


Unconventional quantum annealing methods for difficult trial problems

Zhijie Tang* and Eliot Kapit†

Department of Physics, Colorado School of Mines, 1500 Illinois Street, Golden, Colorado 80401, USA

 (Received 24 November 2020; accepted 16 February 2021; published 22 March 2021)

We consider a range of unconventional modifications to quantum annealing (QA), applied to an artificial trial problem with continuously tunable difficulty. In this problem, inspired by “transverse field chaos” in larger systems, classical and quantum methods are steered toward a false local minimum. To go from this local minimum to the global minimum, all N spins must flip, making this problem exponentially difficult to solve. We numerically study this problem by using a variety of new methods from the literature: inhomogeneous driving, adding transverse couplers, and other types of coherent oscillations in the transverse field terms (collectively known as RFQA, which stands for Random Field Quantum Annealing or Radio Frequency Quantum Annealing). We show that all of these methods improve the scaling of the time to solution (relative to the standard uniform sweep evolution) in at least some regimes. Comparison of these methods could help identify promising paths towards a demonstrable quantum speedup over classical algorithms in solving some realistic problems with near-term quantum annealing hardware.

DOI: [10.1103/PhysRevA.103.032612](https://doi.org/10.1103/PhysRevA.103.032612)

I. INTRODUCTION

Quantum annealing (QA) [1–6] is a promising method to solve optimization problems with noisy quantum hardware, with applications in machine learning, artificial intelligence [7–12], and many other topics. The time-dependent Hamiltonian of QA is engineered to encode the solutions of classical optimization problems in its ground state. By initializing the system in the ground state of a trivial driver Hamiltonian and evolving the system sufficiently slowly, QA can find the ground state of the target (classical) problem Hamiltonian. However, it is notoriously difficult to predict the performance of QA for realistic problems. Conclusive proof of a quantum speedup over classical methods for real problems remains elusive, with the possible exception of a frustrated magnet systems [13], where an empirical scaling advantage over classical path-integral quantum Monte Carlo (QMC) algorithms was shown. To help address this challenge, in this paper we theoretically survey a range of promising extensions to QA applied to a difficult trial problem, and identify a number of potential routes to a quantum speedup.

In the setup for QA, the total Hamiltonian in the standard uniform sweep evolution is a combination of a driving Hamiltonian H_0 and problem Hamiltonian H_p ,

$$H(t) = (1 - s(t))H_0 + s(t)H_p, \quad (1)$$

with the ground state of H_0 being easy to prepare. H_0 and H_p do not commute and the time-dependent annealing parameter $s(t)$ controls the time evolution of the system. The standard uniform sweep starts from $s(0) = 0$ and ends with $s(t_f) = 1$. Different functional choices for $s(t)$ may vary the efficiency of finding the ground state, such as a “reverse annealing schedule” [14]. It is intuitive to see that slowing down the annealing

process in the vicinity of minimum gap can help increase the success probability, and this tuning is required to recover the quantum speedup in the QA formulation of Grover’s search problem [15]. However, the instantaneous minimum gap value and location are not knowable in most realistic problems, and such fine tuning is often frustrated by noise, so we will study the simplest form, a linear schedule, $s(t): s(t) = \frac{t}{t_f}$ throughout this paper.

In QA, the ground state of a problem Hamiltonian, H_p , encodes the optimization problem solution. Experimentally realistic formulations of quantum annealing are typically arranged to solve quadratic unconstrained binary optimization (QUBO) problems, where the problem Hamiltonian is given by the Ising model,

$$H_{\text{ising}} = \sum_{\langle i,j \rangle} J_{ij} \sigma_i^z \sigma_j^z + \sum_{i=1}^N h_i \sigma_i^z, \quad (2)$$

the ground state of which can be encoded as the solution space of some NP-hard problems [16], and given enough additional qubits, any NP-complete problem can be expressed in this form. To find the ground state of the problem Hamiltonian, we first prepare the system in the ground state of H_0 , which is chosen as a uniform transverse field Hamiltonian,

$$H_0 = - \sum_{i=1}^N \sigma_i^x. \quad (3)$$

The initial ground state is a uniform superposition state in the computational basis. The quantum adiabatic theorem states that as long as the annealing evolution is slow enough, the system remains in the instantaneous eigenstate of the time-dependent Hamiltonian at all times. This theorem also provides a widely used criterion that, with the linear annealing schedule, the total adiabatic evolution time t_f to find the ground state with high probability has an inverse

*ztang@mines.edu

†ekapit@mines.edu

minimum-gap-squared dependence,

$$t_f \propto \frac{W}{\Delta_{\min}^2}. \quad (4)$$

Here, Δ_{\min} is the minimum energy gap, and W scales as the total energy change of the final ground state $|0\rangle$ over the entire evolution: $W \sim \langle 0|H_0|0\rangle - \langle 0|H_p|0\rangle$. Note that this result is a *worst-case* scaling estimate of the time to solution, and a variety of diabatic effects can substantially increase performance. We will encounter a number of examples of this later in this work.

In cases where the system undergoes a first-order transition [17], Δ_{\min} typically decreases exponentially with the system size N , and the corresponding evolution time t_f (and thus, time to solution) grows exponentially. For hard optimization problems which suffer from such phase transitions, many new schemes have been proposed to accelerate QA, such as the use of nonstoquastic Hamiltonians [18–23], inhomogeneous driving of the transverse field [24,25], and oscillatory transverse fields (RFQA, which stands for Random Field Quantum Annealing or Radio Frequency Quantum Annealing) [26]. We study a range of examples drawn from these works.

To investigate a number of new methods from the literature, we make our own artificially difficult problem Hamiltonian, partially inspired by previous studies of “spike problems” [18,27], rather than studying QUBO problems directly as was done in [20]. In this problem, which we call the asymmetric magnetization problem (AMP), local searches and QA steer the system toward a false minimum. This “wrong-way steering” makes finding the true ground state exponentially difficult. The difficulty exponent associated with the AMP can be continuously tuned for further investigation of the performance of these alternative QA methods with problem hardness. Further, the energy landscape depends only on total magnetization m , making it easier to study analytically. We do not consider noise in this paper, as none of the methods we consider require fine tuning. We expect these methods to be resilient to noise described by the empirical noise model for superconducting flux qubits. There is strong theoretical evidence for this resiliency in the case of RFQA [26], and both theoretical and experimental evidence for inhomogeneous driving [24,25,28,29].

The rest of the paper is organized as follows. In Sec. II, we introduce our trial problem; in Sec. III, we provide an analytical prediction of its minimum gap. Section IV discusses the performance of the standard uniform sweep applied to this problem, against which we benchmark all other methods. We then introduce inhomogeneous driving, transverse couplers, and RFQA, comparing their behavior with the standard uniform sweep routine in Secs. V and VI separately. The final section summarizes our results and provides comments on the performance of these methods.

II. ASYMMETRIC MAGNETIZATION PROBLEM

While conventional Ising Hamiltonians can encode nearly any combinatorial optimization problem, we choose an artificial toy problem model to better study a variety of new methods to find its solution. This is in part because extracting exponential difficulty scaling is notoriously difficult and unre-

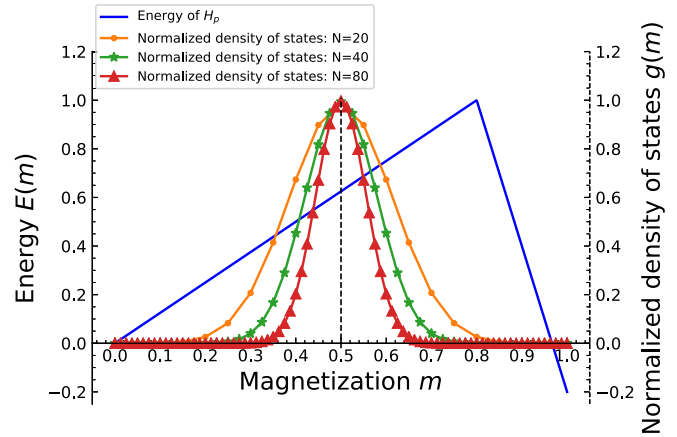


FIG. 1. Density-of-states distribution and energy of the problem Hamiltonian in the AMP model. The x axis is the total magnetization m , and the y axis represents both the energy spectrum of the problem Hamiltonian (left) and the density-of-states distribution (right). The lines with markers are the density of states for system sizes of $N = 10, 11, 12$. The distributions become narrower with increasing system size. The solid blue line is the energy landscape of H_p . In our model, the distribution follows a Gaussian distribution, centered at $m = 0.5$. Comparing the distribution of density of states with the energy spectrum of the problem Hamiltonian, we see the peaks of the density of states are distributed behind the global maximum. The asymmetry of the H_p results in an exponentially difficult problem.

