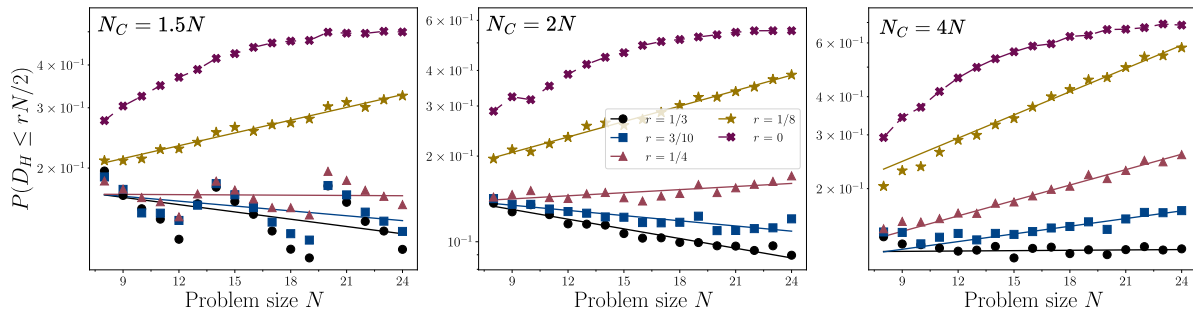


Figure 2. Exponential to polynomial scaling transition for IST-SAT seeded with random initial phase patterns P_0 , with fraction $(1-r)$ of the bits matching the planted solution. For constraint densities $N_C/N = 1.5, 2, 4$ (left, middle, right), the plot shows the probability of returning a solution within $D_H \leq rN/2$ flips of the nearest solution. For probability calculations with $D_H \leq rN$, refer to Appendix A.

3. Performance results of IST-SAT

The results from state-vector simulation of IST-SAT shown in figure 2, demonstrate that an appropriately chosen phase pattern dramatically accelerates a collective rearrangement of spins at the transition from the initial paramagnetic state to ground states of the problem Hamiltonian. Of course, this choice of phase pattern is not obvious since solutions to the problem are generally not known, and the purpose of the algorithm is to find it. If the phases are chosen randomly, the protocol in equation 5 shows no scaling advantage over the uniform field protocol in equation 2 with a modest prefactor disadvantage. Conversely, if we let $\varphi_j = \pi b_j$, where $b_j \in \{0, 1\}$ is the value of bit j in the planted solution G , $P_{GS}(t_f)$ is empirically *constant* with linear runtime $t_f \propto N$ (see the $r = 0$ data in figure 2), compared to exponential decay in the uniform field case shown in Appendix B, with an exponent predicted analytically in [33]. Making this choice requires knowledge of G , and thus a solution to the problem, so it's not obvious that this discovery will help.

The observation that supports IST-SAT is as follows: suppose the phases φ_j are guessed correctly with some probability $r \geq 1/2$, based on the values of each bit in G . We note that the relative phases in P are most crucial. As r increases toward 1, the probabilities of finding any approximate solution nearest in Hamming distance monotonically increase at equivalent total evolution time, as shown in figure 2. Therefore, using the bitstrings from previous iterations of IST-SAT to construct the phase pattern for subsequent iterations yields solutions that are progressively closer to the ground state, enabling the algorithm to converge rapidly.

The results obtained from simulation in this work consistently present an empirical critical Hamming radius (CHR) r_c , which depends on N_C/N and the unsatisfied fraction ϵ in the planted solution. We define the CHR r_c in the following way: suppose each local phase is guessed correctly (relative to the planted solution) with probability $1-r$. Then, for linearly growing t_f , r_c is defined to be the largest value of r such that the probability of returning states with a Hamming distance $\leq rN$ is constant with increasing N . For

$N_C/N = \{1.5, 2, 4\}$ and $\epsilon = 0.1$, we respectively observe $r_c \simeq \{1/4, 1/4, 1/3\}$, which demonstrates that problems with larger constraint densities are solved with comparably worse initial approximations. In simulations of smaller systems with larger constraint densities (data not shown), the CHR are observed to be consistent with $r_c \simeq 1/3$.

Therefore, if the approximate seed strings are within some $r < r_c$, the bitstrings returned each round in IST-SAT will get closer and closer to the global optimum; with r_c defined in this way, we expect convergence from a string less than $r_c N$ flips to the planted solution in a polynomial number of algorithm iterations. If we start from strings which are more than r_c flips away, iterating phase patterns are still expected to converge to a global optima faster than the TAQC algorithm IST-SAT uses as a starting point. This is demonstrated by a smaller scaling exponent which can be inferred from the smaller slopes for curves with $r_c < r < 1/2$ in Appendix B as compared to the TAQC data, though the quantitative analysis of scaling is more complex in that case.

