

Tunneling and Speedup in Quantum Optimization for Permutation-Symmetric Problems

Siddharth Muthukrishnan,^{1,2} Tameem Albash,^{1,2,3} and Daniel A. Lidar^{1,2,4,5}

¹*Department of Physics and Astronomy, University of Southern California, Los Angeles, California 90089, USA*

²*Center for Quantum Information Science and Technology, University of Southern California, Los Angeles, California 90089, USA*

³*Information Sciences Institute, University of Southern California, Marina del Rey, California 90292, USA*

⁴*Department of Electrical Engineering, University of Southern California, Los Angeles, California 90089, USA*

⁵*Department of Chemistry, University of Southern California, Los Angeles, California 90089, USA*
(Received 10 May 2015; revised manuscript received 11 May 2016; published 21 July 2016)

Tunneling is often claimed to be the key mechanism underlying possible speedups in quantum optimization via quantum annealing (QA), especially for problems featuring a cost function with tall and thin barriers. We present and analyze several counterexamples from the class of perturbed Hamming weight optimization problems with qubit permutation symmetry. We first show that, for these problems, the adiabatic dynamics that make tunneling possible should be understood not in terms of the cost function but rather the semiclassical potential arising from the spin-coherent path-integral formalism. We then provide an example where the shape of the barrier in the final cost function is short and wide, which might suggest no quantum advantage for QA, yet where tunneling renders QA superior to simulated annealing in the adiabatic regime. However, the adiabatic dynamics turn out not to be optimal. Instead, an evolution involving a sequence of diabatic transitions through many avoided-level crossings, involving no tunneling, is optimal and outperforms adiabatic QA. We show that this phenomenon of speedup by diabatic transitions is not unique to this example, and we provide an example where it provides an exponential speedup over adiabatic QA. In yet another twist, we show that a classical algorithm, spin-vector dynamics, is at least as efficient as diabatic QA. Finally, in a different example with a convex cost function, the diabatic transitions result in a speedup relative to both adiabatic QA with tunneling and classical spin-vector dynamics.

DOI: [10.1103/PhysRevX.6.031010](https://doi.org/10.1103/PhysRevX.6.031010)

Subject Areas: Computational Physics, Quantum Physics

I. INTRODUCTION

The possibility of a quantum speedup for finding the solution of classical optimization problems is tantalizing, as a quantum advantage for this class of problems would provide a wealth of new applications for quantum computing. The goal of many optimization problems can be formulated as finding an n -bit string x_{opt} that minimizes a given cost function $f(x)$, which can be interpreted as the energy of a classical Ising spin system whose ground state is x_{opt} . Finding the ground state of such systems can be hard if, e.g., the system is strongly frustrated, resulting in a complex energy landscape that cannot be efficiently explored with any known algorithm due to the presence of many local minima [1]. This can occur, e.g., in classical simulated annealing (SA) [2], when the system's state is trapped in a local minimum.

Thermal hopping and quantum tunneling provide two starkly different mechanisms for solving optimization

problems, and finding optimization problems that favor the latter continues to be an open theoretical question [3,4]. It is often stated that quantum annealing (QA) [5–9] uses tunneling instead of thermal excitations to escape from local minima, which can be advantageous in systems with tall but thin barriers that are easier to tunnel through than to thermally climb over [4,9,10]. It is with this potential tunneling-induced advantage over classical annealing that QA and the quantum adiabatic algorithm [11] were proposed. Our goal in this work is to address the question of the role played by tunneling in providing a quantum speedup, and to elucidate it by studying a number of illustrative examples. We demonstrate that the role of tunneling is significantly more subtle than what might be expected on the basis of the “tall and thin barrier” picture.

In order to make progress on this question, the potential with respect to which tunneling occurs must be clearly specified. Tunneling, in quantum mechanics, is always defined with respect to a potential that delineates classically allowed and forbidden regions. This potential is usually obtained by the semiclassical (SC) limit of the quantum dynamics: i.e., by the potential energy term of the action obtained in a path-integral representation of the quantum dynamics. In QA, one typically initializes the system in the

Published by the American Physical Society under the terms of the Creative Commons Attribution 3.0 License. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

known ground state of a simple Hamiltonian and evolves the system towards a Hamiltonian representing the final cost function. We argue that when one takes a natural semiclassical limit, the semiclassical potential does not become the final cost function. Instead, one obtains a potential appearing in the action of the spin-coherent path-integral representation of the quantum dynamics. This potential, which here we call the spin-coherent potential, was first considered in the setting of adiabatic quantum computation by Refs. [12,13]. It has also been used profitably before in Refs. [14,15].

We analyze the spin-coherent potential for several examples from a well-known class of problems known as perturbed Hamming weight oracle (PHWO) problems. These are problems for which instances can be generated where QA either has an advantage over classical random search algorithms with local updates, such as SA [12,16], or has no advantage [16,17]. Moreover, because PHWO problems exhibit qubit permutation symmetry, their quantum evolutions are easily classically simulatable, and furthermore, their spin-coherent potential is one dimensional. Tunneling becomes clear and explicit for these problems when using the spin-coherent potential. While earlier works had already indicated that tunneling can be understood via this potential, here we provide comprehensive evidence for this claim. In the process, we provide a list of intuitive criteria that defines what one means by tunneling in the examples studied (see Sec. III).

Having obtained a clear picture of tunneling, we focus on a particular PHWO problem that has a plateau in the final cost function (henceforth, “the fixed plateau”). This problem offers a counterexample to two commonly held views: (1) QA has an advantage, due to tunneling, over SA only on problems where the barrier in the final cost function is tall and thin, and (2) tunneling is necessary for a quantum speedup in QA. We refute the first statement by showing that for the fixed plateau, which is a short and wide cost function, QA significantly outperforms SA by using tunneling. Indeed, we find numerically that adiabatic QA (AQA) needs a time of $\mathcal{O}(n^{0.5})$ to find the ground state, where n is the number of spins or qubits. Moreover, using the spin-coherent potential, we observe the presence of tunneling during the quantum anneal. On the other hand, we prove that single-spin update SA takes a time of $\mathcal{O}(n^{\text{plateau width}})$. Thus, we have essentially an arbitrary polynomial tunneling speedup of QA over SA on a cost function that is neither tall nor thin. We remark that this result about SA’s performance is also a rigorous proof of a result due to Reichardt [16] that classical local search algorithms will fail on a certain class of PHWO problems and is of independent interest.

We refute the second statement by showing that, for the fixed plateau, it is actually optimal to run QA diabatically (henceforth, DQA for diabatic quantum annealing). The system leaves the ground state, only to return through a

sequence of diabatic transitions associated with avoided-level crossings. In this regime, the run time for QA is $\mathcal{O}(1)$. Moreover, in this regime, we do not observe any of the standard signatures of tunneling. We show that this feature—that the optimal evolution time t_f for QA is far from being adiabatic—is present in a few other PHWO problems and that this optimal evolution involves no multiqubit tunneling.

Given that the optimal evolution involves no tunneling, we are inspired to investigate a classical algorithm, spin-vector dynamics (SVD), which can be interpreted as a semiclassical limit of the quantum evolution with a product-state approximation. We observe that SVD evolves in an almost identical manner to DQA, and is able to recover the speedup by DQA. Thus, in these problems, we show that what may be suspected to be a highly quantum-coherent process—diabatic transitions—can be mimicked by a quantum-inspired classical algorithm.

The structure of this paper is as follows. In Sec. II, we list the PHWO problems we study. In Sec. III, we use these problems to present evidence that tunneling can be understood with respect to the spin-coherent potential. In Sec. IV, we focus on the fixed plateau PHWO problem and exhaustively analyze the performance of various algorithms for this problem. In particular, we numerically characterize AQA (Sec. IV A), provide a rigorous proof of SA’s performance (Sec. IV B), and numerically analyze DQA (Sec. IV C), SVD (Sec. IV D), and a quantum Monte Carlo algorithm (Sec. IV E). We conclude in Sec. V by discussing the implications of our work and possible directions for future work. Additional background information and technical details can be found in the Appendixes.

II. PERTURBED HAMMING WEIGHT OPTIMIZATION PROBLEMS AND THE EXAMPLES STUDIED

The cost function of a PHWO problem is defined as

$$f(x) = \begin{cases} |x| + p(|x|) & l < |x| < u \\ |x| & \text{elsewhere,} \end{cases} \quad (1)$$

where $|x|$ denotes the Hamming weight [18] of the bit string $x \in \{0, 1\}^n$. Since $|x|$ is invariant under permutation of bits, so is $f(x) = f(|x|)$. For SA, this is the cost function. For QA, this will be proportional to the energy of the final Hamiltonian. More precisely, we define QA as the closed-system quantum evolution governed by the time-dependent Hamiltonian,

$$H(s) = \frac{1}{2}(1-s) \sum_i (\mathbb{1} - \sigma_i^x) + s \sum_x f(x) |x\rangle\langle x|, \quad (2)$$

where we have chosen the standard transverse field “driver” Hamiltonian $H(0)$ that assumes no prior knowledge of the form of $f(x)$, and a linear interpolating schedule,

with $s \equiv t/t_f$ being the dimensionless time parameter. The initial state is the ground state of $H(0)$.

Below, we list several of PHWO examples that we study in greater detail. We refer to the case with $p = 0$ as the plain Hamming weight problem.

(1) Fixed plateau:

$$f(x) = \begin{cases} u - 1 & l < |x| < u \\ |x| & \text{otherwise.} \end{cases} \quad (3)$$

Clearly, this forms a plateau in Hamming weight space. We take $u, l = \mathcal{O}(1)$. Since the location of the plateau does not change with n , we refer to it as “fixed.” An instance of this cost function with $l = 3$ and $u = 8$ is illustrated in Fig. 1. By numerical diagonalization, we find that QA has a constant gap for this cost function.

(2) Reichardt:

$$f(x) = \begin{cases} |x| + h(n) & l(n) < |x| < u(n) \\ |x| & \text{otherwise,} \end{cases} \quad (4)$$

with $h[(u - l)/\sqrt{l}] = o(1)$. For this case, Reichardt [16] proved a constant lower bound on the minimum spectral gap during the quantum anneal. In Appendix A, we provide a pedagogical review of this proof and fill in some details not explicitly provided in the original proof.

(3) Moving plateau:

$$f(x) = \begin{cases} u - 1 & l(n) < |x| < u(n) \\ |x| & \text{otherwise,} \end{cases} \quad (5)$$

with $l(n) = n/4$ and $u(n) = \mathcal{O}(1)$. This is termed “moving” since the location of the plateau changes with n . Note that this is a special case from the Reichardt class.

(4) Grover:

$$f(x) = \begin{cases} n & |x| \geq 1 \\ 0 & |x| = 0. \end{cases} \quad (6)$$

This is a minor modification of the standard Grover problem: the marked state is the all-zeros string with energy 0, and the energy of all the other states is n . Scaling the energy by n keeps the maximum energy of all the PHWO problems we consider comparable.

(5) Spike:

$$f(x) = \begin{cases} n & |x| = n/4 \\ |x| & \text{otherwise.} \end{cases} \quad (7)$$

This was studied by Farhi *et al.* in Ref. [12], where it was argued that the quantum minimum gap scales

as $\mathcal{O}(n^{-1/2})$ and that SA will take exponential time to find the ground state. However, we show below (Fig. 8) that SVD is more efficient than QA for this problem.

(6) Precipice:

$$f(x) = \begin{cases} -1 & |x| = n \\ |x| & \text{otherwise.} \end{cases} \quad (8)$$

This was studied by van Dam *et al.* in Ref. [17], where it was proved that the quantum minimum gap for this problem scales as $\mathcal{O}(2^{-n/2})$.

(7) α rectangle:

$$f(x) = \begin{cases} |x| + n^\alpha & \frac{n}{4} - \frac{1}{2}cn^\alpha < |x| < \frac{n}{4} + \frac{1}{2}cn^\alpha \\ |x| & \text{otherwise.} \end{cases} \quad (9)$$

We call this an α rectangle because the width of the perturbation (cn^α) is c times the height. This was studied in Ref. [19], where evidence for the following conjecture for the scaling of the quantum minimum gap g_{\min} was presented:

$$g_{\min} = \begin{cases} \text{const} & \alpha < \frac{1}{4} \\ 1/\text{poly}(n) & \frac{1}{4} < \alpha < \frac{1}{3} \\ 1/\exp(n) & \alpha > \frac{1}{3}. \end{cases} \quad (10)$$

Note that $\alpha < 1/4$ is a member of the Reichardt class, and thus the constant lower bound on the minimum gap is a theorem, and not a conjecture. We restrict ourselves to the case of $c = 1$.