liable in random structured problems, at least when N is small enough for exact classical simulation. Our artificial problem, the “AMP,” has the problem Hamiltonian defined as function of total magnetization m ,

$$H_p = f(m); \quad m = \frac{1}{2N} \sum_{i=1}^N (1 + \sigma_i^z), \quad (5)$$

where N is the system size, σ_i^z is the Pauli matrix with discrete eigenvalues ± 1 , and $f(m)$ is designed to have two competing minima at $m = 0$ and $m = 1$ with all spins down and all spins up. The form of $f(m)$ is controlled by two free parameters, A and x_p :

$$f(x) = \begin{cases} \frac{x}{x_p} & \text{if } x < x_p \\ 1 - (1 + A) \frac{x - x_p}{1 - x_p} & \text{if } x \geq x_p. \end{cases} \quad (6)$$

Here, $f(1)$ is the true ground state, and $f(0)$ is the false minimum. x_p defines the location of the global maximum, and A defines the energy difference of the two competing states: $f(0) - f(1) = A$. By adjusting these two parameters, we can continuously tune the difficulty of the problem.

The difficulty of finding the global minimum in our model is strongly related to the distribution of the density of states. The density of states follows a Gaussian distribution as shown in Fig. 1, with the most probable initial state centered at $m = 0.5$ where half of the spins are flipped from the true ground state, and a global maximum is distributed at $m = x_p > 0.5$. We see the system has a large tendency to get stuck in the local minimum, since the possible initial state is settled behind the global maximum. The wrong-way guidance is the generic failure mechanism for classical and quantum

TABLE I. Summary of scaling exponents for each method. Fitting the time to solution $T_{its}(N)$ of each method to $2^{\beta+\gamma N}$, the table lists the exponential scaling coefficient “ γ ” value for each method. “S” represents the standard uniform sweep method, “I” represents the inhomogeneous driving method, “C_{F,A,M}” represents the transverse couplers method, in which “F,” “A,” and “M” correspond to ferromagnetic, antiferromagnetic, and mixed couplers, respectively. “M” represents the RFQA-M method, “CM” the RFQA-M with couplers method, “SyncM” the synchronized RFQA-M method, “SyncMC” the synchronized RFQA-M with couplers method, and “D” the RQFA-D method.

Problem set	$1/\Delta_{\min}^2$	S	I	C _{F,A,M}	M	CM	SyncM	SyncMC	D
$A = 0.2, x_p = 0.8$	2.25	2.12	0.79	1.77, 2.09, 2.50	1.48	1.31	1.28	0.86	1.56
$A = 0.28, x_p = 0.7$	1.48	1.39	0.70	1.25, 1.23, 1.67	0.89	0.86	0.84	0.64	1.05
$A = 0.3, x_p = 0.64$	1.04	1.06	0.69	0.87, 0.77, 1.12	0.62	0.62	0.64	0.62	0.50
$A = 0.34, x_p = 0.59$	0.61	0.52	0.69	0.48, 0.44, 0.54	0.45	0.45	0.45	0.51	0.40

optimization algorithms [30], as if local guidance from a random initial state tends to point toward the true solution of a problem it can be solved trivially. But if local guidance points toward false minima, then the problem can quickly become hard, and in more realistic problems at large N there are often exponentially many local minima. Multiqubit tunneling between well-separated minima—exactly the process we simulate here—has been identified as a critical bottleneck in many realistic problems [31].

In the AMP problem, we set the global maximum at $m = x_p$ where $x_p > 0.5$, with two competing minima at each end. From the density-of-states distribution in Fig. 1, we can tell the system has a large tendency to be steered toward the false minimum. In classical algorithms such as simulated annealing, the system will easily get stuck in the false minimum since the possible initial state is mostly distributed around $m = 0.5$. The possibility of finding the global minimum is large if the initial instantaneous state of H_p happens to be guessed beyond the global maximum at a position that $m > x_p$, but if the initial state is located at any $m < x_p$, the possibility of climbing the hill is exponentially small. Cost functions similar to the AMP model have been studied in [18,27], and classical simulated annealing was shown to be inefficient for solving such problems. We will show that the AMP problem is also exponentially difficult to solve with quantum annealing, and we focus on how various modifications to QA compare with a homogeneous transverse field and uniform sweep (the “default” quantum annealing method) in solving the AMP problem.

When applying QA to the AMP, performance is bottlenecked by an exponentially small gap at a first-order transition [17,32,33]. As shown in Fig. 2, the magnetization is entropically steered toward zero as the system evolves, and all N spins must simultaneously flip to reach the true ground state. The difficulty scaling of the problem model can be tuned by A and x_p ; smaller A corresponds to a smaller energy gap between the two competing minima, which intuitively increases the difficulty level of the problem. Similarly, larger x_p moves the peak further away from the center of the density of states, and the system then has a larger tendency to be steered to the false local minimum, which also increases problem difficulty. We make an ensemble of problem models with different A and x_p so that we can investigate the relationship between the performance of different methods with the difficulty of the problem models. We make modifications to the traditional QA method and evaluate their performance by numerically calculating the time to solution, and show

that the AMP is exponentially difficult to solve with quantum algorithms, but modifications to the traditional QA method can lead to substantial improvements in the scaling of the time to solution. The exponential scaling coefficients for each method are listed in Table I.

Although this toy model is just a simplified artificial problem without a realistic implementation, as described above, it captures the basic bottleneck of most classical and quantum optimization problems. So any method which accelerates finding a solution in the AMP is likely broadly applicable to more realistic cases.

III. STANDARD UNIFORM SWEEP ROUTINE

We first investigate the performance of the standard uniform sweep method, for system sizes N ranging from 5 to 18 spins. In this method, the driving Hamiltonian is

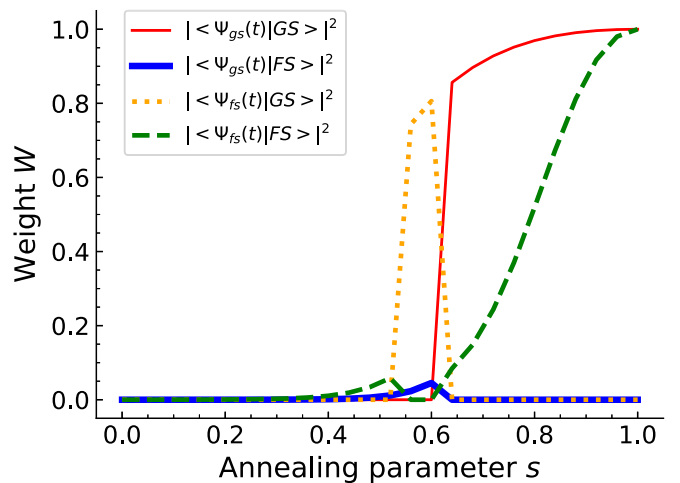


FIG. 2. The instantaneous overlap of the ground and first excited states with the true and false ground states of the classical problem as a function of the annealing parameter s , with system size $N = 18$, and the problem Hamiltonian is defined with parameters $A = 0.2$ and $x_p = 0.8$. The x axis is the annealing parameter s , and the y axis denotes the probability of getting a specific state. The red (thin) and blue (thick) solid lines are the overlap of the instantaneous ground state with the true and false ground states of H_p , the orange (dotted) and green (dashed) lines are the overlap of the instantaneous first excited state with the true and false ground states of H_p . The comparison shows that the system is steered toward the false minimum first and all N spins have to flip to reach the true ground state.

a homogeneous transverse field: $H_0 = -\sum_{i=1}^N \sigma_i^x$. The total Hamiltonian is a combination of the drive Hamiltonian and problem Hamiltonian H_p :

$$H(s) = -(1-s) \frac{1}{N} \sum_{i=1}^N \sigma_i^x + sH_p. \quad (7)$$

The initial ground state of the system is the ground state of H_0 , which is a uniform superposition of states corresponding to all possible assignments of bit values with equal weights. As the system evolves, the Hamiltonian linearly interpolates between the transverse field Hamiltonian and the problem Hamiltonian, i.e., starting as H_0 and ending in H_p . As long as the system stays in the instantaneous ground state, the system will be steered toward the false minimum first, but at some critical s_c the true and false ground states cross and all N spins must flip, as shown in Fig. 2 with red and blue solid lines. For this problem, there is only one avoided crossing in the standard uniform sweep method (though we find multiple crossings when inhomogeneous driving is employed). Since the gap at the phase transition point is exponentially small in N , unless the evolution is performed extremely slowly, the avoided crossing will be diabatically missed and the probability of finding the true ground state will be suppressed, while the probability of finding the false ground state dominates, as shown in Fig. 2 with orange and green dashed lines.