Throughout this paper, we define the time-to-solution (TTS) as the exponent b obtained from exponential fits to numerical data presented in this work. We use the function $f(N) = a2^{bN}$, where a is a pre-factor constant, and b is the exponent which indicates the expected number of shots it takes to reach the a ground state of the problem. Hence, when $b \geq 0$, IST-SAT returns solutions scaling as a polynomial in N , while $b < 0$ implies an exponentially increasing with number of iterations to reach a global minima.

Table 1. Inferred time-to-solution (TTS) to reach a global optima for various algorithms: a semi-greedy classical (SGC) algorithm [37], TAQC [10], SGC seeded with TAQC, and IST-SAT seeded with TAQC. The TTS (defined above in section 3) is derived from the scaling exponent b in numerical fits to simulation data.

Algorithm set-up	$N_C/N = 1.5$	$N_C/N = 2$	$N_C/N = 4$
SGC	-0.092	-0.15	-0.050
TAQC	-0.11	-0.17	-0.20
TAQC→SGC	-0.073	-0.081	-0.025
TAQC→IST-SAT	-0.026	-0.053	-0.027

For the TAQC→SGC sequence in table 1, we report the combination with the least exponentially decaying exponent from Appendix D. For the TAQC→IST-SAT sequence, we report the TAQC exponent from Appendix B corresponding to the approximation distance equal to the critical Hamming radius ($d = r_c$) of IST-SAT. The full set of exponents for other distances $d \neq r_c$ may be found in tables in the referenced appendices. The inferred TTS demonstrates that IST-SAT obtains a significant polynomial speed-up over SGC and TAQC at all constraint densities. The performance of TAQC→IST-SAT out-performs TAQC→SGC at $N_C/N = \{1.5, 2\}$, while demonstrating comparatively worse performance at $N_C/N = 4$. We attribute this difference due to the warm-started SGC algorithm not reaching expected asymptotic exponential decay at $N_C/N = 4$ until system sizes beyond our simulation capabilities

($N > 200$), while TAQC and IST-SAT demonstrate the expected asymptotic scaling in small problem sizes. Further performance results of the warm-started SGC algorithm may be found in Appendix D.




Figure 3. Average energy $\langle E \rangle$ returned from IST-SAT with a lower frequency drive $\omega 10\pi \ln N$ (Appendix E) and TAQC, normalized by the ground state energy E_{GS} . Energy is obtained using H_P as the cost function. Here, we run a single iteration of IST-SAT, and use 10,000 samples (bitstrings) from the wave function in exact state vector simulation. The legend in the center figure specifies the fraction of incorrect bits rN used in the seed state protocol.

To examine the *quality of strings* returned from IST-SAT, we plot the average energy $\langle E \rangle$ (normalized by E_{GS}) in figure 3. We note that the data reported in figure 3 were obtained using a lower frequency of $\omega = 10\pi \ln N$, compared to the data reported in figure 2. Interestingly, the rapid convergence to any global optima is based only on Hamming distance D_H and remains uncorrelated with the energies E of the returned states. However, we still observe monotonically increasing average string qualities as IST-SAT is seeded with better phase patterns. Our very first formulation of IST-SAT used a variation of the quasi-greedy algorithm in [37] to gather the initial seed states, and we observed no improvement over random guessing in making this choice for reasons discussed below. This is supported by the result that, unless r is very close to one, the average energy returned by IST-SAT is somewhat worse than that of TAQC for all other parameters equal (see the supplemental information for figures). We attribute this observation to there being many local minima with energies close to E_{GS} , and the effect of the high frequency oscillating drive is to steer the evolving state away from them and toward $|G\rangle$ and its local excitations, which may be higher in energy even if they are much closer in Hamming distance.

While naive classical seeding was unsuccessful, matching the observed $r_c \simeq \{1/4, 1/4, 1/3\}$ with the corresponding distance probabilities for TAQC yields significantly reduced asymptotic time to solution if strings from TAQC (or equivalently, random phase patterns) are used for initial seeding. Selecting fractions of the TAQC strings based on their energies yielded no benefit over randomly sampling the TAQC output distribution.

Interestingly, the results demonstrate that even when the initial seed string is chosen relative to a single solution (the planted solution), the probability of reaching

other equally valid solutions remains robust. This result is clearly observed at the smallest constraint density ($N_C/N = 1.5$), where each problem instance tends to have several global optima. At larger constraint densities of $N_C/N = \{2, 4\}$, the probability of reaching any solution more closely matches the probability of reaching the planted solution, which is unique with higher probability. In Appendix A, we discuss the choice of phase pattern, convergence to the planted solution in particular, and the exact solution statistics of the problem instances used in this work.

4. Conclusion

We have introduced IST-SAT, a quantum algorithm that uses a non-classical steering mechanism based on high frequency oscillating drives that guides the algorithm towards the global minima of hard optimization problems. IST-SAT is a unique feedback-based quantum algorithm which does not require the calculation of any gradients or averages used in the original variational QAOA proposal, or more recent versions such as recursive QAOA [12]. These costly methods require hundreds to thousands of shots to calculate accurately. The core mechanism for how the phase pattern appropriately guides the optimization is only understood by analogy to other work in a simpler system [32], where we expect a careful analytical derivation of the speedups observed in this work to inform further algorithm innovations.