Notice that the different examples listed above supply different types of barriers to single-spin update SA. Some of them are purely energetic barriers: e.g., the spike. Some of them are purely entropic barriers [20]: e.g., the plateau. And some of them are a combination of both: e.g., the α -rectangle problems. We show in our analysis that when viewed via the SC potential, these entropic barriers get translated into energy barriers for the adiabatic dynamics.




FIG. 1. Fixed plateau cost function with $l = 3$, $u = 8$.

We remark that all the problems listed above are representative members of a large family of problems: if the input bit string to any of the above problems is transformed by an XOR mask, then all of our analysis below will hold. For QA, the XOR mask can be represented as a unitary transformation: $\bigotimes_{i=1}^n (\sigma_i^x)^{a_i}$, with $a \in \{0, 1\}^n$ being the mask string. As this unitary commutes with the QA Hamiltonian at all times, none of our subsequent analysis is affected. Similar arguments go through for SA and all the other algorithms that we consider.

We note that PHWO problems are strictly toy problems since these problems are typically represented by highly nonlocal Hamiltonians (see Appendix B) and thus are not physically implementable, in the same sense that the adiabatic Grover search problem is unphysical [21,22]. Nevertheless, these problems provide us with important insights into the mechanisms behind a quantum speedup (or lack thereof), and illuminate the relative advantages of one computational model over another. In our case, among other lessons, these problems show us the ability of quantum annealing to tunnel, which SA is unable to do. Analysis of these problems could also be viewed as stepping stones towards realizable quantum algorithms.

III. SEMICLASSICAL POTENTIAL AND TUNNELING

In order to study tunneling, we need a potential arising from a semiclassical limit, which defines classically allowed and forbidden regions. One approach to writing a semiclassical potential for quantum Hamiltonians is to use the spin-coherent path-integral formalism [23]. This

semiclassical potential has been used profitably in various QA studies, e.g., Refs. [12–15], and we extend its applications here. For the quantum evolution, since the initial state [the ground state of $H(0)$] is symmetric under permutations of qubits and the unitary dynamics preserves this symmetry [it is a symmetry of $H(s)$ for all s], we can consistently restrict ourselves to spin-1/2 symmetric coherent states $|\theta, \phi\rangle$:

$$|\theta, \phi\rangle = \bigotimes_{i=1}^n \left[\cos\left(\frac{\theta}{2}\right) |0\rangle_i + \sin\left(\frac{\theta}{2}\right) e^{i\phi} |1\rangle_i \right]. \quad (11)$$

The spin-coherent potential is then given by

$$V_{\text{SC}}(\theta, \phi, s) = \langle \theta, \phi | H(s) | \theta, \phi \rangle. \quad (12)$$

We show that for all the examples defined above except the Reichardt class (we address this below), this potential captures important features of the quantum Hamiltonian [Eq. (2)] and reveals the presence of tunneling. Specifically, as follows:

- (1) *The spin-coherent potential displays a degenerate double well almost exactly at the point of the minimum gap.* In Fig. 2(a), we plot, for the fixed plateau, the potential near the minimum gap. The potential transitions from having a single minimum on the right to a single minimum on the left. In between, it becomes degenerate and displays a degenerate double well. Since the instantaneous ground state corresponds to the position of the global minimum, which exhibits a discontinuity, the degeneracy point is where tunneling should be most helpful. In Fig. 3(a), we show that the location

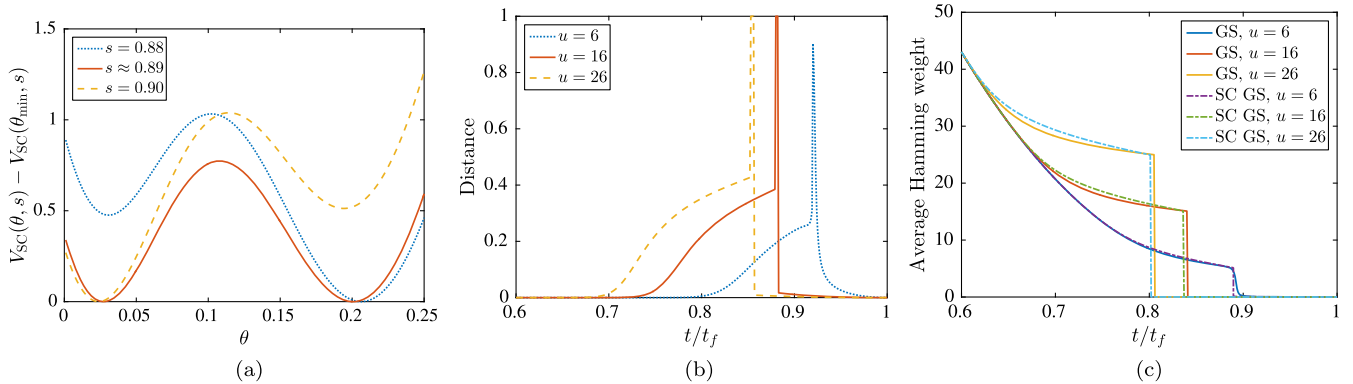


FIG. 2. Results for the fixed plateau problem with $l = 0$ and $n = 512$. (a) The semiclassical potential with $u = 6$ exhibits a double-well degeneracy at the position $s \approx 0.89$ (solid line), but is nondegenerate before and after this point (dotted and dashed lines), thus leading to a discontinuity in the position of its global minimum. The same is observed for the other PHWO problems we study (not shown). (b) The trace-norm distance $\sqrt{1 - |\langle \psi_{\text{GS}} | \psi_{\text{SC GS}} \rangle|^2}$ between the quantum ground state $|\psi_{\text{GS}}\rangle$ (obtained by numerical diagonalization) and the spin-coherent state $|\psi_{\text{SC GS}}\rangle$ corresponding to the instantaneous global minimum in V_{SC} , as a function of t/t_f . The peak corresponds to the location of the tunneling event, at which point the semiclassical approximation breaks down. (c) $\langle \text{HW} \rangle$ in the instantaneous quantum ground state (GS) and the instantaneous ground state as predicted by the semiclassical (SC GS) potential, as a function of t/t_f . The sharp drop in the GS and SC GS curves is due to a tunneling event wherein $\sim u$ qubits are flipped, which occurs at the degeneracy point observed in the spin-coherent potential.

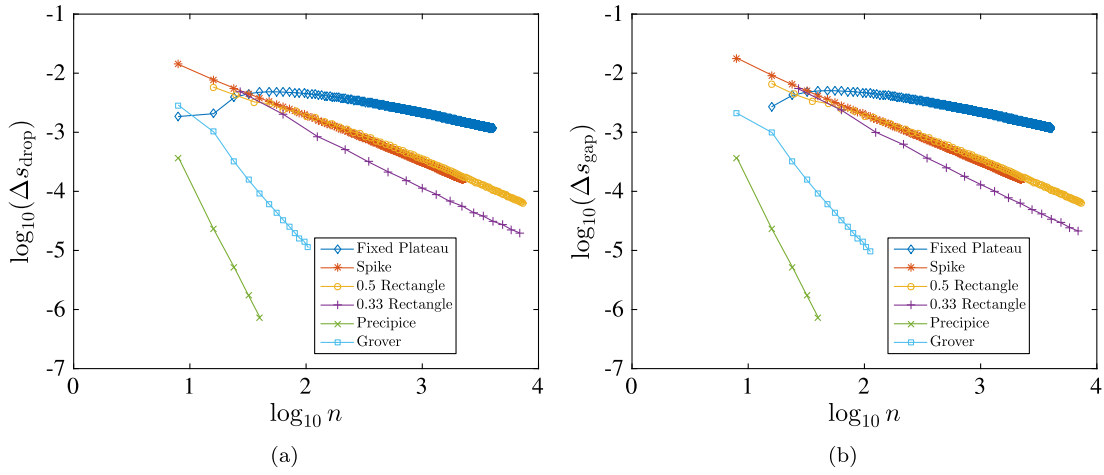


FIG. 3. (a) The difference in the position of the minimum gap from exact diagonalization and the position of the double-well degeneracy [as shown in Fig. 2(a) for the fixed plateau] from the semiclassical potential, as a function of n for the fixed plateau, the spike, the 0.5-rectangle, the 0.33-rectangle, the precipice, and Grover problems (log-log scale). (b) The difference in the position of the sudden drop in Hamming weight [as seen in Fig. 2(c)] and the position of the double-well degeneracy from the semiclassical potential [as seen in Fig. 2(a)], as a function of the number of qubits n for the fixed plateau, the spike, the 0.5-rectangle, the 0.33-rectangle, the precipice, and Grover problems (log-log scale).

of the minimum gap of the quantum evolution is very close to the location of the degenerate double well in the spin-coherent potential.

- (2) *The ground state predicted by the spin-coherent potential is a good approximation to the quantum ground state except near the degeneracy point.* As expected from a potential that arises in a semi-classical limit, the ground state predicted by the spin-coherent potential (i.e., the spin-coherent state corresponding to the instantaneous global minimum in V_{SC}) agrees well with the quantum ground state, except where tunneling is important. In particular, delocalization when the spin-coherent potential is a degenerate double well (or is close to being one) should imply that approximating the ground state with a wave function localized in one of the wells fails. Indeed, we find this to be the case. We illustrate this for the fixed plateau in Fig. 2(b); similar results hold for the other examples we consider.
- (3) *There is a sharp change in the ground state of the adiabatic quantum evolution at the degeneracy point.* Tunneling should be accompanied by a sharp change in the properties of the ground state at the degeneracy point as the state shifts from being localized in one well to the other. We quantify this change by calculating the expectation value of the Hamming weight operator, defined as $\text{HW} = \frac{1}{2} \sum_{i=1}^n (1 - \sigma_i^z)$. We expect a discontinuity in the spin-coherent ground-state expectation value $\langle \text{HW} \rangle$, because the spin-coherent ground state changes discontinuously at the degeneracy point. We find that there is a nearly identical change in the quantum ground-state expectation value $\langle \text{HW} \rangle$ for all of the examples listed above.

This is illustrated explicitly for the fixed plateau in Fig. 2(c). In Fig. 3(b), we show that there is close and increasing agreement (as a function of n) between the position of the sudden drop in $\langle \text{HW} \rangle$ and the position of the degeneracy point, for all of the problems considered.

- (4) *The scaling of the barrier height in the spin-coherent potential is positively correlated with the scaling of the minimum gap of the quantum Hamiltonian.* In Fig. 4, we see that as the barrier height increases, the inverse of the quantum minimum gap also increases.

Note that the Reichardt class is absent from the discussion above. The reason is that, for these problems, the barrier in the spin-coherent potential is very small, which makes its numerical detection difficult. Fortunately, we can make some analytical claims about this class of problems. By adapting Reichardt's proof (reviewed in Appendix A) that these problems have a constant minimum gap, we are able to prove that the barrier height in the spin-coherent potential for this class vanishes as $n \rightarrow \infty$. More precisely, we can show, for any perturbed Hamming weight problem that,

$$\begin{aligned}
 V_{SC}^{\text{pert}} - V_{SC}^{\text{unpert}} &= s \sum_{l < k < u} f(k) \binom{n}{k} p(\theta)^k [1 - p(\theta)]^{n-k} \\
 &= \mathcal{O}\left(h \frac{u-l}{\sqrt{l}}\right),
 \end{aligned} \tag{13}$$

where the unperturbed case refers to $h(n) = 0$ in Eq. (4). Recall that $h[(u-l)/\sqrt{l}] = o(1)$ for the Reichardt class. Thus, asymptotically, the spin-coherent potential for this class approaches the spin-coherent potential of the

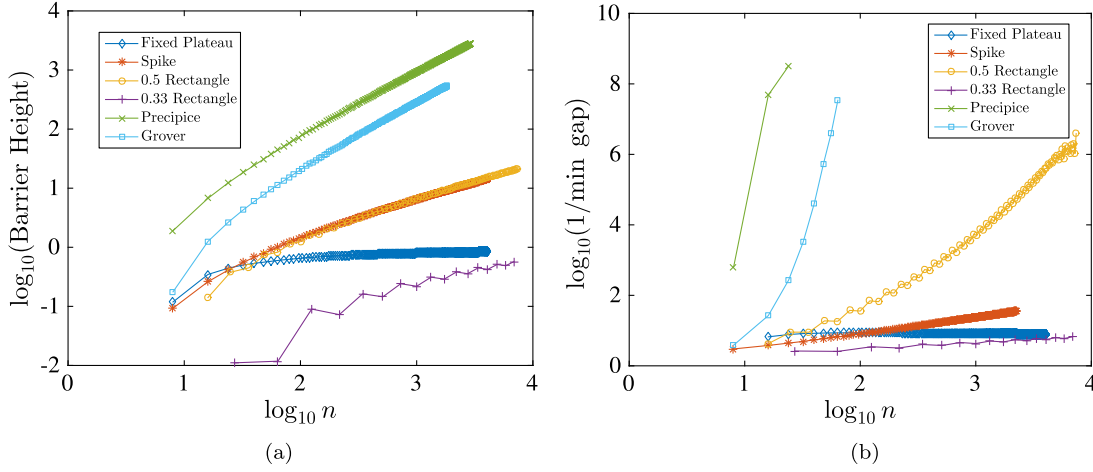


FIG. 4. (a) The height of the barrier between the two wells at the degeneracy point of the spin-coherent potential [as seen in Fig. 2(a)], as a function of n for the fixed plateau, the spike, the 0.5-rectangle, the 0.33-rectangle, the precipice, and Grover problems (log-log scale). (b) The inverse of the minimum gap as a function of n for the fixed plateau, the spike, the 0.5-rectangle, the 0.33-rectangle, the precipice, and Grover problems (log-log scale). The important thing to note is that the *order* of scaling is preserved in both plots; that is, the steeper the scaling of the barrier height, the steeper the scaling of the inverse minimum gap.

unperturbed Hamming weight problem. It is easy to check that the latter has a single minimum throughout the evolution, and hence no barriers.