We evaluate the performance of the standard uniform sweep algorithm by computing the time to solution (T_{ts}) over a range of system sizes. The T_{ts} measures the time needed to find the ground state with 99% success probability [30]:

$$T_{ts} \propto t_f \frac{\ln(1-0.99)}{\ln[1-p(t_f)]}, \quad (8)$$

where $p(t_f)$ is the success probability in a single-trial with runtime t_f ; for small $p(t_f)$, $T_{ts} \propto t_f/p(t_f)$. A short t_f taking advantage of diabaticity might not show how different methods change the scaling of multiqubit tunneling at large N , so we assume polynomially increasing t_f with N in our work, as one would likely want to use in a real annealer at much larger N , and the t_f is empirically chosen as $20(\frac{N}{7})^2$ for the AMP model. As mentioned in the Introduction, the evolution time needed to find the ground state increases as Δ_{\min}^{-2} . To explore the performance of the standard uniform sweep method under different problem models, we choose four sets of parameters, $\{A = 0.2, x_p = 0.8\}$, $\{A = 0.28, x_p = 0.7\}$, $\{A = 0.3, x_p = 0.64\}$, and $\{A = 0.34, x_p = 0.59\}$, forming an ensemble of problem models with descending difficulty level. These parameters are chosen to approximately set $\Delta_{\min} \propto \{2^{-N}, 2^{-3N/4}, 2^{-N/2}, 2^{-N/4}\}$, respectively. As mentioned previously, increasing A or decreasing x_p toward $1/2$ both decrease the difficulty exponent, and moving either parameter in the opposite direction makes the problem harder.

As expected by an exponentially closing gap, the time needed to find the solution exponentially increases with system size. This is confirmed in Fig. 3, where the corresponding time to solution exponentially increases with the system size in all problem sets, and the difficulties of the four sets are well separated from each other. With the minimum gap and performance of the standard uniform sweep method rigorously understood, we now apply other methods from the literature




FIG. 3. Time to solution to find the true ground state in four problem model sets using the standard uniform sweep method, computed from the final success probability for a runtime polynomially increasing with N . The difficulty level of the four models is arranged in descending order: $\{A = 0.2, x_p = 0.8\}$, $\{A = 0.28, x_p = 0.7\}$, $\{A = 0.3, x_p = 0.64\}$, and $\{A = 0.34, x_p = 0.59\}$. Different markers are data with the standard uniform sweep method of four problem model sets, solid lines are best-fit curves of the numerical data, and dashed lines are the inverse of the square of the numerically estimated minimum gap from Fig. 9. The time to solution closely tracks the average minimum gap squared as expected. In all cases, the N -spin tunneling bottleneck in this problem leads to an exponentially increasing time to solution.

to compare their performance with it and investigate their abilities of providing a quantum speedup.

IV. SUMMARY OF RESULTS

Before we proceed to the detailed investigation of alternative QA methods, we compile a summary in Table I that lists the exponential fitting results of T_{ts} for each method. We fit $T_{ts}(N)$ to $2^{\beta+\gamma N}$ and extract the γ value to determine the difficulty scaling for each method. We find that in the harder problem sets where $x_p = 0.8, A = 0.2$ and $x_p = 0.7, A = 0.28$, synchronized RFQA-M (with transverse couplers) and inhomogeneous driving show the best performance, but in the relatively easier problem sets where $x_p = 0.64, A = 0.3$ and $x_p = 0.59, A = 0.34$, the RFQA-D method shows the best scaling advantage. The details are illustrated and discussed in the following sections; we include this table as a central reference point for the results of all of our studies.

V. MODIFIED ADIABATIC ANNEALING STRATEGIES: INHOMOGENEOUS DRIVING AND TRANSVERSE COUPLERS

A wide range of modifications to quantum annealing have shown significant promise in theoretical studies [19,20,23,25,34–38]. In this section, we begin applying methods from the literature to our AMP model and assess their performance by computing the T_{ts} as in Eq. (8). We

begin by considering the inhomogeneous driving method and transverse couplers. Inhomogeneous driving and the ferromagnetic transverse couplers belong to the class of stoquastic Hamiltonians, while the antiferromagnetic couplers and mixed-sign couplers have nonstoquastic Hamiltonians. A stoquastic Hamiltonian has real and nonpositive off-diagonal matrix elements in the computational basis [39], and can often (but not always [40]) be efficiently simulated by sign-problem-free QMC. Nonstoquastic Hamiltonians, on the other hand, suffer from a sign problem and thus cannot be efficiently simulated in QMC in general, though some particular nonstoquastic Hamiltonians can be simulated in QMC by clever schemes to avoid the sign problem [41]. Amenability (or not) to QMC is a critical issue in QA, as in recent studies, QMC displays comparable exponential scaling to the physical incoherent tunneling rate in quantum annealers [40,42–45]. It is thus intuitive to infer that the efficiencies of QMC and quantum annealers are similar in solving many problems, making it difficult to realize a genuine quantum speedup. Nonstoquastic Hamiltonians do not suffer from this issue, and have demonstrated significant benefits in some theoretical work [19,21–23,46].

A. Inhomogeneous driving

In inhomogeneous driving, the transverse fields are ramped down at different rates from one site to the next, as first described in [24,25]. In the original proposal [24], the magnitude of the transverse field applied to the N spins is turned off sequentially with a set of time-dependent amplitudes $\Gamma_i(s)$. In that work, the inhomogeneous driving transverse field circumvents the first-order quantum phase transition and provides an exponential quantum speedup in a p -body interacting mean-field-type model. A more careful analysis [25] which included noise and disorder found the exponential speedup to be somewhat fragile, but showed that a consistent polynomial speedup persisted given these more realistic assumptions. Further, there is experimental evidence that inhomogeneous driving is effective in real hardware [29].

Inspired by the performance improvements offered by inhomogeneous driving of the transverse field Hamiltonian, we apply it to the four AMP problem sets as follows:

$$H(t) = -\frac{1}{N} \sum_{i=1}^N \Gamma_i(s) \sigma_i^x + s H_p,$$

$$\Gamma_i(s) = \begin{cases} 1 & \text{if } s < s_i \\ N(1-s') + (1-i) & \text{if } s_i \leq s \leq s_{i-1} \\ 0 & \text{if } s_{i-1} < s. \end{cases} \quad (9)$$

The scaling of $\Gamma_i(s)$ shown in Fig. 4 is what suggested in [24]. Interestingly, the measurement of T_{tts} in Fig. 5 shows that the inhomogeneous driving method has a difficulty scaling which is very weakly dependent on the control parameters A and x_p , with the T_{tts} scaling virtually identically in each case. It consequently outperforms the standard uniform sweep method for the harder problem regimes, but actually shows worse performance for the easiest parameter sets. While we cannot predict its performance analytically in this case (the perturbation theory we use to calculate Δ_{\min} is not well de-




FIG. 4. Transverse field strength $\Gamma_i(s)$ in the inhomogeneous driving method at each i . This scaling form is drawn from [24].

finer for some of the transverse fields set equal to strictly zero), a clue to the origin of this behavior is found in a numerical analysis of the level structure, as we now describe.