Regarding seed algorithms for IST-SAT, it may be the case that other classical or quantum algorithms could provide further benefits to for finding good initial approximations. For example, an algorithm that breaks H_P into exactly solvable sub-problems may obtain sets of relative phase patterns which potentially yield significant improvements. Future work may also test versions of IST-SAT by assigning each group of spins a different frequency (not just a different phase pattern), for better averaging, as in the earlier multi-frequency AC optimization schemes [43, 19, 44, 31] that inspired this work.

While all phase patterns were set using discrete offsets of $\{0, \pi\}$ in this work, we expect that more sophisticated phase pattern selection protocols would further improve IST-SAT. An interesting future direction could consider using semi-definite programs to selecting phases in the H_{ST} drive in a continuous interval $\varphi_j = [0, \pi]$, building off recent improvements to a variational quantum algorithm [45] for quantum Max-Cut. We expect that the core mechanism of IST-SAT is broadly applicable to other problems in optimization, such as Max-Cut or low auto-correlated binary sequences (LABS), which we leave for future exploration. Finally, it is expected that this IST-SAT is not single-error fragile, in contrast to recently proposed quantum algorithms for MAX-3-XORSAT [33], thus motivating future experiments on near-term quantum devices.

Acknowledgments

We would like to thank Joao Basso and Ojas Parekh for insightful discussions. We would also like to thank Takuto Komatsuki and Joey Liu for support with quantum simulations on the Fujitsu Quantum Simulator. This work was supported by the DARPA Reversible Quantum Machine Learning and Simulation program under contract HR00112190068, as well as by National Science Foundation grants PHY-1653820, PHY-2210566, DGE-2125899, and by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Superconducting Quantum Materials and Systems Center (SQMS) under contract number DE-AC02-07CH11359. The SQMS Center supported EK's advisory role in this project, as well as his time improving and fine tuning the algorithms. Part of this research was performed while BB was visiting the Institute for Pure and Applied Mathematics (IPAM), which is supported by the National Science Foundation (Grant No. DMS-1925919). This work was performed in part at Aspen Center for Physics, which is supported by National Science Foundation grant PHY-2210452. The Flatiron Institute is a division of the Simons Foundation.

[†]BB and JS contributed equally to this work.

Appendix

Appendix A. Further details on performance of IST-SAT

Appendix A.1. Phase pattern selection and convergence to the planted solution

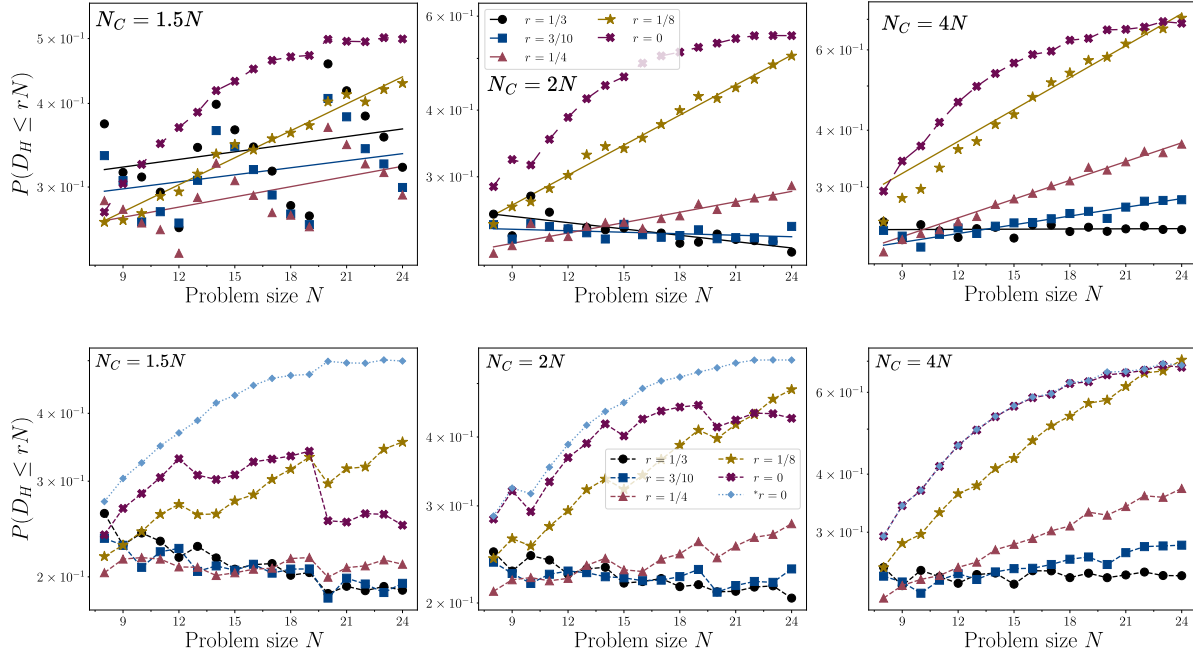


Figure A1. Convergence of IST-SAT to the (top row) nearest solution in Hamming distance and (bottom row) the planted solution. Initial phase patterns P_0 were chosen from random seed strings with $(1 - r)$ of the bits correctly chosen relative to the planted solution. We include the $r = 0$ case from the top row as the $r^* = 0$ case in the bottom row (blue diamonds) to show how IST-SAT targets the planted solution at high constraint densities.