Taken together, these observations indicate that the spin-coherent potential (not the cost function alone) is the appropriate potential with respect to which tunneling is to be understood for these problems.

IV. FIXED PLATEAU: PERFORMANCE OF ALGORITHMS

Having motivated the spin-coherent potential for understanding tunneling, we now exhaustively analyze the fixed plateau. We choose this problem because it forces us to confront some intuitions about the performance of certain algorithms. Considering the final cost function, the fixed plateau has neither local minima nor a barrier going from large to small $|x|$: it just has a long, flat section before the ground state at $|x| = 0$. This might suggest that it is easy for an algorithm such as SA, and is not a candidate for a quantum speedup. Moreover, given the absence of a barrier, one might suspect that the quantum evolution would not even involve multiqubit tunneling. In other words, given that the final cost function presents only an entropic obstacle to SA and tunneling typically needs an energetic obstacle, one might suspect that there is no tunneling present in this problem.

We dispel both of these intuitions and summarize our findings first. In the previous section, we provided evidence that tunneling is unambiguously present for this problem. The spin-coherent potential involves energy barriers, despite their absence in the final cost function, and the adiabatic quantum evolution is forced to tunnel in order to follow the ground state. By numerically solving the Schrödinger equation, we find that AQA needs a time of

$\mathcal{O}(n^{0.5})$ in order to reach a given success probability (see Sec. IV A). Therefore, the adiabatic algorithm, via tunneling, is able to solve this problem efficiently. Moreover, we see that an entropic barrier in the final cost function can get translated into an energetic barrier during the quantum evolution.

Turning to SA, an algorithm that performs a local stochastic search on the final cost function, we prove that simulated annealing with single-spin updates takes time $\mathcal{O}(n^{u-l-1}) = \mathcal{O}(n^{\text{plateau width}})$ to find the ground state (see Sec. IV B). This result is due to the fact that a random walker on the plateau has no preferred direction and becomes trapped there. More precisely, the probability of a leftward transition while on the plateau is proportional to the probability of flipping one of a constant number of bits (given by the Hamming weight) out of n , which scales as $\sim 1/n$ if $l, u = \mathcal{O}(1)$. And since the walker needs to make as many consecutive leftward transitions as the width of the plateau in order to fall off the plateau, the time taken for this to happen scales as $\mathcal{O}(n^{\text{plateau width}})$. Consequently, we obtain a polynomial speedup of AQA over SA that can be made as large as desired. Therefore, using the fixed plateau, we are able to demonstrate that a quantum speedup over SA is possible via tunneling in the adiabatic regime.

However, is the adiabatic evolution optimal? In order to find the optimal evolution time, we employ the optimal time to solution (TTS_{opt}), a metric that is commonly used in benchmarking studies [24] (also see Appendix C). It is defined as the minimum total time such that the ground state is observed at least once with desired probability p_d :

$$\text{TTS}_{\text{opt}} = \min_{t_f > 0} \left\{ t_f \frac{\ln(1 - p_d)}{\ln[1 - p_{\text{GS}}(t_f)]} \right\}, \quad (14)$$

where t_f is the duration (in QA) or the number of single-spin updates (in SA) of a single run of the algorithm, and $p_{\text{GS}}(t_f)$ is the probability of finding the ground state in a single such run. The use of TTS_{opt} allows for the possibility that multiple short runs of the evolution, each lasting an optimal annealing time $(t_f)_{\text{opt}}$, result in a better scaling than a single long (adiabatic) run with an unoptimized t_f . The quantum evolution that gives the optimal annealing time relative to this cost function is actually DQA, with an asymptotic scaling of $\mathcal{O}(1)$. Importantly, this diabatic evolution does not contain any of the signatures of tunneling discussed in the previous section. Therefore, for the fixed plateau, tunneling does not give rise to the optimal quantum performance.

Motivated by the fact that the optimal quantum evolution involves no multiqubit tunneling, we consider spin-vector dynamics [25] (see also Refs. [26,27]), a model that evolves according to the spin-coherent potential in Eq. (12). SVD can be derived as the saddle-point approximation to the path-integral formulation of QA in the spin-coherent basis [27]. The SVD equations are equivalent to the Ehrenfest equations for the magnetization under the assumption that the density matrix is a product state; i.e., $\rho = \otimes_{i=1}^n \rho_i$, where ρ_i denotes the state of the i th qubit. This algorithm is useful since it is derived under the assumption of continuity of the angles (θ, ϕ) , so tunneling, which here would amount to a discrete jump in the angles, is absent.

We also consider a quantum Monte Carlo based algorithm, often called simulated quantum annealing (SQA) [28,29]. We show that SQA has a scaling that is better than SA's. Indeed, this is consistent with the fact that SQA thermalizes not just relative to the final cost function, but also during the evolution.

We provide further details of our implementations of each of these algorithms in Appendix D. We now turn to each of the algorithms individually and detail their performance for the fixed plateau problem.

A. Adiabatic dynamics

In order to study the scaling of adiabatic dynamics, we consider the minimum time τ_0 required to reach the ground state with some probability p_0 , where we choose p_0 to ensure that we are exploring a regime close to adiabaticity for QA. We call this benchmark metric the “threshold criterion,” and set $p_0 = 0.9$. As shown in Fig. 5, we observe a scaling for AQA that is $\sim n^{0.5}$. As is to be expected given that the tunneling for the fixed plateau problem is controlled by the width of the plateau, which is constant (does not scale with n), we find that τ_0 scales in the same way for the fixed plateau and the plain Hamming weight problems (see Appendix A). This suggests that the dominant contribution to the scaling at large n is not the time associated with tunneling but rather the time associated with the plain Hamming weight problem.

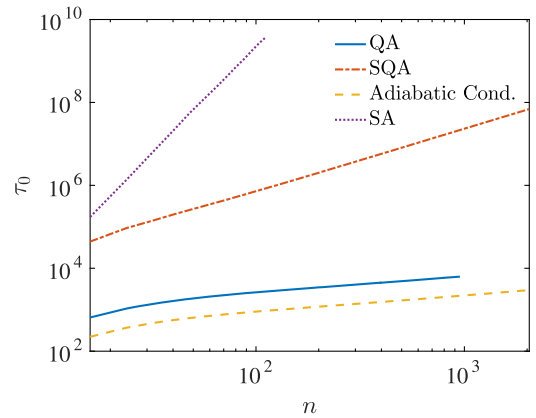


FIG. 5. Performance of different algorithms for the fixed plateau problem with $l = 0$ and $u = 6$. Shown is a log-log plot of the scaling of the time to reach a threshold success probability of 0.9, as a function of system size n for AQA, SQA ($\beta = 30$, $N_\tau = 64$) and SA ($\beta_f = 20$). The time for SQA and SA is measured in single-spin updates (for SQA this is N_τ times the number of sweeps times the number of spins, whereas for SA this is the number of sweeps times the number of spins), where both are operated in “solver” mode as described in Appendix D. Also shown is the scaling of the numerator of the adiabatic condition as defined in Eq. (15). The scaling for AQA and the adiabatic condition extracted by a fit using $n \gtrsim 10^2$ is approximately $n^{0.44}$. However, the true asymptotic scaling is likely to be $\sim n^{0.5}$ since the scaling for the fixed plateau problem is clearly lower bounded by the plain Hamming weight problem, for which we have verified $\tau_0 \sim n^{0.5}$ (see Appendix A), and we expect the effect of the plateau to become negligible in the large n limit. SQA scales more favorably ($\sim n^{1.5}$) than SA ($\sim n^5$). We have checked that the scaling of SQA does not change even if we double the number of Trotter slices N_τ and keep the temperature $1/\beta$ fixed.

As also shown in Fig. 5, we find that the textbook adiabatic criterion [30] given by

$$t_f \gtrsim \max_{s \in [0,1]} \frac{|\langle \epsilon_0(s) | \partial_s H(s) | \epsilon_1(s) \rangle|}{\text{gap}(s)^2} \quad (15)$$

serves as an excellent proxy for the scaling of AQA [31]. The scaling of AQA is matched by the scaling of the numerator of the adiabatic condition, which is explained by the fact that we find a constant minimum gap for the case $l, u = \mathcal{O}(1)$. This numerator turns out to be well approximated in our case by the matrix element of $H(s)$ between the ground and first excited states, leading to $t_f \sim n^{0.5}$ in the adiabatic limit. Note that calculating this matrix element can easily be done for arbitrarily large systems, and is hence much easier to check directly than the scaling of AQA.

B. Simulated annealing using random spin selection

We consider a version of SA with random spin selection as the rule that generates candidates for Metropolis updates. Our main motivation is to understand the behavior of a

local, stochastic search algorithm that has access only to the final cost function. We note that our analysis below is general for any plateau problem, and is not limited to the fixed plateau or the moving plateau.

If we pick a bit string at random, then for large n we start with very high probability at a bit string with Hamming weight close to $n/2$. The plateau may be to the left or to the right of $n/2$; if the plateau is to the right, then the random walker is unlikely to encounter it and can quickly descend to the ground state. Thus, the more interesting case is when the random walker arrives at the plateau from the right. We proceed to analyze these two cases separately.

1. Walker starts to the right of the plateau

In this case, how much time would it take, typically, for the walker to fall off the left edge? It is intuitively clear that traversing the plateau will be the dominant contribution to the time taken to reach the ground state, as after that the random walker can easily walk down the potential. We show below (for the walker that starts to the left of the plateau) that this time can be at most $\mathcal{O}(n^2)$ if $\beta = \Omega(\log n)$.

To evaluate the time to fall off the plateau, note that the perturbation is applied on strings of Hamming weight $l+1, l+2, \dots, u-1$, so the width of the plateau is $w \equiv u-l-1$. Consider a random walk on a line of $w+1$ nodes labeled $0, 1, \dots, w$. Node i represents the set of bit strings with Hamming weight $l+i$, with $0 \leq i \leq w$. We may assume that the random walker starts at node w . Only nearest-neighbor moves are allowed and the walk terminates if the walker reaches node 0 .

Our analysis provides a lower bound on the actual time to fall off the left edge, because in the actual PHWO problem one can also go back up the slope on the right, and in addition, we disallow transitions from strings of Hamming weight l to $l+1$. This is justified because the Metropolis rule exponentially (in β) suppresses these transitions.