In Fig. 6 we show the energy difference of the higher-order excited states with the ground state in the hardest problem class with $N = 10$. In contrast to a uniform sweep, we find two avoided crossings in the annealing process, a generic feature of inhomogeneous driving in this system that we observed for other parameter sets as well (data not shown). The presence of two crossings is likely what is responsible for the performance boost observed in the harder problems, and why it seems to have the same scaling for different parameters. A similar phenomenon is observed in the glued trees problem [5], where constructive interference of diabatically missing two avoided




FIG. 5. Time to find the true ground state in four problem model sets using the inhomogeneous driving method, computed from the final success probability for a runtime polynomially increasing with N . Larger markers are data with the inhomogeneous driving method, smaller markers are data with the standard uniform sweep method, dashed lines are best-fit curves of the inhomogeneous driving method, and solid lines are best-fit curves of the standard uniform sweep method for comparison purposes. The inhomogeneous driving method helps in the harder cases but in the easier cases it is less efficient than the standard uniform sweep.




FIG. 6. Energy difference of the higher-order excited states with the ground state in the inhomogeneous driving method. This is the hardest problem where $\{A = 0.2, x_p = 0.8\}$, $N = 10$. We simulated the energy differences up to the 20th excited state.

crossings leads to an exponential speedup. However, unlike the glued trees problem, there is no clear separation between the two competing minima and the higher excited states in the AMP model. It is clear from Fig. 6 that there also exists an overlap region of the higher-order excited states with the first excited state. As A and x_p are varied to make the difficulty scaling decrease, the two avoided crossings move closer together, and the distance from higher levels also shrinks and becomes exponentially small. Consequently, this effect does not result in an exponential speedup here, and shows worse performance than a uniform sweep in the easiest cases.

B. Transverse couplers

Adding two-body transverse coupling to QA [19,21,23,46,47] is often considered to be a promising route to a quantum speedup. For instance, Hormozi *et al.* [23], constructed a stoquastic Hamiltonian by inserting the ferromagnetically coupled term H_I^F into the traditional Ising model, and a nonstoquastic Hamiltonian by inserting the antiferromagnetically coupled term H_I^A or mixed coupled term H_I^M as follows:

$$\begin{aligned}
 H_I^F &= - \sum_{(i,j)}^N \sigma_i^x \sigma_j^x, \\
 H_I^A &= + \sum_{(i,j)}^N \sigma_i^x \sigma_j^x, \\
 H_I^M &= \sum_{(i,j)}^N r_{ij} \sigma_i^x \sigma_j^x,
 \end{aligned}
 \tag{10}$$

where r_{ij} is randomly chosen from $\{-1, 1\}$ to include both ferromagnetic and antiferromagnetic cases. In that work, they found that both stoquastic and nonstoquastic Hamiltonians showed an advantage over a uniform transverse field for a class of long-range Ising spin glass problems, with the nonstoquastic methods generally showing better performance. This

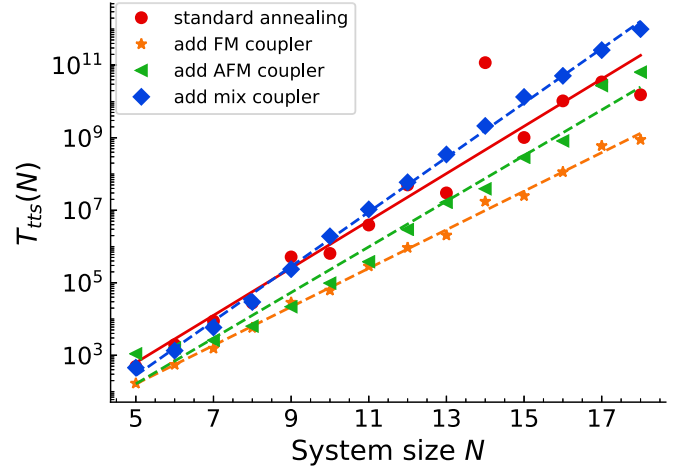


FIG. 7. Time to find the true ground state in the hardest problem set, $\{A = 0.2, x_p = 0.8\}$, using the transverse coupler methods, computed from the final success probability for a run time polynomially increasing with N . Data for adding a ferromagnetic coupler, an antiferromagnetic coupler, and mixed couplers are given by orange, green, and blue markers, respectively. Red dots are data of the standard uniform sweep method for comparison purposes. The solid red line is the best-fit curve of the standard uniform sweep method for comparison. Other dashed lines are best-fit curves for the transverse coupler methods. Adding ferromagnetic and antiferromagnetic couplers to the conventional standard uniform sweep routine shows obvious quantum speedup, although adding the mixed couplers reduces the advantage to some extent.

motivated us to investigate the same method in our AMP model. We add transverse couplers into our model and choose a path of the form [23,48]

$$H(s) = (1 - s) \frac{1}{N} H_0 + s(1 - s) \frac{1}{N} H_I + s H_p. \tag{11}$$

We apply the transverse coupler Hamiltonian to our four problem sets, plotting the results for the hardest scaling choice as an example in Fig. 7. It is straightforward to see that adding ferromagnetic or antiferromagnetic coupler has a clear scaling advantage over a standard uniform sweep, but mixed couplers actually lead to decreased performance. The quantum speedup from coupler terms is probably because the couplers can flip two spins simultaneously, so the tunneling process from one configuration to the other can occur at lower order than with a uniform transverse field (where it occurs at N th order in this model). The ferromagnetic coupler increases the minimum gap and thus provides a quantum speedup over the standard uniform sweep method; the same effect is observed in [23]. The antiferromagnetic couplers actually decreased the minimum gap but still show a scaling advantage, so the reason for the increased performance from the antiferromagnetic ones remains elusive. The behavior of the transverse coupler methods in the other three problem sets are shown in Table I.

VI. REVERSE ANNEALING AND COLD BATHS

Reverse annealing [11,49–51], where the system is initialized in a local classical minimum and the transverse field is ramped up and down to search for other minima, unfortunately

provides no benefit for the AMP. Reverse annealing was shown in [51] to provide benefit for the p -spin ferromagnet problem, if one is able to guess an initial state sufficiently close to the true ground state. However, in the AMP there are only two minima to choose from, separated by N spin flips. The only sensible choice (without dramatically modifying H_p) is thus to initialize the system in the false minimum. We simulated the reverse annealing protocol (data not shown) by initializing the system in the false minimum, ramping the transverse field up to a finite value guessed randomly from an $O(1)$ range enclosing the phase transition point, evolving from that point for $O(N^2)$ time, then ramping it back down to zero. With sufficient averaging over the location of the pause point (which is not knowable precisely in real problems), we found a time to solution which scaled nearly identically to the standard uniform sweep method for all parameters studied. Thus, we found no benefit in applying reverse annealing to this problem.

The influence of a cold bath on this system is more subtle. It is well known [36,52–58] that coupling a quantum spin glass to a cold bath can improve the process of finding its low-energy states. So let us consider coupling the AMP to a low-temperature bath during annealing. Importantly, we here assume that T is small compared to the single qubit excitation energy, but it may still be large compared to the (exponentially small) minimum gap. How much can such a bath improve the time to solution?

Unfortunately, numerical simulation of such a system is prohibitively expensive [59] given the complexity of the Lindblad operators used to represent the finite temperature bath. We can, however, estimate the relaxation rate from the bath by appealing to the matrix element scaling conjecture (MSCALE conjecture) [26]. This conjecture states that, for few-body operators, the scaling (with problem size N) of matrix elements of these operators between competing minima of quantum spin glasses near a phase transition is the same as the scaling of the minimum gap itself. This conjecture is true by inspection for the AMP, since the gap can be computed accurately using the modified forward approximation in the Appendix. If we assume that each spin couples to a cold bath independently, then the rate of mixing near the phase transition scales as $N\Gamma_B^2/W$, where $\Gamma_B \propto \Delta_{\min}$ is the matrix element from a local spin operator and W is the energy range swept over. This produces a factor-of- N enhancement relative to the closed system, but does not change the scaling exponent as the other methods do. The cold bath may, however, improve performance in a real system by relaxing few-body excitations back toward the ground state, correcting “errors” induced by other channels.