We report the convergence of IST-SAT to the nearest solution in Hamming distance and the planted solution in particular shown in figure A1. Notably, the results for convergence to any solution at $N_C/N = 1.5$ demonstrate a positive exponent for all fractions r of the random approximate seed string. All exponential fits for the top row of figure A1 can be found in Appendix A.2. When considering the convergence to the planted solution in particular, the performance of IST-SAT notably degrades at $r = 0$. We attribute this behavior to the problem instances containing several global minima (multiple ground states) at lower constraint density, which is supported by exact solution statistics provided in Appendix A.3. As discussed in the main text, while performance of reaching the planted solution degrades, IST-SAT remains robust to finding other global optima equal in energy to the planted solution.

While low constraint densities have more ground states, the number of solutions is typically very small at larger constraint density ($N_C/N = 4$), where the convergence to any solution and the convergence to the planted solution are nearly identical (see right

column of figure A1). This alignment in convergence suggests that IST-SAT targets not only the solution to which the pattern was pre-selected for, but all equally valid solutions. Therefore, in the problem instances we tested in this work, a pattern selected relative to one solution remains robust to finding other global minima. These results are supported by exact solution statistics in Appendix A.3 for the MAX-3-XORSAT problem instances we consider in this work.

Appendix A.2. Performance of IST-SAT with random approximate seed strings

Table A1. Exponential fits to the numerical data presented for IST-SAT seeded with random approximate seed strings for probabilities $P(D_H \leq rN)$ (see top row of figure A1) and $P(D_H \leq rN/2)$ (see figure 2). In this table, we report the exponent b from fitting the numerical data to $a2^{bN}$. The exponents corresponding to the critical Hamming radius are highlighted in bold text.

(a)			
Guessing error r	$N_C/N = 1.5$	$N_C/N = 2$	$N_C/N = 4$
1/3	0.013	-0.013	0.00045
3/10	0.012	-0.0031	0.021
1/4	0.017	0.022	0.046
1/8	0.045	0.062	0.077
(b)			
Guessing error r	$N_C/N = 1.5$	$N_C/N = 2$	$N_C/N = 4$
1/3	-0.024	-0.038	0.0013
3/10	-0.016	-0.020	0.028
1/4	-0.00087	0.012	0.056
1/8	0.041	0.060	0.084

Appendix A.3. Exact solution statistics at small problem sizes




Figure A2. (a) Percentage of problem instances (from the total 1000 instances) with more than one ground state, including the PPS. For each system size $N = 8 - 24$, we computed exact ground states using brute-force methods. (b) Average number of ground states per problem instance. (c) Pairwise Hamming distance D_H within the set of ground states. The $N = 8, 9, 12$ cases have only one ground state, which is the PPS.

To produce figure A2, we calculated the exact number of exact solutions via brute force search to $N = 24$, for 1000 problem instances at each system size. In figure A2 (a-b), the problem instances at $N_C/N = \{1.5, 2\}$ tend to have more than one ground state, while at $N_C/N = 4$, the planted solution tends to be a unique ground state. These statistics support the results shown in figure A1, where the convergence to the any solution does not show any exponential decay for all approximate seed state distances. For the largest constraint density we present in this work ($N_C/N = 4$), the planted solution is nearly always the unique global optima.

When a given instance has multiple ground states, we further would like to understand how these solutions are correlated. In figure A2 (c) we report the pairwise Hamming distance between ground states, which demonstrates that even when a problem contains several solutions, the solutions tend to very close to each other measured in Hamming distance. This strong correlation demonstrates that solutions tend to be only a few bit flips away from each other on average. These statistics support the claim that when the phase patterns are selected only relative to one solution (the planted solution, which is known), IST-SAT is robust to finding other global minima due a small pairwise Hamming distance within the set of optimal solutions. Future work may consider other problems with instances in which there are no such strong correlations are expected.