The transition probabilities $p_{i \rightarrow j}$ for this problem can be written as a $(w+1) \times (w+1)$ row-stochastic matrix $p_{ij} = p_{i \rightarrow j}$. Here, p is a tridiagonal matrix with zeros on the diagonal, except at p_{00} and p_{ww} . Consider $1 \leq i \leq w-1$. If the walker is at node i , then the transition to node $i+1$ (which has Hamming weight $l+i+1$) occurs with probability $[n-(l+i)/n]$ (the chance that the bit picked had the value 0). Similarly, for $1 \leq i \leq w$, the Hamming weight will decrease to $l+i-1$ with probability $[(l+i)/n]$ (the chance that the bit picked had the value 1). Combining this with the fact that a walker at node 0 stays put, we can write

$$b_i \equiv p_{i \rightarrow i} = \begin{cases} 1 & \text{if } i = 0 \\ 0 & \text{if } 1 \leq i \leq (w-1) \\ 1 - \frac{l+w}{n} & \text{if } i = w, \end{cases} \quad (16a)$$

$$c_i \equiv p_{i-1 \rightarrow i} = \begin{cases} 0 & \text{if } i = 1 \\ 1 - \frac{l+i-1}{n} & \text{if } i = 2, \dots, w, \end{cases} \quad (16b)$$

$$a_i \equiv p_{i \rightarrow i-1} = \frac{l+i}{n} \quad \text{if } i = 1, 2, \dots, w. \quad (16c)$$

Let $X(t)$ be the position of the random walker at time step t . The random variable measuring the number of steps the random walker starting from node r would need to take to reach node s for the first time is

$$\tau_{r,s} \equiv \min\{t > 0 : X(t) = s, X(t-1) \neq s | X(0) = r\}. \quad (17)$$

The quantity we are after is $\mathbb{E}\tau_{w,0}$, the expectation value of the random variable $\tau_{w,0}$, i.e., the mean time taken by the random walker to fall off the plateau. Since only nearest-neighbor moves are allowed, we have

$$\mathbb{E}\tau_{w,0} = \sum_{r=1}^w \mathbb{E}\tau_{r,r-1}. \quad (18)$$

Stefanov [35] (see also Ref. [36]) has shown that

$$\mathbb{E}\tau_{r,r-1} = \frac{1}{a_r} \left(1 + \sum_{s=r+1}^w \prod_{t=r+1}^s \frac{c_t}{a_t} \right), \quad (19)$$

where $c_{w+1} \equiv 0$. Evaluating the sum term by term, we obtain

$$\mathbb{E}\tau_{w,w-1} = \frac{n}{l+w}, \quad (20a)$$

⋮

$$\mathbb{E}\tau_{w-k,w-k-1} = \frac{n}{l+w-k} \left[1 + \frac{n-(l+w-k)}{l+w-(k-1)} + \dots + \frac{n-(l+w-k)}{l+w-(k-1)} \times \dots \times \frac{n-(l+w-2)}{l+w-1} \times \frac{n-(l+w-1)}{l+w} \right]. \quad (20b)$$

Now consider the following cases:

- (1) *Fixed plateau*, $l, u = \mathcal{O}(1)$.—Here, using the fact that $k = \mathcal{O}(w) = \mathcal{O}(1)$, we conclude that $\mathbb{E}\tau_{w-k,w-k-1} = \mathcal{O}(n^{k+1})$. Since the leading-order term is $\mathbb{E}\tau_{w-(w-1),w-w} = \mathbb{E}\tau_{1,0}$, the time to fall off the plateau is $\mathcal{O}(n^w) = \mathcal{O}(n^{u-l-1})$. This result about SA's performance is confirmed numerically in Fig. 5.
- (2) In order for Reichardt's bound (see Appendix A) to give a constant lower bound to the quantum problem, we need $u = l + o(l^{1/4})$. Since at most we can

have $l = \mathcal{O}(n)$, we can conclude $\mathbb{E}\tau_{w-k, w-k-1} = \mathcal{O}(n/l)^{k+1}$. Therefore, the time to fall off becomes $\mathbb{E}\tau_{w,0} = \mathcal{O}[w(n/l)^w]$.

- (i) *Moving plateau.*—If $l = \Theta(n)$ and $w = \mathcal{O}(1)$, we can see that $\mathbb{E}\tau_{w,0} = \mathcal{O}(1)$, which is a constant time scaling.
- (ii) *Moving plateau with changing width.*—If $l = \Theta(n)$ and $w = \mathcal{O}(n^a)$, where $0 < a < 1/4$, then $\mathbb{E}\tau_{w,0} = \mathcal{O}(n^a \mathcal{O}(1)^{n^a})$, which is superpolynomial.
- (iii) *Most general plateau in the Reichardt class.*—More generally, if $l = \mathcal{O}(n^b)$, with $b \leq 1$ and $w = \mathcal{O}(n^a)$, where $0 \leq a < b/4$, then we get the scaling $\mathbb{E}\tau_{w,0} = \mathcal{O}(n^a \mathcal{O}(n^{1-b})^{n^a})$.

2. Walker starts to the left of the plateau

Note that this case is equivalent to the unperturbed Hamming weight problem, which is a straightforward gradient descent problem. We may therefore consider a simple fixed temperature version of SA (i.e., the standard Metropolis algorithm). We show that the performance of SA on this problem provides an upper bound of $\mathcal{O}(n^2)$ on the time for a random walker to arrive at the plateau and on the time for a random walker to reach the ground state after descending from the plateau. Moreover, our analysis provides a lower bound of $\mathcal{O}(n \log n)$ on the efficiency of such algorithms.

For this problem, the transition probabilities are

$$c_i \equiv p_{i-1 \rightarrow i} = \frac{n-i+1}{n} e^{-\beta}, \quad (21a)$$

$$a_i \equiv p_{i \rightarrow i-1} = \frac{i}{n}, \quad (21b)$$

with $i = 1, 2, \dots, n$ denoting strings of Hamming weight i , and β is the inverse temperature. Using the Stefanov formula Eq. (19), we can write (after much simplification)

$$\mathbb{E}\tau_{n-k, n-k-1} = \frac{n}{n-k} \binom{n}{k}^{-1} \sum_{l=0}^k e^{-l\beta} \binom{n}{k-l}. \quad (22)$$

We bound

$$\mathbb{E}\tau_{n,0} = \sum_{k=0}^{n-1} \frac{n}{n-k} \binom{n}{k}^{-1} \sum_{l=0}^k e^{-l\beta} \binom{n}{k-l}, \quad (23)$$

the expected time to reach the all-zeros string starting from the all-ones string. This is the worst-case scenario, as we are assuming that we are starting from the string farthest from the all-zeros string. Note again that if we start from a random spin configuration, then with overwhelming probability we will pick a string with Hamming weight close to $n/2$. Thus, most probably, $\mathbb{E}\tau_{n/2,0}$ will be the time to hit the ground state.

We first show that $\beta = \mathcal{O}(1)$ will lead to an exponential time to hit the ground state, irrespective of the walker's starting string. Toward that end,

$$\mathbb{E}\tau_{1,0} = \mathbb{E}\tau_{n-(n-1), n-n} \quad (24a)$$

$$= \sum_{l=0}^{n-1} e^{-l\beta} \binom{n}{n-1-l} \quad (24b)$$

$$= e^\beta [(e^{-\beta} + 1)^n - 1], \quad (24c)$$

which is clearly exponential in n if $\beta = \mathcal{O}(1)$.

Next, let $\beta(n) = \log n$; i.e., we decrease the temperature logarithmically in system size. In this case,

$$\mathbb{E}\tau_{1,0} = n \left[\left(1 + \frac{1}{n} \right)^n - 1 \right] \leq n(e-1) = \mathcal{O}(n). \quad (25)$$

Now it is intuitively clear that $\mathbb{E}\tau_{1,0} > \mathbb{E}\tau_{r,r-1}$ for all $r > 1$, which implies that $n\mathbb{E}\tau_{1,0} \geq \mathbb{E}\tau_{n,0}$. Thus, if $\beta = \log n$, then $\mathbb{E}\tau_{n,0} = \mathcal{O}(n^2)$ at worst.

To obtain a lower bound on the performance of the algorithm, we take $\beta \rightarrow \infty$. Thus, for each k in Eq. (23), only the $l=0$ term will survive. Hence,

$$\lim_{\beta \rightarrow \infty} \mathbb{E}\tau_{n,0} = \sum_{k=0}^{n-1} \frac{n}{n-k} = n \sum_{i=1}^n \frac{1}{i} \approx n(\log n + \gamma), \quad (26)$$

for large n , with γ being the Euler-Mascheroni constant. The scaling here is $\mathcal{O}(n \log n)$. This is the best possible performance for single-spin update SA with random spin selection on the plain Hamming weight problem. Therefore, if $\beta = \Omega(\log n)$, the scaling will be between $\mathcal{O}(n \log n)$ and $\mathcal{O}(n^2)$. Of course, this cost needs to be added to the time taken for the walker starting to the right of the plateau.

Two clarifications are in order regarding the comparison between our theoretical bound on SA's performance and the associated numerical simulations we have presented. First, while Fig. 5 displays the time to cross a threshold probability, our theoretical bound of $\mathcal{O}(n^{u-l-1})$ is on the expected time for the random walker to hit the ground state [Eq. (18)]. However, we find that both metrics show identical scaling. Second, while the SA data in Fig. 5 are generated using sequential spin updates, the theoretical bound assumes random spin updates (see Appendix D 1 for more details on the update schemes). However, we find that the asymptotic scaling for both cases is nearly identical in the long-time regime, and thus plot only the former.

Finally, we remark that SA as we consider here has only oracle access to the cost function and thus cannot exploit the symmetry of the cost function. This is the fair way to compare its performance to QA because QA too has only oracle access to the final Hamiltonian. QA's dynamics are

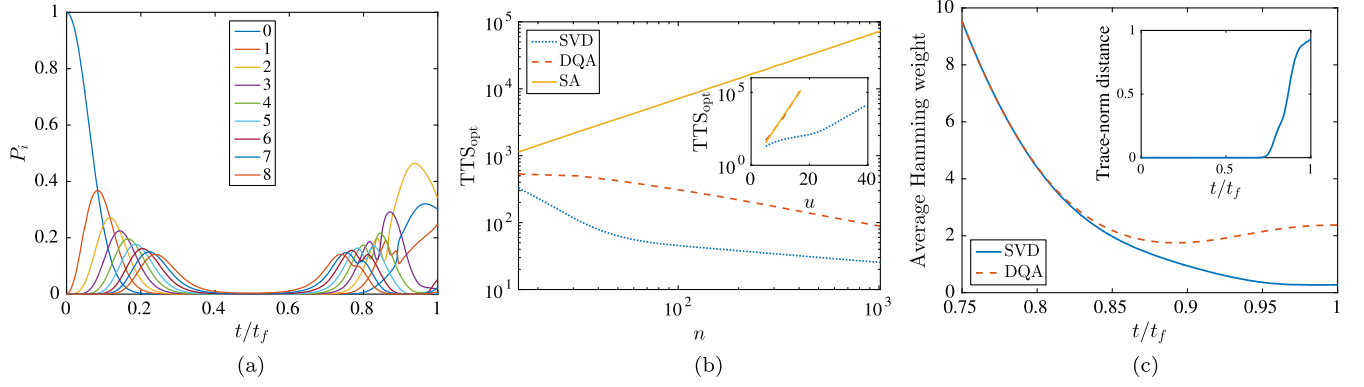


FIG. 6. Diabatic QA versus SA and SVD for the fixed plateau problem with $l = 0$. (a) Population P_i in the i th energy eigenstate along the diabatic QA evolution at the optimal TTS for $n = 512$ and $u = 6$. Excited states are quickly populated at the expense of the ground state. By $t/t_f = 0.5$ the entire population is outside the lowest nine eigenstates. In the second half of the evolution the energy eigenstates are repopulated in order. This kind of dynamics occurs due to a lining up of avoided-level crossings as seen in Fig. 7. (b) Scaling of the optimal TTS with n for $u = 6$, with an optimized number of single-spin updates for SA, and equal $(t_f)_{\text{opt}}$ for DQA and SVD. SA scales as $\mathcal{O}(n)$, a consequence of performing sequential single-spin updates. DQA and SVD both approach $\mathcal{O}(1)$ scaling as n increases. Here, we set $p_d = 0.7$ in Eq. (14), in order to be able to observe the saturation of SVD's TTS to the point where a single run suffices, i.e., $\text{TTS}_{\text{opt}} = (t_f)_{\text{opt}}$. The conclusion is unchanged if we increase p_d : this moves the saturation point to larger n for both SVD and DQA, and we have checked that SVD always saturates before DQA. Inset: Scaling as a function of u for $n = 1008$. SVD is again seen to exhibit the best scaling, while for this value of n the scaling of DQA and SA is similar (DQA's scaling with n improves faster than SA's as a function of n , at constant u). (c) $\langle \text{HW} \rangle$ of the QA wave function and the SVD state (defined as the product of identical spin-coherent states) for $n = 512$ and $u = 6$. The behavior of the two is identical up to $t/t_f \approx 0.8$, when they begin to differ significantly, but neither displays any of the sharp changes observed in Fig. 2(c) for the instantaneous ground state. Inset: The trace-norm distance between the DQA and SVD states, showing that they remain almost indistinguishable until $t/t_f \approx 0.8$.

automatically restricted to the symmetric subspace. This is *not* an unfair advantage for QA because the user of either algorithm—QA or SA—need not be aware of the symmetry of the problem.