VII. RFQA

Stoquastic or not, the previous sections all explored “DC” schemes involving slow variations of transverse field and coupler terms. In this section, we consider an AC alternative, called RFQA [26]. In RFQA, the traditional transverse field driver Hamiltonian is modified by independently oscillating either the magnitude (RFQA-M) or direction (RFQA-D) of each transverse field term (M and D refer to magnitude and direction, respectively). As we will describe shortly, the

qualitative explanation for a quantum speedup in RFQA is an exponential proliferation of weak many-spin processes, leading to accelerated mixing near first-order quantum phase transitions. The total Hamiltonian in RFQA is given by

$$H(t) = (1-s)H_{M/D}(t) + sH_p, \quad (12)$$

where the driving fields in RFQA-M and RFQA-D are defined as follows:

$$\begin{aligned} H_M(t) &= -\kappa \sum_{i=1}^N [1 + \bar{\alpha}_i \sin(2\pi f_i t)] \sigma_i^x, \\ H_D(t) &= -\kappa \sum_{i=1}^N [\cos(\bar{\alpha}_i \sin(2\pi f_i t)) \sigma_i^x \\ &\quad + \sin(\bar{\alpha}_i \sin(2\pi f_i t)) \sigma_i^y]. \end{aligned} \quad (13)$$

Here, $\bar{\alpha}_i$ is the amplitude of each oscillation, the frequencies f_i of the field are randomly chosen between f_{\min} and f_{\max} , and κ is the magnitude of the transverse field. To avoid uncontrolled heating, both f_{\min} and f_{\max} have inverse polynomial scaling in N . To estimate the performance of RFQA, we average the success probability $p(t_f)$ over hundreds of random choices of the $\{f_i\}$ when computing time to solution. The RFQA methods all rely on finite frequency dynamics that are not captured by QMC, making them promising candidates for producing a quantum speedup. The two methods are straightforward to implement in flux qubit hardware, by applying oscillating magnetic fields as described in [26].

As described in the original work, the qualitative speedup mechanism from RFQA is complex and arises from an exponential proliferation of weak multiphoton transitions. As the system nears a phase transition point, whenever the energy of the two ground states crosses a combination of m oscillating frequencies there is an m th-order driving process that (very weakly) mixes the two states. In general, the Rabi frequency of such a process decreases exponentially in m , but there are $2^m \binom{N}{m}$ such terms and the combination of all of them dramatically accelerates the phase transition. If the m th-order resonance is smaller than the base tunneling rate $\Omega_0 = \Delta_{\min}/2$ by a factor Λ^m , then the total transition rate is expected to scale approximately as

$$\begin{aligned} \Gamma_T &= \frac{\sum_{i=1}^N |\Omega_i|^2}{W} \simeq \frac{\Omega_0^2}{W} \sum_{l=1}^N \Lambda^{2l} \binom{N}{l} 2^l \\ &\simeq \frac{\Omega_0^2}{W} (1 + 2\Lambda^2)^N. \end{aligned} \quad (14)$$

Predicting Λ is a subtle challenge and something we will leave for future work; we restrict our study of RFQA to purely numerical simulations here.

A. RFQA-M

In RFQA-M the magnitudes of the transverse field terms coherently oscillate with time as the global amplitude is ramped down toward zero, so that an individual transverse field term κ is replaced with $\kappa(1 + \bar{\alpha} \sin 2\pi f_i t)$. In our simulations we used $\bar{\alpha} = 0.9$, and magnitude of frequencies f_i is randomly chosen between $\{\frac{0.3}{N^{1.5}}, \frac{0.6}{N^{1.5}}\}$; the signs of the f_i are also randomly chosen. This is superficially similar to

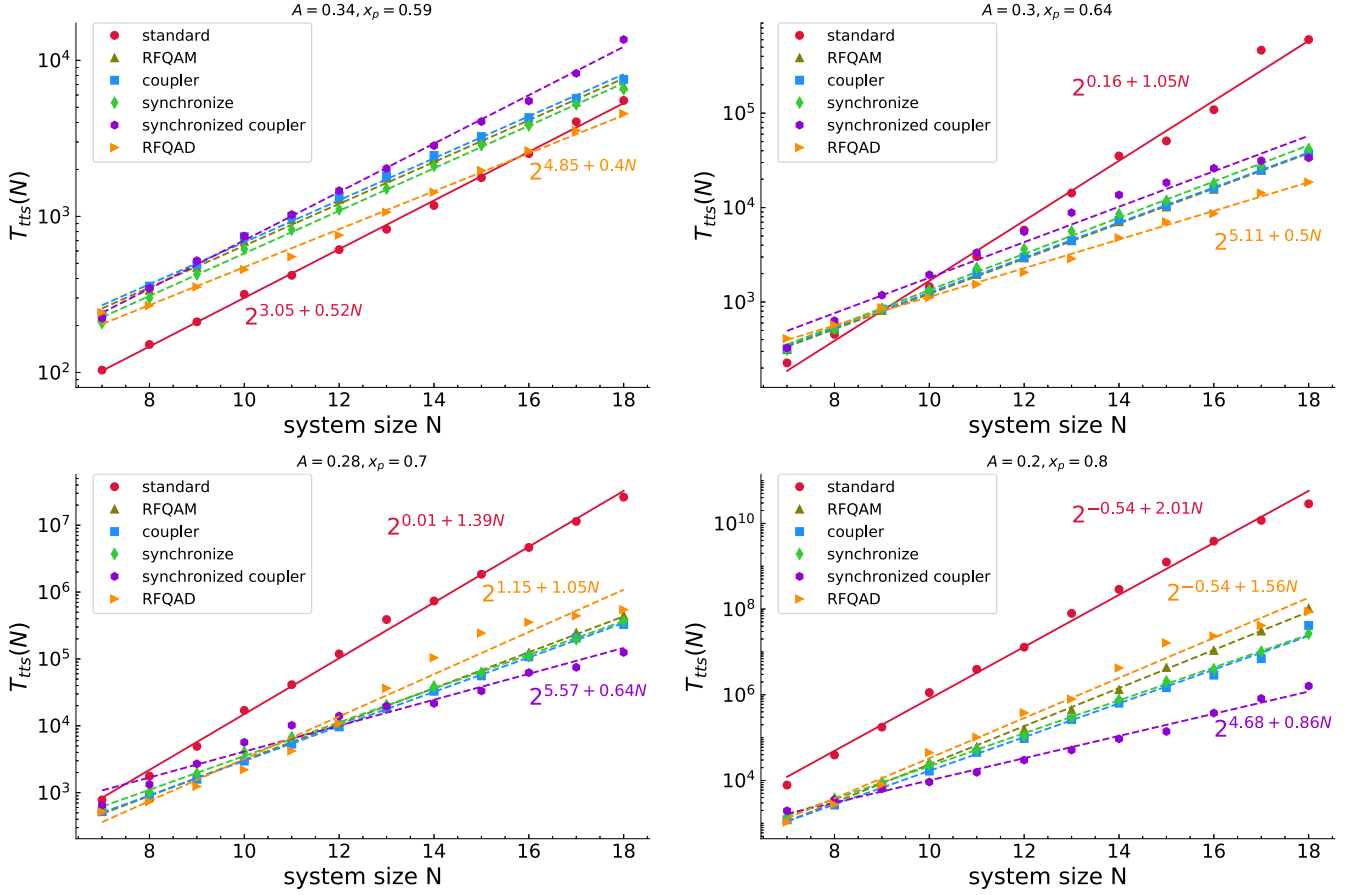


FIG. 8. Time to find the true ground state in four problem model sets using the RFQA method, computed from the final success probability for a run time polynomially increasing with N . Data of the RFQA methods are given by different markers. Red dots are data of the standard uniform sweep method, and the solid red line is the best-fit curve of the standard uniform sweep method for comparison purposes. Other dashed lines are best-fit curves of the RFQA methods.

inhomogeneous driving, but the coherent oscillations lead to nonmonotonic changes of κ_i with time and very different scaling as a result. We also considered a few additional variations of RFQA-M. In one set of simulations, we explored a partially synchronized RFQA-M method, in which N spins are broken into k groups; instead of generating different random frequencies for each site, the transverse fields in each group are all oscillated in phase with the same frequencies. In this work, we only divided the N spins into two groups, but other arrangements are possible. We also explored adding transverse couplers to the RFQA-M method, where all transverse couplers and fields are independently oscillated in magnitude. The total Hamiltonian in this method is defined as follows:

$$H(s) = (1-s)H_{M/D} + sH_p + s(1-s)\kappa_r \sum_{\langle i,j \rangle} \sin(2\pi r_{ijt})\sigma_i^x \sigma_j^x, \quad (15)$$

where κ_r is the magnitude of the coupler terms, and r_{ij} are the oscillating frequencies of the coupler, randomly chosen between $\{r_{\min}, r_{\max}\}$. The magnitude r is defined to polynomially decrease with N , and the frequencies r_{ij} are also inverse polynomial in N . Finally, we looked at partially synchronized RFQA-M with transverse couplers, where the transverse couplers are also synchronized into groups.