Appendix B. Seeding with trotterized adiabatic quantum computation

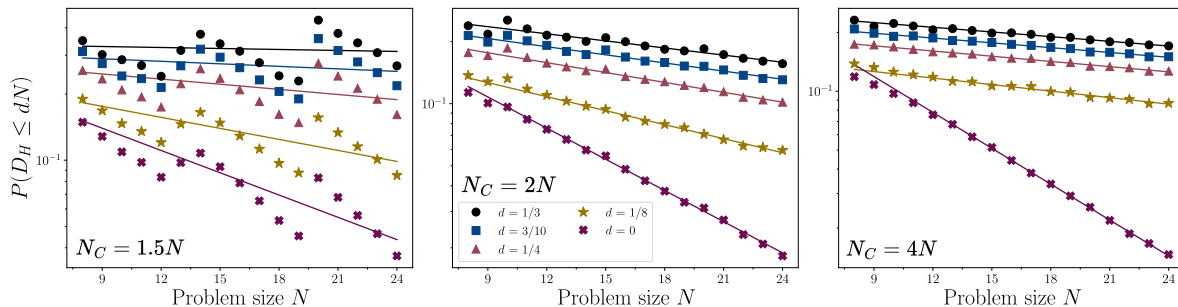


Figure B1. Probabilities of TAQC returning approximate strings within a Hamming distance $D_H \leq dN$ from the nearest solution for (left to right) $N_C/N = \{1.5, 2, 4\}$. Data are fit to lines $a2^{bN}$, where the exponents b are reported in Appendix B.

In this section, we consider the capabilities of TAQC as a seed algorithm for IST-SAT. Specifically, we report the approximation performance of the TAQC and SGC algorithms in terms of Hamming distance to any ground state. In figure B1, we show the probabilities of reaching approximate solutions, measured by an approximation in Hamming distance $D_H \leq dN$ to the nearest solution, where d is some fraction in the range $[0, 1/2]$. Exact solutions are measured by $d = 0$, while $d = 1/2$ is equivalent to random guessing, and $0 < d < 1/2$ are approximate solutions.

Table B1. Exponential fits to the numerical data presented for TAQC in figure B1. In this table, we report the parameters b obtained from fitting the function $a2^{bN}$ to the probabilities $P(D_H \leq dN)$ of finding states close in Hamming distance D_H for different distances dN or fewer flips away from the nearest solution. We include the fitting results for the set of constraint densities $N_C/N = \{1.5, 2, 4\}$. We denote the fraction d for each constraint density associated to r_c in bold face text.

Fractional Hamming distance d	$N_C/N = 1.5$	$N_C/N = 2$	$N_C/N = 4$
1/3	-0.0050	-0.038	-0.027
3/10	-0.013	-0.044	-0.027
1/4	-0.026	-0.053	-0.030
1/8	-0.056	-0.075	-0.037
$0(P_{GS})$	-0.11	-0.17	-0.20

Appendix C. Performance of the semi-greedy classical algorithm




Figure C1. Probabilities of SGC returning approximate strings within a Hamming distance $D_H \leq dN$ from the *planted solution* for (left to right) $N_C/N = \{1.5, 2, 4\}$. Data is plotted as $P(D_H \leq dN) + P_{GS}$ to estimate the probability of finding any ground state. Data are fit to lines $a2^{bN}$, where the exponents b are reported in Appendix C.

Table C1. Exponential fits to $a2^{bN}$ for the SGC data presented in figure C1.

Fractional Hamming distance d	$N_C/N = 1.5$	$N_C/N = 2$	$N_C/N = 4$
1/3	-0.053	-0.049	-0.040
3/10	-0.065	-0.059	-0.044
1/4	-0.084	-0.077	-0.048
1/8	-0.11	-0.12	-0.050
$0(P_{GS})$	-0.092	-0.15	-0.050

Appendix D. Warm starting the semi-greedy classical algorithm

To further approximate the critical Hamming radius for our PPSPs we applied a classical semi-greedy descent to random problem instances with constraint densities $N_C/N = \{1.5, 2, 4\}$. The algorithm is a modified version of the simple greedy algorithm

introduced in [37] and later applied to hypergraphs in [33]. To approximate the critical Hamming radius, we ran the semi-greedy algorithm on random hypergraphs for the stated constraint densities. The initial states started from distances rN flips away from the planted solution to show the SGC algorithm performance when starting from approximate solutions nearby the planted solution.