C. Optimal QA via diabatic transitions

Having established that for the fixed plateau AQA enjoys a quantum speedup over local search algorithms such as SA via tunneling, we are motivated to ask, Is tunneling *necessary* to achieve a quantum speedup on these problems? In order to answer this question, we demonstrate using the optimal TTS criterion defined in Eq. (14) that the optimal annealing time for QA is far from adiabatic. Instead, as shown in Fig. 6(a), the optimal TTS for QA is such that the system leaves the instantaneous ground state for most of the evolution and returns to the ground state only towards the end. The cascade down to the ground state is mediated by a sequence of avoided energy level crossings, as shown in Fig. 7. We consider this a diabatic form of QA (DQA) and call this mechanism through which DQA achieves a speedup a *diabatic cascade*.

As n increases for fixed u , repopulation of the ground state improves for fixed $(t_f)_{\text{opt}}$, hence causing TTS_{opt} to decrease with n , as shown Fig. 6(b), until it saturates to a constant at the lowest possible value, corresponding to a single run at $(t_f)_{\text{opt}}$. At this point, the problem is solved in constant time $(t_f)_{\text{opt}}$, compared to the $\sim \mathcal{O}(n^{0.5})$ scaling of

the adiabatic regime. Moreover, as shown in Fig. 6(c), there are no sharp changes in $\langle \text{HW} \rangle$, suggesting that the non-adiabatic dynamics do not entail multiqubit tunneling events, unlike the adiabatic case. Thus, this establishes that we may have speedups in QA that do not involve multiqubit tunneling.

One may worry that for this diabatic evolution to be successful the optimal annealing time may need to be very finely tuned. We address this concern in Appendix E, where we show that if ϵ is the precision desired in p_{GS} , we need only have a precision of $\text{polylog}(1/\epsilon)$ in setting t_f , which means that the diabatic speedup is robust.

Figure 8 shows that the speedup of DQA and SVD over AQA exists for three other PHWO problems: the moving plateau, the spike, and the 0.5-rectangle problems. Importantly, DQA and SVD have an exponential speedup over AQA for the 0.5-rectangle problem. We do not observe a diabatic speedup for the precipice or Grover problems.

One might ask why it is that we see tunneling with respect to the spin-coherent potential for AQA but not for DQA. The reason for this is simple: In AQA, the adiabatic dynamics keep the state of the quantum evolution close to the instantaneous ground state. Thus, the wave function is localized around the instantaneous global minimum in the spin-coherent potential. Thus, in order to follow the true ground state, the system is forced to tunnel. On the other hand, in DQA there is no restriction that the dynamics need to follow the instantaneous ground state; indeed, the




FIG. 7. The eigenenergy spectrum along the evolution for the fixed plateau with $n = 512$, $l = 0$, and $u = 16$. Note the sequence of avoided-level crossings that unmistakably line up in the spectrum to reach the ground state. This is the pathway through which DQA is able to achieve a speedup over AQA.

dynamics extensively populate higher energy levels. Thus, the problem is solved without needing to tunnel.

D. Spin-vector dynamics

Given the absence of tunneling in the time-optimal quantum evolution, we are motivated to consider the behavior of spin-vector dynamics, which arise in a

semiclassical limit (see Appendix D 3 for an overview of this algorithm). As we show in Fig. 6(b), the scaling of SVD's optimal TTS also saturates to a constant time, i.e., $(t_f)_{\text{opt}}$. Moreover, it reaches this value earlier (as a function of problem size n) than DQA, thus outperforming DQA for small problem sizes, while for large enough n , both achieve $\mathcal{O}(1)$ scaling. As seen in the inset, SVD's advantage persists as a function of u at constant n .

The dynamics of DQA are well approximated by SVD until close to the end of the evolution, as shown in Fig. 6(c): the trace-norm distance between the instantaneous states of DQA and SVD is almost zero until $t/t_f \approx 0.8$, after which the states start to diverge. This suggests that SVD is able to replicate the DQA dynamics up to this point, and deviates only because it is more successful at repopulating the ground state than DQA.

In Fig. 8, we show that SVD's speedup over AQA is replicated for the spike, moving plateau, and 0.5-rectangle problems as well. Remarkably, while the 0.5-rectangle problem has an exponentially small gap [see Eq. (10) and Fig. 4(b)], SVD and DQA both achieve $\mathcal{O}(1)$ scaling, and hence the diabatic cascades provides an exponential speedup relative to AQA.

It is important to note that SVD is ineffective if one desires to simulate the adiabatic evolution. In the absence of unitary dynamics (which allow for tunneling) or thermal

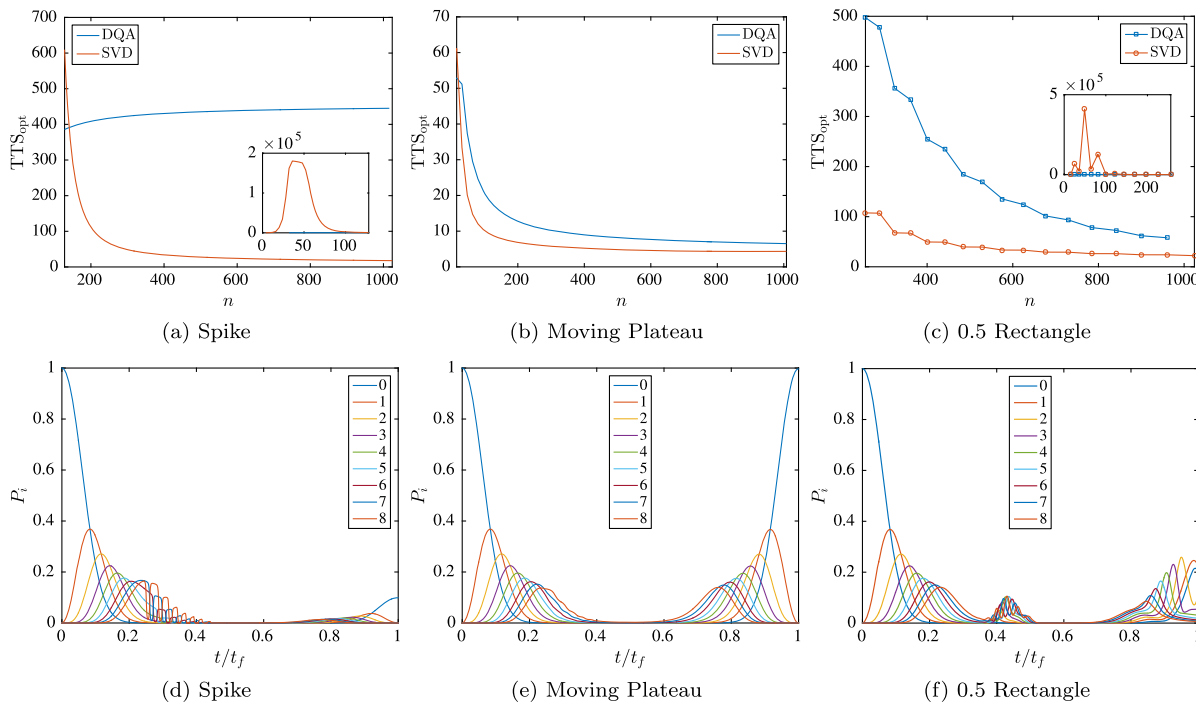


FIG. 8. (a)–(c) The optimal TTS for the spike, moving plateau, and 0.5-rectangle problems, respectively. Inset for (a) and (c): The optimal TTS for small problem sizes, where we observe SVD at first scaling poorly. However, as n grows, this difficulty vanishes and it quickly outperforms DQA. (d)–(f) Population P_i in the i th energy eigenstate along the diabatic QA evolution at the optimal TTS. We observe similar diabatic transitions for these problems (shown are the cases with $n = 512$ and $t_f = 9.85$ for the spike, $n = 512$ and $t_f = 10$ for the moving plateau, and $n = 529$ and $t_f = 9.8$ for the 0.5-rectangle problems) as we observed for the fixed plateau problem [Fig. 6].

activation (to thermally hop over the barrier), SVD gets trapped behind the barrier that forms in the semiclassical potential separating the two degenerate minima [see Fig. 2(a)] and is unable to reach the new global minimum. In this sense, SVD does not enjoy the guarantee provided by the quantum adiabatic theorem for the unitary evolution [32–34], that for sufficiently long t_f dictated by the adiabatic condition, the ground state can be reached with any desired probability.

Likewise, it is important to keep in mind the distinction between a classical algorithm being able to match, or sometimes outperform, a quantum algorithm (as SVD does here) and the classical algorithm approximating the evolution or instantiating the physics of the quantum algorithm (as SVD fails to do here). Indeed, in both the diabatic and adiabatic regimes, SVD provides a poor approximation to the instantaneous quantum state. For example, in the diabatic regime, it is clear from Fig. 6(c) that the trace-norm distance between the instantaneous SVD state and the instantaneous quantum state starts to increase significantly for $s \gtrsim 0.8$. In the same spirit, consider the instantaneous semiclassical ground state, i.e., the spin-coherent state evaluated at the minimum of the spin-coherent potential, which may be suspected to provide a good approximation to the instantaneous quantum ground state, but does not, as shown in Fig. 2(b). Thus, the unentangled semiclassical ground state also fails to provide a good approximation to the quantum ground state.

E. Simulated quantum annealing

Simulated quantum annealing is a quantum Monte Carlo algorithm performed along the annealing schedule (see Appendix D 4 for further details). It is often used as a benchmark against which QA is compared (however, see Ref. [4] for caveats). SQA scales better than SA for the fixed plateau problem using the threshold criterion (see Fig. 5). In order to understand why SQA enjoys an advantage over SA using this benchmark metric, it is useful to study the behavior of the state of SQA along the annealing schedule. We show the behavior of $\langle HW \rangle$ for SQA in Fig. 9, where we observe that SQA at the optimal number of sweeps (sweeps are defined in Appendix D 1; the case of 1500 sweeps shown in Fig. 9) does *not* follow the instantaneous ground state. Instead, it reaches the threshold success probability by thermally relaxing to the ground state *after* the minimum gap point (and tunneling event) of the quantum Hamiltonian. Therefore, SQA’s advantage over SA stems from the fact that it thermalizes in a different energy landscape than SA.

We also contrast the behavior of SQA and AQA using the threshold criterion. While SQA is able to follow the instantaneous ground state for a sufficiently large number of sweeps and thus mimic the tunneling of AQA (see Fig. 9), this is not the optimal way for it to reach the threshold criterion. For a fixed threshold success probability, the process of thermal relaxation after the minimum




FIG. 9. The expectation value of the Hamming weight operator for the quantum ground state, SQA, and AQA for the fixed plateau problem with $n = 512$, $l = 0$, and $u = 6$ and annealing time chosen so as to reach a success probability of 0.9. The expectation value for SQA ($\beta = 30$, $N_\tau = 64$) at a given t/t_f is calculated by averaging over the Hamming weight of the N_τ imaginary-time states at that time and over 10^5 independent trials. The case of 1500 sweeps is the minimum number of sweeps required for SQA (in “annealer” mode) to reach the threshold ground-state probability of $p_0 = 0.9$, and similarly for the annealing time value of $t_f = 4931.16$ for AQA. While AQA is able to approximately follow the quantum ground state (i.e., the evolution is very close to being adiabatic), the optimal SQA evolution (i.e., that requires the fewest sweeps) for achieving the threshold criterion involves *not* following the ground state at the minimum gap point and instead thermally relaxing towards the ground state after this point. As shown using the higher N_{sw} values, only after increasing the number of sweeps by more than 2 orders of magnitude does SQA follow the instantaneous ground state closely.

gap point uses fewer sweeps (and hence is more efficient) than following the instantaneous ground state closely throughout the anneal [37]. This is in contrast to AQA, where tunneling is the only means for it to reach a high success probability and, nevertheless, is more efficient than SQA, as shown in Fig. 5.

We note that SQA’s threshold criterion advantage over SA does not carry over to the optimal TTS criterion. In fact, we find that using the optimal TTS criterion, SQA scales as $O(n^{1.5})$, while SA scales as $O(n)$, as shown in Fig. 6(b). The reason for the latter scaling is that the optimal number of sweeps for SA is 1, simply because there is a small but nonzero probability that in the first sweep all the 1’s are flipped to 0’s.

V. DISCUSSION

It is often assumed that the shape of the final cost function determines how hard it is for QA to solve the problem (in fact, this is partly the motivation for the spike problem in Ref. [12]), and that potentials with tall and thin barriers should be advantageous for AQA, since this is

where tunneling dominates over thermal hopping (see, e.g., Ref. [4], p. 215; Ref. [9], p. 1062; and Ref. [10], p. 226). It is then assumed that problems where the final potential has this feature are those for which there should be a quantum speedup. We have given several counterexamples to such claims, and have shown that tunneling is not necessary to achieve the optimal TTS. Instead, the optimal trajectory may use diabatic transitions to first scatter completely out of the ground state and return via a sequence of avoided-level crossings. That diabatic transitions can help speed up quantum algorithms has also been noted and advantageously exploited in Refs. [38–41]. Moreover, we show that the instantaneous semiclassical potential provides important insight into the role of tunneling, while the final cost function can be rather misleading in this regard.