We compare the T_{tts} of the various implementations of RFQA-M with the standard uniform sweep method in Fig. 8. The results show that RFQA-M and its adaptations can provide quantum speedup over the standard uniform sweep routine, with a scaling advantage which is particularly obvious in harder problem sets.

B. RFQA-D

In RFQA-D the direction of each transverse field term oscillates in time, tipping back and forth in the x - y plane. This can be engineered through oscillating z biases (which can be shown to be equivalent to a tipping transverse field through a time-dependent unitary transformation) and has the elegant property that the instantaneous spectrum of the system is preserved in the evolution, so the oscillations have no “steering” effect whatsoever (unlike RFQA-M, where changing transverse field magnitudes can change the relative energies of competing minima, in addition to any AC effects). Any performance advantage from RFQA thus comes directly from the proliferation of weak transitions described above.

As shown in Fig. 8, we see that RFQA-D does provide an obvious quantum speedup over a uniform sweep. In all studied cases, RFQA-D reduced the exponent for $T_{tts}(N)$ relative to the standard uniform sweep, and for the two easier

difficulty regimes, it outperformed all other studied methods. We expect that these results will carry over to the larger class of optimization problems that experience wrong-way steering towards false minima.

VIII. CONCLUSION

In this work, we defined a simple toy model—the asymmetric magnetization problem—with two competing minima separated by a global peak, and used it to benchmark a variety of modifications to quantum annealing in the literature. The problem is exponentially difficult to solve due to its exponentially closing gap, and the entropic steering toward a false minimum responsible for its difficulty is a generic bottleneck mechanism for a huge array of optimization problem classes. Thus, methods to accelerate finding the solution in it should prove beneficial in much broader contexts.

We studied an ensemble of problem model sets with descending difficulty: $\{A = 0.2, x_p = 0.8\}$, $\{A = 0.28, x_p = 0.7\}$, $\{A = 0.3, x_p = 0.64\}$, and $\{A = 0.34, x_p = 0.59\}$, and assessed a variety of new quantum methods by evaluating the scaling of the time to solution (T_{its}). To have a straightforward view of the performance of each method, we fit their T_{its} to exponential functions and extracted the exponential scaling value, summarized in Table I. The standard uniform sweep method shows inverse gap-squared dependence as expected. In contrast, in the inhomogeneous driving approach, the T_{its} has a nearly constant difficulty scaling of $2^{\sim 7N/10}$, roughly independent of the tuning parameters. It likely means that the problem is steered less toward the false minimum in this case than it is for uniform driving, but that is not enough to avoid a first-order transition and the resulting lack of guidance becomes counterproductive when the problem is easier.

While the problem we studied is homogeneous (in that energy is a function of total magnetization only), we do not expect disorder to significantly change the results for the standard uniform sweep, couplers, and RFQA. If we modeled disorder as a simple random Z term with magnitude $1/N$ for each spin [recall that the problem energy is $O(1)$] added to H_p , it would change the relative energies of the ground states by $1/\sqrt{N}$ for the energy scale we have chosen. This will move the transition point s_c around from one instance to the next, but by an amount that vanishes as $N \rightarrow \infty$. Given that our modified forward approximation calculation predicts Δ_{\min} fairly accurately, examination of those equations shows that this change should not affect the scaling exponent at large N . However, it might affect inhomogeneous driving more significantly, as has been seen in other problems [5,24,25,60].

For the transverse coupler method, both ferromagnetic and antiferromagnetic coupler terms can provide obvious improvements, but adding a mixture of ferromagnetic and antiferromagnetic coupler terms proved counterproductive. We expect that the speedup from the coupler terms arises from creating more tunneling paths between the two competing minima, since each coupler flips two spins simultaneously. Interestingly, we saw very similar scaling benefits for both ferromagnetic (stoquastic) and antiferromagnetic (nonstoquastic) coupler methods; there is no obvious connection between nonstoquasticity and increased performance in this problem.

Among the RFQA methods, synchronized RFQA-M with the added couplers provided the greatest quantum speedup in the two hardest problem sets, while RFQA-D showed the best scaling in the easier problem instances. The speedup mechanism for both couplers and RFQA methods is due to an amplification of the tunneling rate and has nothing to do with local energetic guidance.

We conclude that although we did not achieve an exponential speedup for this problem, all of the methods can provide a quantum speedup over the standard uniform sweep routine: inhomogeneous driving provides a significant boost in harder problem sets; transverse couplers added to the standard uniform sweep routine create more tunneling paths between two competing minima and help decrease the time needed to find the solution; and the introduction of oscillating fields in the RFQA methods can help to stimulate multitone transitions, providing more possibilities for the two competing minima to mix. Undoubtedly, the relative advantages of the methods will certainly depend on the problem class, but since our AMP model exhibits such a generic bottleneck mechanism, any methods to accelerate finding the solution in it should be very widely applicable. Given that all three of inhomogeneous driving, RFQA-M, and RFQA-D only require modification to the control circuitry and not the qubit hardware itself, we see them as the most promising and cost-effective routes to a near-term benefit. It would be worthy to continue to study these methods on realistic problems at larger scales, and it requires minimal changes to existing hardware to verify their potential in experiment.

ACKNOWLEDGMENTS

Z.J.T. thanks Nicholas Materise and David Rodriguez Perez for helpful discussions. This work was supported by the National Science Foundation through Grant No.




FIG. 9. Numerical values of the minimum gap for four parameter sets in the standard uniform sweep method, with the system size N ranging from 5 to 12 spins. The markers are data of numerically estimated minimum gap energies, and least-squares fits of the minimum gap are red, orange, green, and blue solid lines. The minimum gaps in the four problem sets all decrease exponentially with increasing system size.




FIG. 10. Ratio of analytical and numerical values of minimum gap for four problem sets in the standard uniform sweep method, system size N ranging from 5 to 12. The analytical values of Δ_{\min} are from the modified forward approximation [Eqs. (A1)–(A3)]. The markers represent the ratio of analytical and numerical values of minimum gap in four problem sets. The dash-dotted line, dashed line, line with point marker, and solid line represent the ratio of analytical and numerical values with a correction term $2\pi/N$ for problem sets $\{A = 0.34, x_p = 0.59\}$, $\{A = 0.3, x_p = 0.64\}$, $\{A = 0.28, x_p = 0.7\}$, and $\{A = 0.2, x_p = 0.8\}$ separately. The corrected analytical predictions in the four sets match well with the numerical values.

PHY-1653820. It was made possible by the high performance computing resources from the Tulane University Cypress platform, and the Colorado School of Mines Wendian platform.

APPENDIX: ANALYTICAL PREDICTION OF THE MINIMUM GAP

The minimum gap Δ_{\min} determines the worst-case difficulty of a problem, so analytically predicting it can help us to better assess the behavior of the quantum annealing algorithm. To compute it, we use a modified form of the “forward approximation” N th-order perturbation theory employed in [61–65].