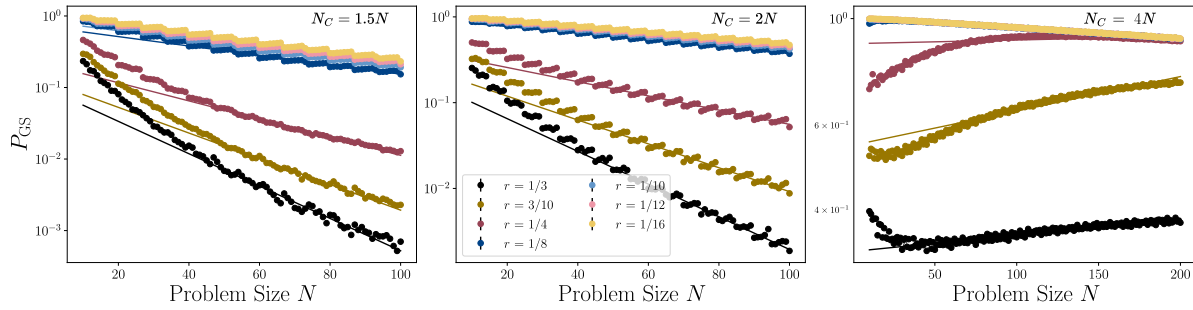


Figure D1. Probabilities of finding the ground state of random hypergraphs for $N_C/N = \{1.5, 2, 4\}$ using 100,000 trials on each constraint density and each value of rN . The data was fit to $a2^{bN}$ using problem sizes from $N = 40$ to $N = 100$. The results for the $4N$ constraint density demonstrated unusual behaviors at smaller N , so we ran it to $N = 200$ to try and find convergent behavior. The fits were calculated on the data points from $N = 100$ to $N = 200$. We found that the seeding the algorithm from farther away resulted in an increasing probability in finding the ground state with N , however it is clear that these trends eventually become an exponential decay in finding the ground state. With all of that considered, the algorithm still performs best when seeded with values close to the ground state.

Table D1. In (a), we report the exponential fits to the warm started SGC algorithm for various guessing errors $r \in [1/16, 1/3]$ shown in figure D1. In (b), we report the inferred performance of the algorithm set-up: TAQC \rightarrow SGC, where TAQC is used to produce a string to warm start the SGC algorithm. The exponents in (b) are calculated by adding the exponent from TAQC in Appendix B to the exponents in (a), at fractions $d = r$, where d is the approximation distance of TAQC, and r is the guessing error for the warm start of the SGC algorithm. We highlight the best combination of exponents in (b) using bold text, which are reported in table 1 in the main text.

(a)			
Guessing error r	$N_C/N = 1.5$	$N_C/N = 2$	$N_C/N = 4$
1/3	-0.075	-0.063	0.00093
3/10	-0.060	-0.046	0.0017
1/4	-0.042	-0.028	-0.00040
1/8	-0.022	-0.013	-0.00077
1/10	-0.021	-0.012	-0.00076
1/12	-0.021	-0.012	-0.00076
1/16	-0.021	-0.011	-0.00075
(b)			
Guessing error r	$N_C/N = 1.5$	$N_C/N = 2$	$N_C/N = 4$
1/3	-0.080	-0.10	-0.026
3/10	-0.073	-0.09	-0.025
1/4	-0.068	-0.081	-0.030
1/8	-0.078	-0.088	-0.038

Appendix E. Performance of IST-SAT with low frequency AC drives

In this section, we report the results obtained for IST-SAT with lower frequency of $\omega = 10\pi \ln N$ and shorter run-time that was simulated to larger system size ($N = 30$). While this parameter set underperforms higher frequency drives and with longer run-times as presented in the main text, the empirical radii of convergence we identify do not change significantly, except at $N_C/N = 1.5$, where the difference in performance is most significant. One may consider a finer-grid search to identify how exactly r_c scales with the constraint density N_C/N and fraction ϵ of unsatisfied constraints in the ground state.