While both adiabatic and diabatic QA outperform SA for the fixed plateau problem, the faster quantum diabatic algorithm is not better than the classical SVD algorithm for this problem. The PHWO problems due to Reichardt [16], which include problems very similar to the fixed plateau, have widely been considered an example where tunneling provides a quantum advantage; we show that this holds if one limits the comparison to SA, but that there is, in fact, no quantum speedup in the problem when one compares the quantum diabatic evolution (which outperforms adiabatic quantum annealing) to SVD.

These results of the diabatic optimal evolution extend beyond the plateau problems: even the spike problem studied in Ref. [12]—which is in some sense the antithesis of the plateau problem since it features a sharp spike at a single Hamming weight—also exhibits the diabatic-beats-adiabatic phenomenon, indicating that tunneling is not required to efficiently solve the problem. Thus, diabatic evolution, especially via diabatic cascades, is an important and relatively unexplored mechanism in quantum optimization that is different from tunneling. The fact that we observe a speedup relative to AQA for several problems, especially an exponential speedup for the 0.5-rectangle problem, motivates the search for algorithms exploiting this mechanism and may yield fruitful results. However, we also already know that diabatic cascades are not generic. For example, we found that this mechanism is absent in the Grover and precipice problems, even though the Grover problem is equivalent to a “giant” plateau problem.

In summary, our work provides a counterargument to the widely made claims that tunneling should be understood with respect to the final cost function, that speedups due to tunneling require tall and thin barriers, and that tunneling is needed for a quantum speedup in optimization problems. Which features of Hamiltonians of optimization problems favor diabatic or adiabatic algorithms remains an open question, as is the understanding of tunneling for non-permutation-symmetric problems.

We finish on a positive note for QA. We have given several examples where SVD outperforms QA, e.g., the




FIG. 10. The optimal TTS for the potential given in Eq. (27). QA outperforms SVD over the range of problem sizes we were able to check. The reason can be seen in the inset, which displays the ground-state probability for SVD and QA for different annealing times t_f , with $n = 512$. The optimal annealing time for SVD occurs at the first peak in its ground-state probability ($t_f \approx 8.98$), whereas the optimal annealing time for QA occurs at the much larger *second* peak in its ground-state probability ($t_f \approx 10.91$).

spike problem [12]. However, we make no claim that SVD will always have an advantage over QA. A simple and instructive example comes from the class of cost functions that are convex in Hamming weight space, which have a constant minimum gap [42]:

$$f(x) = \begin{cases} 2 & |x| = 0 \\ |x| & \text{otherwise.} \end{cases} \quad (27)$$

We observe similar diabatic transitions for this problem as for the fixed plateau (not shown), but find that DQA outperforms SVD, as shown in Fig. 10. This results because the optimal TTS for QA occurs at a slightly higher optimal annealing time, i.e., there is an advantage to evolving somewhat more slowly, though still far from adiabatically. Thus, this provides an example of a “limited” quantum speedup [24].

ACKNOWLEDGMENTS

Special thanks to Ben Reichardt for insightful conversations and for suggesting the plateau problem, and to Bill Kaminsky for inspiring talks [43,44]. We also thank Itay Hen, Joshua Job, Iman Marvian, Milad Marvian, and Rolando Somma for useful comments. Additionally, we thank the anonymous referees for suggesting improvements to the paper. The computing resources were provided by the USC Center for High Performance Computing and Communications and by the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. This work was supported under ARO Grant No. W911NF-12-1-0523 and ARO MURI Grant No. W911NF-11-1-0268.

APPENDIX A: REVIEW OF THE HAMMING WEIGHT PROBLEM AND REICHARDT'S BOUND FOR PHWO PROBLEMS

Here, we closely follow Ref. [16].

1. Hamming weight problem

We review the analysis within QA of the minimization of the Hamming weight function $f_{\text{HW}}(x) = |x|$, which counts the number of 1's in the bit string x . This problem is of course trivial, and the analysis we give here is done in preparation for the perturbed problem.

For the adiabatic algorithm, we start with the driver Hamiltonian,

$$H_D = \frac{1}{2} \sum_{i=1}^n (\mathbb{1}_i - \sigma_i^x) = \sum_{i=1}^n |-\rangle_i \langle -|, \quad (\text{A1})$$

which has $|+\rangle^{\otimes n}$ as the ground state.

The final Hamiltonian for the cost function $f_{\text{HW}}(x)$ is

$$H_P = \frac{1}{2} \sum_{i=1}^n (\mathbb{1}_i - \sigma_i^z) = \sum_{i=1}^n |1\rangle_i \langle 1|, \quad (\text{A2})$$

which has $|0\rangle^{\otimes n}$ as the ground state.

We interpolate linearly between H_D and H_P :

$$H(s) = (1-s)H_D + sH_P; \quad s \in [0, 1] \quad (\text{A3})$$

$$= \sum_{i=1}^n \frac{1}{2} \begin{pmatrix} 1-s & -(1-s) \\ -(1-s) & 1-s \end{pmatrix}_i + \begin{pmatrix} 0 & 0 \\ 0 & s \end{pmatrix}_i \quad (\text{A4})$$

$$= \frac{1}{2} \sum_{i=1}^n [\mathbb{1} - (1-s)\sigma_i^x - s\sigma_i^z] \equiv H_i(s). \quad (\text{A5})$$

We note that $H_i(s)$ in Eq. (A5) is similar to a variant of the Landau-Zener (LZ) Hamiltonian with finite coupling duration [45,46], for which the Schrödinger equation has an analytical solution, except that there it is assumed that the σ^x term is constant and only the σ^z term has a (linear) time dependence over a finite interval. The analytical solution of the problem obtained in Ref. [45] is rather complicated, and for our purposes a simpler approach suffices.

Since there are no interactions between the qubits, the adiabatic problem can be solved exactly by diagonalizing the Hamiltonian acting on each qubit separately. For each term, we have the energy eigenvalues $E_{\pm}(s)$,

$$E_{\pm}(s) = \frac{1}{2} [1 \pm \Delta(s)]; \quad \Delta(s) \equiv \sqrt{1 - 2s + 2s^2}, \quad (\text{A6})$$

and associated eigenvectors,

$$|v_{\pm}(s)\rangle = \frac{1}{\sqrt{2\Delta(\Delta \mp s)}} [\mp(\Delta \mp s)|0\rangle + (1-s)|1\rangle]. \quad (\text{A7})$$

The ground state of $H(s)$ is

$$|\psi_{\text{GS}}(s)\rangle = |v_{-}(s)\rangle^{\otimes n}. \quad (\text{A8})$$

The gap is given by

$$\text{gap}[H(s)] = H(s)|v_{+}(s)\rangle \otimes |v_{-}(s)\rangle^{\otimes(n-1)} - H(s)|v_{-}(s)\rangle^{\otimes n} \quad (\text{A9a})$$

$$= E_{+} + (n-1)E_{-} - nE_{-} \quad (\text{A9b})$$

$$= E_{+} - E_{-} \quad (\text{A9c})$$

$$= \Delta(s). \quad (\text{A9d})$$

The gap is minimized at $s = \frac{1}{2}$ with minimum value $\Delta(\frac{1}{2}) = (1/\sqrt{2})$. The minimum gap is independent of n and hence does not scale with problem size. Therefore, we can predict an adiabatic run time to be given by

$$t_f = \mathcal{O}\left(\frac{\|\partial_s H\|}{\Delta^2}\right) = \mathcal{O}(n), \quad (\text{A10})$$

where the n dependence is solely due to $\|\partial_s H\|$ (see Appendix D 2). However, this is actually a loose upper bound. We next provide separate numerical and analytical arguments to demonstrate that the actual scaling for AQA is $\mathcal{O}(n^{0.5})$.

a. Numerical argument

Suppose the adiabatic algorithm runs long enough so as to attain a desired success probability p_0 . Let this time be t_f . Using the fact that the quantum evolution of the plain Hamming weight problem is the evolution of n noninteracting qubits, we can express the global ground-state probability in terms of the ground-state probabilities of single qubits. So, if the single qubit ground-state probability for this run time is $p_{\text{GS}}(t_f)$, then we must have $p_0 = p_{\text{GS}}(t_f)^n$.

We find numerically (see Fig. 11) that $p_{\text{GS}}(t_f)$ has an envelope that is excellently approximated by

$$p_{\text{GS}}(t_f) = 1 - \frac{1}{t_f^2} + \mathcal{O}(t_f^{-3}), \quad (\text{A11})$$

for sufficiently large t_f . We therefore can write

$$\ln p_0 = n \ln p_{\text{GS}}(t_f) \approx n \ln \left(1 - \frac{1}{t_f^2}\right), \quad (\text{A12})$$

and upon expanding the \ln , we extract a tighter scaling for our adiabatic time,




FIG. 11. Ground-state probability for a single qubit for different total time t_f , evolving under the plain Hamming weight Hamiltonian in Eq. (A5).

$$t_f = \mathcal{O}(n^{1/2}). \quad (\text{A13})$$

b. Analytical argument

Here, we invoke a result due to Boixo and Somma [47]. This result states the following.

Theorem 1 (Ref. [47]).—To adiabatically prepare a final eigenstate using a Hamiltonian evolution $H(s)$ requires time that scales at least as $\mathcal{O}(L/\Delta)$. Here, L is the eigenpath length,

$$L \equiv \int_0^1 \|\partial_s \psi(s)\| ds, \quad (\text{A14})$$

where $|\psi(s)\rangle$ is the eigenpath traversed to reach the final eigenstate.

We analytically compute L for the ground-state path in the plain Hamming weight problem, and show that it scales as $\mathcal{O}(\sqrt{n})$. Since we know that in this case $\Delta = \mathcal{O}(1)$, we conclude that the adiabatic algorithm will require at least $\mathcal{O}(\sqrt{n})$ time.

Recall that the instantaneous ground state is [Eq. (A8)] $|\psi_{\text{GS}}(s)\rangle = \bigotimes_{i=1}^n |v_-^i(s)\rangle$, where $|v_-^i(s)\rangle = \sqrt{1-q(s)}|0\rangle_i + \sqrt{q(s)}|1\rangle_i$, with [Eq. (A7)],

$$q(s) = \frac{(1-s)^2}{2\Delta(\Delta+s)}. \quad (\text{A15})$$

Differentiating,

$$\frac{d}{ds} |\psi_{\text{GS}}(s)\rangle = \sum_{i=1}^n \left(\bigotimes_{j \neq i} |v_-^j(s)\rangle \otimes \frac{d}{ds} |v_-^i(s)\rangle \right), \quad (\text{A16})$$

so that

$$\|\partial_s \psi_{\text{GS}}(s)\|^2 \equiv \langle \partial_s \psi_{\text{GS}}(s) | \partial_s \psi_{\text{GS}}(s) \rangle \quad (\text{A17})$$

$$\begin{aligned} &= n \left\| \frac{d}{ds} |v_-^i(s)\rangle \right\|^2 + n(n-1) \\ &\quad \times \langle v_-^i(s) | \frac{d}{ds} |v_-^i(s)\rangle|^2. \end{aligned} \quad (\text{A18})$$

The term $\|(d/ds)|v_-^i(s)\rangle\|$ does not have any scaling with n , and the second term vanishes because it is equal to $\frac{1}{2}(d/ds)\langle v_-^i(s) | v_-^i(s) \rangle = 0$, where we use the fact that $|v_-^i(s)\rangle$ is real valued and normalized. Thus, taking the square root on both sides and integrating from 0 to 1, we obtain the \sqrt{n} scaling of L .

If we desire to fix the constant in front of L , a straightforward calculation shows that

$$\int_0^1 \left\| \frac{d}{ds} |v_-^i(s)\rangle \right\| ds = \pi/4. \quad (\text{A19})$$

2. Reichardt's bound for PHWO problems

Here, we review Reichardt's derivation of the gap lower bound for general PHWO problems and provide additional details not found in the original proof [16].

We use the same initial Hamiltonian [Eq. (A1)] and linear interpolation schedule as before, $\tilde{H}(s) = (1-s)H_D + s\tilde{H}_P$, and choose the final Hamiltonian to be

$$\tilde{H}_P = \sum_{x \in \{0,1\}^n} \tilde{f}(x) |x\rangle\langle x|, \quad (\text{A20})$$

where

$$\tilde{f}(x) = \begin{cases} |x| + p(x) & l < |x| < u \\ |x| & \text{elsewhere,} \end{cases} \quad (\text{A21})$$

where $p(x) \geq 0$ is the perturbation. Note that here we do *not* assume that the perturbation $p(x)$ respects qubit permutation symmetry.