In this approximation, the minimum gap is predicted to be

$$\Delta_{\min} = N! \frac{\prod_{i=1}^N \kappa_i}{\prod_{n=1}^{N-1} U_n}, \quad (\text{A1})$$

where κ_i is the transverse field strength on each site i (evaluated at the critical point κ_c), and U_n^{-1} is the average of the inverse of the energy difference to flip n spins from either ground state:

$$\frac{1}{U_n} = \left\langle \frac{1}{\epsilon_n + \delta_{\epsilon_n} - \epsilon_0 - \delta_{\epsilon_0}} \right\rangle. \quad (\text{A2})$$

Here, the ϵ terms are the classical energies defined in the problem Hamiltonian and the δ terms are their perturbative corrections from the transverse field, which act to increase the excitation energies in this case. Including these corrections in the energy denominators (which is effectively a resummation scheme) is vital to obtaining relatively accurate predictions; explicitly, for the AMP

$$\begin{aligned} U_n &\simeq n \left(\frac{1}{x_p} + 2\kappa_c^2 x_p \right), \quad \{n \leq x_p N\}, \\ &\simeq (N - n) \left(\frac{1 + A}{1 - x_p} + 2\kappa_c^2 \frac{1 - x_p}{1 + A} \right), \quad \{n > x_p N\}. \end{aligned} \quad (\text{A3})$$

From this expression, it is straightforward to predict the minimum gap in our problem. As shown in Fig. 9, the exponential fittings of numerical Δ_{\min} as a function of N for descending-difficulty problem sets are $2^{-0.4-1.12N}$, $2^{-0.79-0.74N}$, $2^{-1.35-0.52N}$, and $2^{-2.04-0.31N}$. Figure 10 indicates that the scaling of our theoretical prediction matches well with the numerical result by multiplying by a factor of $2\pi/N$. Equation (A1) appears to overestimate the true gap by a factor of $\sim N/2\pi$; the reason for this is unclear. Some level of disagreement is expected, however, particularly in the easiest of the four parameter sets. If the coefficient of the problem Hamiltonian is 1, the phase transition for those parameters occurs at $\kappa_c \simeq 1.73$. At such a large value of κ_c the ratio of κ to the single spin excitation energy approaches unity and thus a perturbative expansion in it may break down. As the log of the minimum gap is an integral of a function of total magnetization and the order in which spins are flipped does not matter, we expect that other Hamiltonians which are simple polynomials of total magnetization (such as p -spin models) would likely show similar physics to our AMP.

-
- [1] A. B. Finnila, M. A. Gomez, C. Sebenik, C. Stenson, and J. D. Doll, Quantum annealing: A new method for minimizing multidimensional functions, *Chem. Phys. Lett.* **219**, 343 (1994).
- [2] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse Ising model, *Phys. Rev. E* **58**, 5355 (1998).
- [3] A. Das and B. K. Chakrabarti, Colloquium: Quantum annealing and analog quantum computation, *Rev. Mod. Phys.* **80**, 1061 (2008).
- [4] M. W. Johnson, M. H. S. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk, E. M. Chapple, C. Enderud, J. P. Hilton, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Oh, I. Perminov, C. Rich, M. C. Thom, E. Tolkacheva, C. J. S. Truncik, S. Uchaikin, J. Wang, B. Wilson, and G. Rose, Quantum annealing with manufactured spins, *Nature (London)* **473**, 194 (2011).
- [5] R. D. Somma, D. Nagaj, and M. Kieferová, Quantum Speedup by Quantum Annealing, *Phys. Rev. Lett.* **109**, 050501 (2012).
- [6] T. Albash and D. A. Lidar, Adiabatic quantum computation, *Rev. Mod. Phys.* **90**, 015002 (2018).
- [7] S. Boixo, T. F. Rønnow, S. V. Isakov, Z. Wang, D. Wecker, D. A. Lidar, J. M. Martinis, and M. Troyer, Evidence for quantum

- annealing with more than one hundred qubits, *Nat. Phys.* **10**, 218 (2014).
- [8] G. Rosenberg, P. Haghnegahdar, P. Goddard, P. Carr, K. Wu, and M. L. de Prado, Solving the optimal trading trajectory problem using a quantum annealer, *IEEE J. Sel. Top. Signal Process.* **10**, 1053 (2016).
- [9] F. Neukart, G. Compostella, C. Seidel, D. von Dollen, S. Yarkoni, and B. Parney, Traffic flow optimization using a quantum annealer, *Front. ICT* **4**, 29 (2017).
- [10] D. Venturelli, M. Do, E. Rieffel, and J. Frank, Compiling quantum circuits to realistic hardware architectures using temporal planners, *Quantum Sci. Technol.* **3**, 025004 (2018).
- [11] A. D. King, J. Carrasquilla, J. Raymond, I. Ozfidan, E. Andriyash, A. Berkley, M. Reis, T. Lanting, R. Harris, F. Altomare, K. Boothby, P. I. Bunyk, C. Enderud, A. Fréchet, E. Hoskinson, N. Ladizinsky, T. Oh, G. Poulin-Lamarre, C. Rich, Y. Sato, A. Yu. Smirnov, L. J. Swenson, M. H. Volkmann, J. Whittaker, J. Yao, E. Ladizinsky, M. W. Johnson, J. Hilton, and M. H. Amin, Observation of topological phenomena in a programmable lattice of 1,800 qubits, *Nature (London)* **560**, 456 (2018).
- [12] D. Venturelli and A. Kondratyev, Reverse quantum annealing approach to portfolio optimization problems, *Quantum Mach. Intell.* **1**, 17 (2019).
- [13] A. D. King, J. Raymond, T. Lanting, S. V. Isakov, M. Mohseni, G. Poulin-Lamarre, S. Ejtemaee, W. Bernoudy, I. Ozfidan, A. Yu. Smirnov, M. Reis, F. Altomare, M. Babcock, C. Baron, A. J. Berkley, K. Boothby, P. I. Bunyk, H. Christiani, C. Enderud, B. Evert, R. Harris, E. Hoskinson, S. Huang, K. Jooya, A. Khodabandelou, N. Ladizinsky, R. Li, P. A. Lott, A. J. R. MacDonald, D. Marsden, G. Marsden, T. Medina, R. Molavi, R. Neufeld, M. Norouzpour, T. Oh, I. Pavlov, I. Perminov, T. Prescott, C. Rich, Y. Sato, B. Sheldan, G. Sterling, L. J. Swenson, N. Tsai, M. H. Volkmann, J. D. Whittaker, W. Wilkinson, J. Yao, H. Neven, J. P. Hilton, E. Ladizinsky, M. W. Johnson, and M. H. Amin, Scaling advantage in quantum simulation of geometrically frustrated magnets, [arXiv:1911.03446](https://arxiv.org/abs/1911.03446).
- [14] J. King, M. Mohseni, W. Bernoudy, A. Fréchet, H. Sadeghi, S. V. Isakov, H. Neven, and M. H. Amin, Quantum-assisted genetic algorithm, [arXiv:1907.00707](https://arxiv.org/abs/1907.00707).
- [15] J. Roland and N. J. Cerf, Quantum search by local adiabatic evolution, *Phys. Rev. A* **65**, 042308 (2002).
- [16] A. Lucas, Ising formulations of many NP problems, *Front. Phys.* **2**, 5 (2014).
- [17] T. Jörg, F. Krzakala, J. Kurchan, A. C. Maggs, and J. Pujos, Energy gaps in quantum first-order mean-field-like transitions: The problems that quantum annealing cannot solve, *Europhys. Lett.* **89**, 40004 (2010).
- [18] E. Farhi, J. Goldstone, and S. Gutmann, Quantum adiabatic evolution algorithms versus simulated annealing, [arXiv:quant-ph/0201031](https://arxiv.org/abs/quant-ph/0201031).
- [19] Y. Seki and H. Nishimori, Quantum annealing with antiferromagnetic fluctuations, *Phys. Rev. E* **85**, 051112 (2012).
- [20] E. Crosson, E. Farhi, C. Yen-Yu Lin, H.-H. Lin, and P. Shor, Different strategies for optimization using the quantum adiabatic algorithm, [arXiv:1401.7320](https://arxiv.org/abs/1401.7320).
- [21] Y. Seki and H. Nishimori, Quantum annealing with antiferromagnetic transverse interactions for the Hopfield model, *J. Phys. A: Math. Theor.* **48**, 335301 (2015).
- [22] H. Nishimori and K. Takada, Exponential enhancement of the efficiency of quantum annealing by non-stoquastic Hamiltonians, *Front. ICT* **4**, 2 (2017).
- [23] L. Hormozi, E. W. Brown, G. Carleo, and M. Troyer, Nonstoquastic Hamiltonians and quantum annealing of an Ising spin glass, *Phys. Rev. B* **95**, 184416 (2017).
- [24] Y. Susa, Y. Yamashiro, M. Yamamoto, and H. Nishimori, Exponential speedup of quantum annealing by inhomogeneous driving of the transverse field, *J. Phys. Soc. Jpn.* **87**, 023002 (2018).
- [25] Y. Susa, Y. Yamashiro, M. Yamamoto, I. Hen, D. A. Lidar, and H. Nishimori, Quantum annealing of the p -spin model under inhomogeneous transverse field driving, *Phys. Rev. A* **98**, 042326 (2018).
- [26] E. Kapit and V. Oganesyan, Noise-tolerant quantum speedups in quantum annealing without fine tuning, *Quantum Science and Technology* (2020).
- [27] L. Kong and E. Crosson, The performance of the quantum adiabatic algorithm on spike Hamiltonians, *Int. J. Quantum Information* **15**, 1750011 (2017).
- [28] M. V. Berry, Transitionless quantum driving, *J. Phys. A: Math. Theor.* **42**, 365303 (2009).
- [29] J. I. Adame and P. McMahon, Inhomogeneous driving in quantum annealers can result in orders-of-magnitude improvements in performance, *Quantum Sci. Technol.* **5**, 035011 (2020).
- [30] T. Albash and D. A. Lidar, Demonstration of a Scaling Advantage for a Quantum Annealer over Simulated Annealing, *Phys. Rev. X* **8**, 031016 (2018).
- [31] S. Knysh, Zero-temperature quantum annealing bottlenecks in the spin-glass phase, *Nat. Commun.* **7**, 12370 (2016).
- [32] J. Dziarmaga, Dynamics of a quantum phase transition and relaxation to a steady state, *Adv. Phys.* **59**, 1063 (2010).
- [33] C. R. Laumann, R. Moessner, A. Scardicchio, and S. L. Sondhi, Quantum Adiabatic Algorithm and Scaling of Gaps at First-Order Quantum Phase Transitions, *Phys. Rev. Lett.* **109**, 030502 (2012).
- [34] D. Sels and A. Polkovnikov, Minimizing irreversible losses in quantum systems by local counterdiabatic driving, *Proc. Natl. Acad. Sci. USA* **114**, E3909 (2017).
- [35] W. Vinci and D. A. Lidar, Non-stoquastic Hamiltonians in quantum annealing via geometric phases, *npj Quantum Inf.* **3**, 38 (2017).
- [36] J. Marshall, D. Venturelli, I. Hen, and E. G. Rieffel, Power of Pausing: Advancing Understanding of Thermalization in Experimental Quantum Annealers, *Phys. Rev. Appl.* **11**, 044083 (2019).
- [37] T. Graß, Quantum Annealing with Longitudinal Bias Fields, *Phys. Rev. Lett.* **123**, 120501 (2019).
- [38] P. Hauke, H. G. Katzgraber, W. Lechner, H. Nishimori, and W. D. Oliver, Perspectives of quantum annealing: Methods and implementations, *Rep. Prog. Phys.* **83**, 054401 (2020).
- [39] S. Bravyi, D. P. Divincenzo, R. I. Oliveira, and B. M. Terhal, The complexity of stoquastic local Hamiltonian problems [arXiv:quant-ph/0606140](https://arxiv.org/abs/quant-ph/0606140).
- [40] E. Andriyash and M. H. Amin, Can quantum Monte Carlo simulate quantum annealing?, [arXiv:1703.09277](https://arxiv.org/abs/1703.09277).
- [41] M. Ohzeki, Quantum Monte Carlo simulation of a particular class of non-stoquastic Hamiltonians in quantum annealing, *Sci. Rep.* **7**, 41186 (2017).