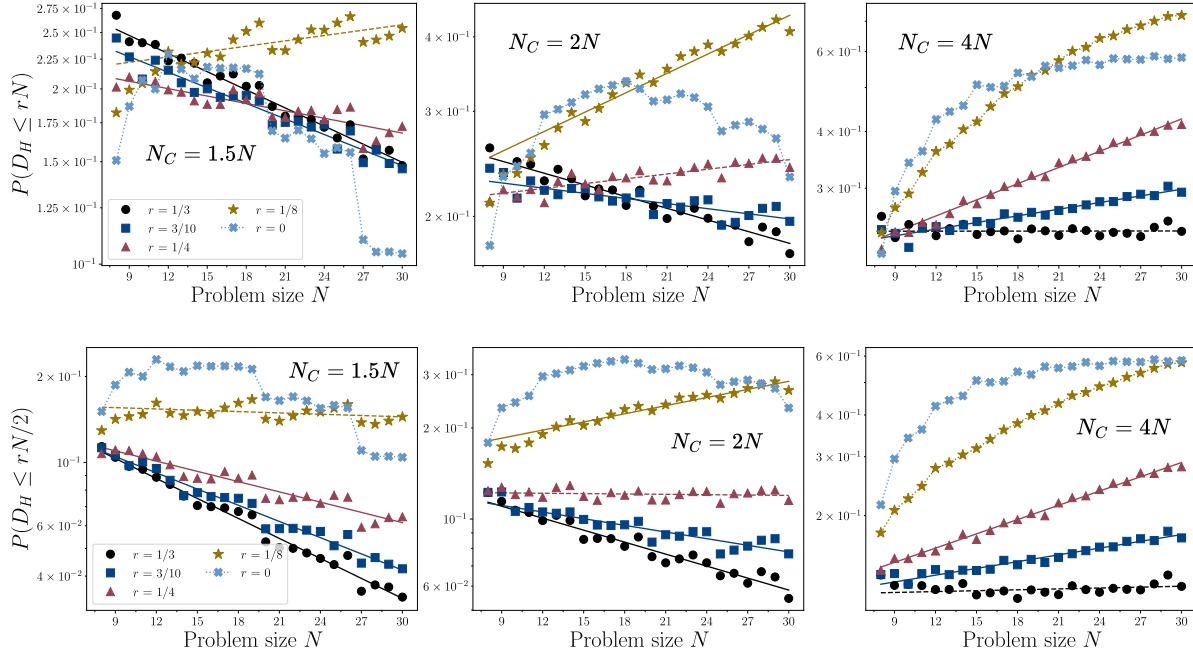


Figure E1. Convergence of IST-SAT to the planted solution. Probabilities of returning a solution with probabilities $P(D_H \leq rN)$ and $P(D_H \leq rN/2)$ are presented in the top and bottom row respectively. A critical Hamming radius r_c can be readily identified in these plots by dashed lines, which are approximately the first fraction r to obtain constant probability with increasing N .

Shown in figure E1, we observe monotonically increasing probabilities of returning solutions that approximate the PPS at different approximation ratios rN in terms of Hamming distance D_H . We identify a critical Hamming radius r_c for each constraint density $N_C/N = \{1.5, 2, 4\}$ respectively as $r_c = \{1/8, 1/4, 1/3\}$.

Table E1. Exponential fits to the low frequency IST-SAT data presented in figure E1. We report the exponents b from fitting data to $a2^{bN}$ where (a) is $P(D_H \leq rN)$ data, and (b) is the fit to $P(D_H \leq rN/2)$ data. The exponents lying at the approximate radii of convergence $r_c = \{1/8, 1/4, 1/3\}$ for $N_c/N = \{1.5, 2, 4\}$ (respectively) are highlighted in bold text. The -- entries represent data where exponential fits could not accurately represent the data.

(a)

Guessing error r	$N_c/N = 1.5$	$N_c/N = 2$	$N_c/N = 4$
1/3	0.034	0.022	-0.000083
3/10	0.030	0.00095	-0.016
1/4	0.014	-0.0089	-0.039
1/8	-0.010	-0.036	–

(b)

Guessing error r	$N_c/N = 1.5$	$N_c/N = 2$	$N_c/N = 4$
1/3	0.078	0.044	-0.030
3/10	0.063	0.024	-0.022
1/4	0.040	0.0013	-0.047
1/8	0.0051	-0.030	–

References

- [1] AB Finnila et al. Quantum annealing: a new method for minimizing multidimensional functions. *Chemical physics letters*, 219(5-6):343–348, 1994.
- [2] Tadashi Kadowaki and Hidetoshi Nishimori. Quantum annealing in the transverse ising model. *Physical Review E*, 58(5):5355, 1998.
- [3] Arnab Das and Bikas K Chakrabarti. Colloquium: Quantum annealing and analog quantum computation. *Reviews of Modern Physics*, 80(3):1061, 2008.
- [4] Mark W Johnson et al. Quantum annealing with manufactured spins. *Nature*, 473(7346):194, 2011.
- [5] Sergio Boixo et al. Evidence for quantum annealing with more than one hundred qubits. *Nature Physics*, 10(3):218, 2014.
- [6] Philipp Hauke et al. Perspectives of quantum annealing: Methods and implementations. *arXiv preprint arXiv:1903.06559*, 2019.
- [7] Andrew D King et al. Quantum critical dynamics in a 5,000-qubit programmable spin glass. *Nature*, pages 1–6, 2023.
- [8] E. Farhi et al. Quantum computation by adiabatic evolution. *arXiv:quant-ph/0001106*, 2000.
- [9] Tameem Albash and Daniel A. Lidar. Adiabatic quantum computation. *Rev. Mod. Phys.*, 90:015002, 2018.
- [10] Edward Farhi et al. A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028*, 2014.
- [11] Linghua Zhu et al. Adaptive quantum approximate optimization algorithm for solving combinatorial problems on a quantum computer. *Physical Review Research*, 4(3):033029, 2022.
- [12] Sergey Bravyi et al. Hybrid quantum-classical algorithms for approximate graph coloring. *Quantum*, 6:678, 2022.
- [13] Ruslan Shaydulin et al. Evidence of scaling advantage for the quantum approximate optimization algorithm on a classically intractable problem. *Science Advances*, 10(22):eadm6761, 2024.
- [14] CL Baldwin and CR Laumann. Quantum algorithm for energy matching in hard optimization problems. *Physical Review B*, 97(22):224201, 2018.
- [15] Vadim N. Smelyanskiy et al. Nonergodic delocalized states for efficient population transfer within a narrow band of the energy landscape. *Phys. Rev. X*, 10:011017, 2020.
- [16] Kostyantyn Kechedzhi et al. Efficient population transfer via non-ergodic extended states in quantum spin glass. *arXiv preprint arXiv:1807.04792*, 2018.
- [17] Vadim N Smelyanskiy et al. Intermittency of dynamical phases in a quantum spin glass. *arXiv preprint arXiv:1907.01609*, 2019.
- [18] J. Roland and N. J. Cerf. Quantum search by local adiabatic evolution. *Phys. Rev. A* **65**, 042308, 2002.
- [19] Eliot Kapit and Vadim Oganessian. Noise-tolerant quantum speedups in quantum annealing without fine tuning. *Quantum Science and Technology*, 6(2):025013, 2021.
- [20] Sergey Bravyi et al. The complexity of stoquastic local hamiltonian problems. *arXiv preprint quant-ph/0606140*, 2006.
- [21] Sergei V Isakov et al. Understanding quantum tunneling through quantum monte carlo simulations. *Physical review letters*, 117(18):180402, 2016.
- [22] Evgeny Andriyash and Mohammad H Amin. Can quantum monte carlo simulate quantum annealing? *arXiv preprint arXiv:1703.09277*, 2017.
- [23] Zhang Jiang et al. Scaling analysis and instantons for thermally assisted tunneling and quantum monte carlo simulations. *Physical Review A*, 95(1):012322, 2017.
- [24] Zhang Jiang et al. Path-integral quantum monte carlo simulation with open-boundary conditions. *Physical Review A*, 96(4):042330, 2017.
- [25] Andrew D. King et al. Scaling advantage over path-integral Monte Carlo in quantum simulation of geometrically frustrated magnets. *Nature Communications*, 12(1):1113, 2021.