We wish to bound the minimum gap of $\tilde{H}(s)$. Unlike the Hamming weight problem $H(s)$, this problem is no longer noninteracting. Define

$$h_k \equiv \max_{|x|=k} p(x); \quad h \equiv \max_k h_k = \max_x p(x). \quad (\text{A22})$$

Lemma 1 (Ref. [16]).—Let $u = \mathcal{O}(l)$ and let $E_0(s)$ and $\tilde{E}_0(s)$ be the ground-state energies of $H(s)$ and $\tilde{H}(s)$, respectively. Then, $\tilde{E}_0(s) \leq E_0(s) + \mathcal{O}[h(u-l)/(\sqrt{l})]$.

Proof.—First note that

$$\tilde{H}(s) - H(s) = s \sum_{x: l < |x| < u} p(x) |x\rangle\langle x|. \quad (\text{A23})$$

Below, we suppress the s dependence of all the terms for notational simplicity. We know that $E_0 = \langle v_-^{\otimes n} | H | v_-^{\otimes n} \rangle$. Using this,

$$\langle \tilde{E}_0 | \tilde{H} | \tilde{E}_0 \rangle \leq \langle \psi | \tilde{H} | \psi \rangle \quad \forall |\psi\rangle \in \mathcal{H} \quad (\text{A24a})$$

$$\Rightarrow \tilde{E}_0 - E_0 \leq \langle v_-^{\otimes n} | \tilde{H} | v_-^{\otimes n} \rangle - E_0 \quad (\text{A24b})$$

$$\leq \langle v_-^{\otimes n} | \tilde{H} - H | v_-^{\otimes n} \rangle \quad (\text{A24c})$$

$$= s \sum_{x: l < |x| < u} p(x) |\langle v_{\underline{n}}^{\otimes n} | x \rangle|^2 \quad (\text{A24d})$$

$$= s \sum_{x: l < |x| < u} p(x) q^{|x|} (1-q)^{n-|x|} \quad (\text{A24e})$$

$$\leq \sum_{k: l < k < u} h_k \binom{n}{k} q^k (1-q)^{n-k}, \quad (\text{A24f})$$

where $\binom{n}{k}$ is the number of strings with Hamming weight k , we use the fact that if we measure in the computational basis, the probability of getting outcome x is $|\langle v_{\underline{n}}^{\otimes n} | x \rangle|^2 = q(s)^{|x|} [1 - q(s)]^{n-|x|}$, and $q(s)$ is given in Eq. (A15).

Consider the partial binomial sum (dropping the h_k 's),

$$\sum_{k: l < k < u} \binom{n}{k} q^k (1-q)^{n-k}. \quad (\text{A25})$$

Using the fact that the binomial is well approximated by the Gaussian in the large n limit [note that this approximation requires that $q(s)$ and $1 - q(s)$ not be too close to zero], we can write

$$\begin{aligned} \sum_{k: l < k < u} \binom{n}{k} q^k (1-q)^{n-k} &\approx \int_l^u d\xi \frac{1}{\sqrt{2\pi}\sigma} e^{-(\xi-\mu)^2/2\sigma^2} \\ &= \frac{1}{\sigma} \int_l^u d\xi \phi\left(\frac{\xi-\mu}{\sigma}\right) \\ &= \int_{(l-\mu)/\sigma}^{(u-\mu)/\sigma} dt \phi(t), \end{aligned} \quad (\text{A26})$$

where $\mu \equiv nq$, $\sigma \equiv \sqrt{nq(1-q)}$, and $\phi(t) \equiv e^{-t^2/2}/\sqrt{2\pi}$. Note that σ and μ depend on n , and also on s via $q(s)$. The parameters l and u are specified by the problem Hamiltonian, and are therefore allowed to depend on n as long as $l(n) < u(n) < n$ is satisfied for all n .

Let us define

$$B[s, n, l(n), u(n)] \equiv \int_{[l(n)-\mu(n,s)]/\sigma(n,s)}^{[u(n)-\mu(n,s)]/\sigma(n,s)} dt \frac{e^{-t^2/2}}{\sqrt{2\pi}}. \quad (\text{A27})$$

We seek an upper bound on this function. We observe that $q(s)$ decreases monotonically from $\frac{1}{2}$ to 0 as s goes from 0 to 1. Thus, the mean of the Gaussian $\mu(n, s) = nq(s)$ decreases from $n/2$ to 0. Depending on the values of $l(n)$, $u(n)$, and $\mu(n, s)$, we thus have three possibilities: (i) $l(n) < \mu(n, s) < u(n)$, (ii) $\mu(n, s) < l(n) < u(n)$, and (iii) $l(n) < u(n) < \mu(n, s)$. Note that (ii) and (iii) are cases where the integral runs over the tails of the Gaussian and so the integral is exponentially small. We focus on (i), as this induces the maximum values of the integral. In this case, the lower limit of the integral Eq. (A27) is negative, while the upper limit is positive. Thus, the integral runs through the center of the standard Gaussian, and we can

upper bound the value of the integral by the area of the rectangle of width $\{[u(n) - l(n)]/\sigma(n, s)\}$ and height $1/\sqrt{2\pi}$. Hence,

$$B[s, n, l(n), u(n)] \leq \frac{1}{\sqrt{2\pi}} \frac{u(n) - l(n)}{\sigma(n, s)} \quad (\text{A28a})$$

$$= \frac{1}{\sqrt{2\pi}} \frac{u(n) - l(n)}{\sqrt{\mu(n)[1 - q(s)]}} \quad (\text{A28b})$$

$$\leq \frac{1}{\sqrt{2\pi}} \frac{u(n) - l(n)}{\sqrt{l(n)[1 - q(s)]}}, \quad (\text{A28c})$$

where we use the fact that $l(n) < \mu(n, s) = nq(s)$.

Thus, we obtain the bound:

$$\tilde{E}_0 - E_0 \leq \mathcal{O}\left(h \frac{u-l}{\sqrt{l}}\right). \quad (\text{A29})$$

Lemma 2 (Ref. [16]).—If $\tilde{H} - H$ is non-negative, then the spectrum of \tilde{H} lies above the spectrum of H . That is, $\tilde{E}_j \geq E_j$ for all j , where \tilde{E}_j and E_j denote the j th largest eigenvalue of \tilde{H} and H , respectively.

This can be proved by a straightforward application of the Courant-Fischer min-max theorem (see, for example, Ref. [48]).

Combining these lemmas results in the desired bound on the gap:

$$\text{gap}[\tilde{H}(s)] = \tilde{E}_1 - \tilde{E}_0 \quad (\text{A30a})$$

$$\geq E_1 - \tilde{E}_0 \quad (\text{A30b})$$

$$= E_1 - E_0 - (\tilde{E}_0 - E_0) \quad (\text{A30c})$$

$$\geq \Delta - \mathcal{O}\left(h \frac{u-l}{\sqrt{l}}\right), \quad (\text{A30d})$$

where in Eq. (A30b) we use Lemma 2 and in Eq. (A30d) we use Lemma 1.

Now, if we choose a parameter regime for the perturbation such that $h[(u-l)/\sqrt{l}] = o(1)$, then the perturbed problem maintains a constant gap. For example, if $l = \Theta(n)$ and $h(u-l) = \mathcal{O}(n^{1/2-\epsilon})$, for any $\epsilon > 0$, then the gap is constant as $n \rightarrow \infty$.

APPENDIX B: (NON-)LOCALITY OF PHWO PROBLEMS

Since the PHWO problems, including the plateau, are quantum oracle problems, they cannot generically be represented by a local Hamiltonian. For completeness, we prove this claim here and also show why the (plain) Hamming weight problem is 1-local.

Let r be a bit string of length n , i.e., $r \in \{0, 1\}^n$, and let

$$\sigma^r \equiv \sigma_1^{r_1} \otimes \sigma_2^{r_2} \otimes \dots \otimes \sigma_n^{r_n}, \quad (\text{B1})$$

with $\sigma_i^0 \equiv I_i$ and $\sigma_i^1 \equiv \sigma_i^z$. This forms an orthonormal basis for the vector space of diagonal Hamiltonians. Thus,

$$H_P = \sum_{r \in \{0,1\}^n} J_r \sigma^r, \quad (\text{B2})$$

with

$$J_r = \frac{1}{2^n} \text{Tr}(\sigma^r H_P) \quad (\text{B3a})$$

$$= \frac{1}{2^n} \sum_{x \in \{0,1\}^n} f(x) \langle x | \sigma^r | x \rangle \quad (\text{B3b})$$

$$= \frac{1}{2^n} \sum_{x \in \{0,1\}^n} f(x) (-1)^{x \cdot r}. \quad (\text{B3c})$$

Note that, generically, J_r will be nonzero for arbitrary-weight strings r , leading to $|r|$ -local terms in H_P , even as high as n -local.

For example, substituting the plateau Hamiltonian [Eq. (3)] into this, we obtain

$$J_r = \frac{1}{2^n} \left[\sum_{|x| \leq l \text{ and } |x| \geq u} |x| (-1)^{x \cdot r} + (u-1) \sum_{l < |x| < u} (-1)^{x \cdot r} \right]. \quad (\text{B4})$$

On the other hand, if $f(x) = |x|$ (i.e., in the absence of a perturbation), the Hamiltonian is only 1-local:

$$H_P = \sum_{x \in \{0,1\}^n} |x| |x\rangle \langle x| \quad (\text{B5a})$$

$$= \sum_{x_1=0}^1 \dots \sum_{x_n=0}^1 (x_1 + x_2 + \dots + x_n) |x_1\rangle \langle x_1| \otimes |x_2\rangle \langle x_2| \otimes \dots \otimes |x_n\rangle \langle x_n| \quad (\text{B5b})$$

$$= \sum_{k=1}^n (x_k |x_k\rangle \langle x_k|) \otimes \left(\sum_{j \neq k}^1 |x_j\rangle \langle x_j| \right) \quad (\text{B5c})$$

$$= \sum_{k=1}^n |1\rangle_k \langle 1| \otimes_{j \neq k} I_j = \sum_{k=1}^n |1\rangle_k \langle 1|. \quad (\text{B5d})$$

APPENDIX C: DERIVATION OF EQ. (14)

Equation (14) is easily derived as follows: the probability of successively failing k times to obtain the ground state is

$[1 - p_{\text{GS}}(t_f)]^k$, so the probability of succeeding at least once after k runs is $1 - [1 - p_{\text{GS}}(t_f)]^k$, which we set equal to the desired success probability p_d ; from here, one extracts the number of runs k and multiplies by t_f to get the time-to-solution TTS. Optimizing over t_f yields TTS_{opt} , which is natural for benchmarking purposes in the sense that it captures the trade-off between repeating the algorithm many times versus optimizing the probability of success in a single run.

A disadvantage of using the TTS_{opt} as a benchmarking metric is that extracting t_f^{opt} is, in general, difficult and is problem dependent. Algorithms that have an adiabatic regime (e.g., SA, QA) alleviate this issue by providing a theoretical guarantee that one will have a certain probability of success provided t_f is sufficiently large. Of course, the adiabatic regime is typically suboptimal.

In our analyses we extract TTS_{opt} by running the algorithm under consideration—i.e., SA, QA, SVD, or SQA—for a wide range of values of t_f , computing the $p_{\text{GS}}(t_f)$ for each value of t_f , using these to compute the $\text{TTS}(t_f)$ for that value of t_f , and, finally, choosing t_f that has the smallest TTS in the range considered. This method is not guaranteed to achieve the true optimum because the true t_f^{opt} may lie outside the range of values that we consider. Nevertheless, in the cases we study, the trends followed by the p_{GS} versus t_f suggest that we are extracting a value very close to the optimal TTS.