- [42] V. S. Denchev, S. Boixo, S. V. Isakov, N. Ding, R. Babbush, V. Smelyanskiy, J. Martinis, and H. Neven, What Is the Computational Value of Finite-Range Tunneling?, *Phys. Rev. X* **6**, 031015 (2016).
- [43] S. V. Isakov, G. Mazzola, V. N. Smelyanskiy, Z. Jiang, S. Boixo, H. Neven, and M. Troyer, Understanding Quantum Tunneling through Quantum Monte Carlo Simulations, *Phys. Rev. Lett.* **117**, 180402 (2016).
- [44] Z. Jiang, V. N. Smelyanskiy, S. V. Isakov, S. Boixo, G. Mazzola, M. Troyer, and H. Neven, Scaling analysis and instantons for thermally assisted tunneling and quantum Monte Carlo simulations, *Phys. Rev. A* **95**, 012322 (2017).
- [45] G. Mazzola, V. N. Smelyanskiy, and M. Troyer, Quantum Monte Carlo tunneling from quantum chemistry to quantum annealing, *Phys. Rev. B* **96**, 134305 (2017).
- [46] Y. Susa, J. F. Jadebeck, and H. Nishimori, Relation between quantum fluctuations and the performance enhancement of quantum annealing in a nonstoquastic Hamiltonian, *Phys. Rev. A* **95**, 042321 (2017).
- [47] Y. Matsuda, H. Nishimori, and H. G. Katzgraber, Ground-state statistics from annealing algorithms: Quantum versus classical approaches, *New J. Phys.* **11**, 073021 (2009).
- [48] E. Farhi, J. Goldstone, and S. Gutmann, Quantum adiabatic evolution algorithms with different paths, [arXiv:quant-ph.0208135](https://arxiv.org/abs/quant-ph/0208135).
- [49] A. Perdomo-Ortiz, S. E. Venegas-Andraca, and A. Aspuru-Guzik, A study of heuristic guesses for adiabatic quantum computation, *Quantum Inf. Process.* **10**, 33 (2011).
- [50] N. Chancellor, Modernizing quantum annealing using local searches, *New J. Phys.* **19**, 023024 (2017).
- [51] M. Ohkuwa, H. Nishimori, and D. A. Lidar, Reverse annealing for the fully connected p -spin model, *Phys. Rev. A* **98**, 022314 (2018).
- [52] M. Keck, S. Montangero, G. E. Santoro, R. Fazio, and D. Rossini, Dissipation in adiabatic quantum computers: Lessons from an exactly solvable model, *New J. Phys.* **19**, 113029 (2017).
- [53] V. N. Smelyanskiy, D. Venturelli, A. Perdomo-Ortiz, S. Knysh, and M. I. Dykman, Quantum Annealing via Environment-Mediated Quantum Diffusion, *Phys. Rev. Lett.* **118**, 066802 (2017).
- [54] L. C. Venuti, T. Albash, M. Marvian, D. Lidar, and P. Zanardi, Relaxation versus adiabatic quantum steady-state preparation, *Phys. Rev. A* **95**, 042302 (2017).
- [55] L. Arceci, S. Barbarino, D. Rossini, and G. E. Santoro, Optimal working point in dissipative quantum annealing, *Phys. Rev. B* **98**, 064307 (2018).
- [56] T. Kadowaki and M. Ohzeki, Experimental and theoretical study of thermodynamic effects in a quantum annealer, *J. Phys. Soc. Jpn.* **88**, 061008 (2019).
- [57] S. Suzuki, H. Oshiyama, and N. Shibata, Quantum annealing of pure and random Ising chains coupled to a bosonic environment, *J. Phys. Soc. Jpn.* **88**, 061003 (2019).
- [58] D. Roberts, L. Cincio, A. Saxena, A. Petukhov, and S. Knysh, Noise amplification at spin-glass bottlenecks of quantum annealing: A solvable model, *Phys. Rev. A* **101**, 042317 (2020).
- [59] D. Jaschke, L. D. Carr, and I. de Vega, Thermalization in the quantum Ising model—approximations, limits, and beyond, *Quantum Sci. Technol.* **4**, 034002 (2019).
- [60] E. J. Crosson and D. A. Lidar, Prospects for quantum enhancement with diabatic quantum annealing, [arXiv:2008.09913](https://arxiv.org/abs/2008.09913).
- [61] F. Pietracaprina, V. Ros, and A. Scardicchio, Forward approximation as a mean-field approximation for the Anderson and many-body localization transitions, *Phys. Rev. B* **93**, 054201 (2016).
- [62] C. L. Baldwin, C. R. Laumann, A. Pal, and A. Scardicchio, The many-body localized phase of the quantum random energy model, *Phys. Rev. B* **93**, 024202 (2016).
- [63] C. L. Baldwin, C. R. Laumann, A. Pal, and A. Scardicchio, Clustering of Nonergodic Eigenstates in Quantum Spin Glasses, *Phys. Rev. Lett.* **118**, 127201 (2017).
- [64] A. Scardicchio and T. Thiery, Perturbation theory approaches to Anderson and many-body localization: Some lecture notes, [arXiv:1710.01234](https://arxiv.org/abs/1710.01234).
- [65] C. L. Baldwin and C. R. Laumann, Quantum algorithm for energy matching in hard optimization problems, *Phys. Rev. B* **97**, 224201 (2018).