- [26] Andrew Arrasmith et al. Variational quantum algorithms. *Nature Reviews Physics*, 3(9):625–644, 2021.
- [27] Michael Ragone et al. A unified theory of barren plateaus for deep parametrized quantum circuits. *arXiv preprint arXiv:2309.09342*, 2023.
- [28] Samson Wang et al. Noise-induced barren plateaus in variational quantum algorithms. *Nature Communications*, 12(1), 2021.
- [29] Ryan Babbush et al. Focus beyond quadratic speedups for error-corrected quantum advantage. *PRX Quantum*, 2(1):010103, 2021.
- [30] Sanjeev Arora and Boaz Barak. *Computational complexity: a modern approach*. Cambridge University Press, 2009.
- [31] Gianni Mossi et al. Embedding quantum optimization problems using ac driven quantum ferromagnets. *arXiv preprint arXiv:2306.10632*, 2023.
- [32] George Grattan et al. Exponential acceleration of macroscopic quantum tunneling in a floquet ising model. *arXiv preprint arXiv:2311.17814*, 2024.
- [33] Eliot Kapit et al. On the approximability of random-hypergraph max-3-xorsat problems with quantum algorithms. *arXiv preprint arXiv:2312.06104*, 2024.
- [34] Naeimeh Mohseni et al. Ising machines as hardware solvers of combinatorial optimization problems. *Nature Reviews Physics*, 4(6):363–379, 2022.
- [35] Johan Hästad. Some optimal inapproximability results. *Journal of the ACM (JACM)*, 48(4):798–859, 2001.
- [36] Olivier Dubois and Jacques Mandler. The 3-xorsat threshold. *Comptes Rendus Mathématique*, 335(11):963–966, 2002.
- [37] Matteo Bellitti et al. Entropic barriers as a reason for hardness in both classical and quantum algorithms. *Phys. Rev. Res.*, 3:043015, 2021.
- [38] Matthew Kowalsky et al. 3-regular three-xorsat planted solutions benchmark of classical and quantum heuristic optimizers. *Quantum Science and Technology*, 7(2):025008, 2022.
- [39] Dmitry A Abanin et al. Exponentially slow heating in periodically driven many-body systems. *Physical review letters*, 115(25):256803, 2015.
- [40] Tomotaka Kuwahara et al. Floquet–magnus theory and generic transient dynamics in periodically driven many-body quantum systems. *Annals of Physics*, 367:96–124, 2016.
- [41] Yasunari Suzuki et al. Qulacs: a fast and versatile quantum circuit simulator for research purpose. *Quantum*, 5:559, 2021.
- [42] Madelyn Cain et al. The qaoa gets stuck starting from a good classical string. *arXiv preprint arXiv:2207.05089*, 2023.
- [43] Eliot Kapit. Systems and methods for passive quantum error correction, March 23 2021. US Patent 10,956,267.
- [44] Zhijie Tang and Eliot Kapit. Unconventional quantum annealing methods for difficult trial problems. *Physical Review A*, 103(3):032612, 2021.
- [45] Robbie King. An Improved Approximation Algorithm for Quantum Max-Cut on Triangle-Free Graphs. *Quantum*, 7:1180, 2023.