APPENDIX D: METHODS

1. Simulated annealing

SA is a general heuristic solver [2], whereby the system is initialized in a high-temperature state, i.e., in a random state, and the temperature is slowly lowered while undergoing Monte Carlo dynamics. Local updates are performed according to the Metropolis rule [49,50]: a spin is flipped and the change in energy ΔE associated with the spin flip is calculated. The flip is accepted with probability P_{Met} :

$$P_{\text{Met}} = \min\{1, \exp(-\beta \Delta E)\}, \quad (\text{D1})$$

where β is the current inverse temperature along the anneal. Note that there could be different schemes governing which spin is to be selected for the update. We consider two such schemes: random spin selection, where the next spin to be updated is selected at random, and sequential spin selection, where one runs through all of the n spins in a sequence. Random spin selection (including just updating nearest neighbors) satisfies detailed balance and thus is guaranteed to converge to the Boltzmann distribution. Sequential spin selection does not satisfy strict detailed balance (since the reverse move of sequentially updating in the reverse order never occurs), but it too converges to the Boltzmann distribution [51]. In sequential updating, a

“sweep” refers to all the spins having been updated once. In random spin selection, we define a sweep as the total number of spin updates divided by the total number of spins. When it is possible to parallelize the spin updates, the appropriate metric of time complexity is the number of sweeps N_{sw} , not the number of spin updates (they differ by a factor of n) [24]. However, in our problem, this parallelization is not possible and hence the appropriate metric is the number of spin updates, and this is what we plot in Fig. 6(b). After each sweep, the inverse temperature is incremented by an amount $\Delta\beta$ according to an annealing schedule, which we take to be linear; i.e., $\Delta\beta = (\beta_f - \beta_i)/(N_{\text{sw}} - 1)$.

We can use SA both as an *annealer* and as a *solver* [52]. In the former, the state at the end of the evolution is the output of the algorithm, and can be thought of as a method to sample from the Boltzmann distribution at a specified temperature. For the latter, we select the state with the lowest energy found along the entire anneal as the output of the algorithm, the better technique if one is interested in finding only the global minimum. We use the latter to maximize the performance of the algorithm.

2. Quantum annealing

Here, we consider the most common version of quantum annealing:

$$H(s) = (1-s) \sum_{i=1}^n \frac{1}{2} (1_i - \sigma_i^x) + s \sum_{x \in \{0,1\}^n} f(x) |x\rangle\langle x|, \quad (\text{D2})$$

where $s \equiv t/t_f$ is the dimensionless time parameter and t_f is the total anneal time. The initial state is taken to be $|+\rangle^{\otimes n}$, which is the ground state of $H(0)$.

The initial ground state and the total Hamiltonian are symmetric under qubit permutations [recall that $f(x) = f(|x|)$ for our class of problems]. It then follows that the time-evolved state, at any point in time, will also obey the same symmetry. Therefore, the evolution is restricted to the $(n+1)$ -dimensional symmetric subspace, a fact that we can take advantage of in our numerical simulations. This symmetric subspace is spanned by the Dicke states $|S, M\rangle$, with $S = n/2, M = -S, -S+1, \dots, S$, which satisfy

$$S^2 |S, M\rangle = S(S+1) |S, M\rangle, \quad (\text{D3a})$$

$$S^z |S, M\rangle = M |S, M\rangle, \quad (\text{D3b})$$

where $S^{x,y,z} \equiv \frac{1}{2} \sum_{i=1}^n \sigma_i^{x,y,z}$, $S^2 = (S^x)^2 + (S^y)^2 + (S^z)^2$. We can denote these states by

$$|w\rangle \equiv \left| \frac{n}{2}, M = \frac{n}{2} - w \right\rangle = \binom{n}{w}^{-1/2} \sum_{x: |x|=w} |x\rangle, \quad (\text{D4})$$

where, $w \in \{0, \dots, n\}$.

In this basis the Hamiltonian is tridiagonal, with the following matrix elements:

$$\begin{aligned} [H(s)]_{w,w+1} &= [H(s)]_{w+1,w} \\ &= -\frac{1}{2}(1-s) \sqrt{(n-w)(w+1)}, \end{aligned} \quad (\text{D5a})$$

$$[H(s)]_{w,w} = (1-s) \frac{n}{2} + sf(w). \quad (\text{D5b})$$

The Schrödinger equation with this Hamiltonian can be solved reliably using an adaptive Runge-Kutte Cash-Karp method [53] and the Dormand-Prince method [54] (both with orders 4 and 5).

If the quantum dynamics is run adiabatically the system remains close to the ground state during the evolution, and an appropriate version of the adiabatic theorem is satisfied. For evolutions with a nonzero spectral gap for all $s \in [0, 1]$, an adiabatic condition of the form

$$t_f \geq \text{const} \sup_{s \in [0,1]} \frac{\|\partial_s H(s)\|}{\text{gap}(s)^2} \quad (\text{D6})$$

is often claimed to be sufficient [55] [however, see the discussion after Eq. (21) in Ref. [32]]. In our case, $\|\partial_s H(s)\| = \|H(1) - H(0)\|$ is upper bounded by n ; since we are considering a constant gap, the adiabatic algorithm can scale at most linearly by condition Eq. (D6). This is true for the plateau problems.

In the main text we show that the following version of the adiabatic condition, known to hold in the absence of resonant transitions between energy levels [33], estimates the scaling we observe very well:

$$\max_{s \in [0,1]} \frac{|\langle \varepsilon_0(s) | \partial_s H(s) | \varepsilon_1(s) \rangle|}{\text{gap}(s)^2} \ll t_f, \quad (\text{D7})$$

where $\varepsilon_0(s)$ and $\varepsilon_1(s)$ are the instantaneous ground and excited states in the symmetric subspace, respectively.

The permutation symmetry is explicitly enforced only in our numerical simulations of the quantum evolution. Since, of course, we do not have quantum hardware that can implement the problems under consideration, we must explicitly enforce this symmetry in order to be able to perform numerical simulations at large problem sizes. Note that even if we were to simulate the quantum system without explicitly imposing this symmetry, the symmetry would be automatically preserved in the dynamics, and we would draw the same lessons about the performance of the quantum algorithm (but our classical simulations would quickly become intractable).

3. Spin-vector dynamics

Starting with the spin-coherent path-integral formulation of the quantum dynamics, we can obtain spin-vector dynamics as the saddle-point approximation (see, for example, Ref. [27], p. 10, or Refs. [25,26]). It can be interpreted as a semiclassical limit describing coherent single qubits interacting incoherently. In this sense, SVD is a well-motivated classical limit of the quantum evolution of QA. SVD describes the evolution of n unit-norm classical vectors under the Lagrangian (in units of $\hbar = 1$):

$$\mathcal{L} = i\langle\Omega(s)|\frac{d}{ds}|\Omega(s)\rangle - t_f\langle\Omega(s)|H(s)|\Omega(s)\rangle, \quad (\text{D8})$$

where $|\Omega(s)\rangle$ is a tensor product of n independent spin-coherent states [56]:

$$|\Omega(s)\rangle = \bigotimes_{i=1}^n \left[\cos\left(\frac{\theta_i(s)}{2}\right)|0\rangle_i + \sin\left(\frac{\theta_i(s)}{2}\right)e^{i\varphi_i(s)}|1\rangle_i \right]. \quad (\text{D9})$$

We can define an effective semiclassical potential associated with this Lagrangian:

$$\begin{aligned} V_{\text{SC}}(\{\theta_i\}, \{\varphi_i\}, s) &\equiv \langle\Omega(s)|H(s)|\Omega(s)\rangle \\ &= (1-s) \sum_{i=1}^n \frac{1}{2} [1 - \cos\varphi_i(s) \sin\theta_i(s)] \\ &\quad + s \sum_{x \in \{0,1\}^n} f(x) \prod_{j:x_j=0} \cos^2\left(\frac{\theta_j(s)}{2}\right) \\ &\quad \times \prod_{j:x_j=1} \sin^2\left(\frac{\theta_j(s)}{2}\right), \end{aligned} \quad (\text{D10})$$

with the probability of finding the all-zero state at the end of the evolution (which is the ground state in our case), as $\prod_{i=1}^n \cos^2(\theta_i(1)/2)$. The quantum Hamiltonian obeys qubit permutation symmetry: $PHP = H$, where P is a unitary operator that performs an arbitrary permutation of the qubits. This implies that our classical Lagrangian obeys the same symmetry:

$$\begin{aligned} \mathcal{L}' &\equiv i\langle\Omega(s)|P\frac{d}{ds}P|\Omega(s)\rangle - t_f\langle\Omega(s)|PH(s)P|\Omega(s)\rangle \\ &= i\langle\Omega(s)|\frac{d}{ds}|\Omega(s)\rangle - t_f\langle\Omega(s)|H(s)|\Omega(s)\rangle = \mathcal{L}, \end{aligned} \quad (\text{D11})$$

where the derivative operator is trivially permutation symmetric. Therefore, the Euler-Lagrange equations of motion derived from this action will be identical for all spins. Thus, if we have symmetric initial conditions, i.e., $[\theta_i(0), \varphi_i(0)] = [\theta_j(0), \varphi_j(0)] \forall i, j$, then the time-evolved state will also be symmetric:

$$\begin{aligned} [\theta_i(s), \varphi_i(s)] &= [\theta_j(s), \varphi_j(s)] \quad \forall i, j \\ \forall s &\in [0, 1]. \end{aligned} \quad (\text{D12})$$

As we show below, under the assumption of a permutation-symmetric initial condition, we need to solve only two (instead of $2n$) semiclassical equations of motion:

$$\frac{n}{2} \sin\theta(s)\theta'(s) - t_f\partial_{\varphi(s)}V_{\text{SC}}^{\text{sym}}[\theta(s), \varphi(s), s] = 0, \quad (\text{D13a})$$

$$-\frac{n}{2} \sin\theta(s)\varphi'(s) - t_f\partial_{\theta(s)}V_{\text{SC}}^{\text{sym}}[\theta(s), \varphi(s), s] = 0, \quad (\text{D13b})$$

where we define the symmetric effective potential $V_{\text{SC}}^{\text{sym}}$ as

$$\begin{aligned} V_{\text{SC}}^{\text{sym}}[\theta(s), \varphi(s), s] &\equiv \langle\Omega^{\text{sym}}(s)|H(s)|\Omega^{\text{sym}}(s)\rangle \\ &= (1-s) \frac{n}{2} [1 - \cos\varphi(s) \sin\theta(s)] \\ &\quad + s \sum_{w=0}^n f(w) \frac{n}{w} \sin^{2w}\left(\frac{\theta(s)}{2}\right) \\ &\quad \times \cos^{2(n-w)}\left(\frac{\theta(s)}{2}\right), \end{aligned} \quad (\text{D14})$$

and $|\Omega^{\text{sym}}(s)\rangle$ is simply $|\Omega(s)\rangle$, with all the θ 's and φ 's set equal. Note that in the main text [see Eq. (12)] we slightly abuse notation for simplicity, and use V_{SC} instead of $V_{\text{SC}}^{\text{sym}}$. The probability of finding the all-zero bit string at the end of the evolution is accordingly given by $\cos^{2n}[\theta(1)/2]$. We would have arrived at the same equations of motion had we used the symmetric spin-coherent state in our path-integral derivation, but that would have been an artificial restriction. In our present derivation the symmetry of the dynamics naturally imposes this restriction.

Note that the object in Eq. (D10) involves a sum over all 2^n bit strings and is thus exponentially hard to compute; on the other hand, the object in Eq. (D14) involves only a sum over n terms and is thus easy to compute. Therefore, just as in the quantum case—where due to permutation symmetry the quantum evolution is restricted to the $(n+1)$ -dimensional subspace of symmetric states instead of the full 2^n -dimensional Hilbert space—given knowledge of the symmetry of the problem, we can efficiently compute the SVD potential and efficiently solve the SVD equations of motion.

We also remark that the computation of the potential in Eq. (D10) is significantly simplified if our cost function $f(x)$ is given in terms of a local Hamiltonian. For example, if $H(1) = \sum_{i,j} J_{ij} \sigma_i^z \sigma_j^z$, then

$$V_{\text{SC}}(\{\theta_i\}, \{\varphi_i\}, 1) = \sum_{i,j} J_{ij} \cos\theta_i \cos\theta_j, \quad (\text{D15})$$

which is easy to compute as it is a sum over poly(n) number of terms.

Let us now derive the symmetric SVD equations of motion Eq. (D13). Without any restriction to symmetric spin-coherent states, the SVD equations of motion, for the pair θ_i, φ_i , read

$$\frac{1}{2} \sin \theta_i(s) \theta_i'(s) - t_f \partial_{\varphi_i(s)} V_{\text{SC}}(\{\theta_i\}, \{\varphi_i\}, s) = 0, \quad (\text{D16a})$$

$$-\frac{1}{2} \sin \theta_i(s) \varphi_i'(s) - t_f \partial_{\theta_i(s)} V_{\text{SC}}(\{\theta_i\}, \{\varphi_i\}, s) = 0. \quad (\text{D16b})$$

As can be seen by comparing Eqs. (D13) and (D16), it is sufficient to show that

$$\frac{\partial}{\partial \theta_i} V_{\text{SC}}|_{\theta_j=\theta, \varphi_j=\varphi \forall j} = \frac{1}{n} \frac{\partial}{\partial \theta} V_{\text{SC}}^{\text{sym}}, \quad (\text{D17})$$

and an analogous statement holds for derivatives with respect to φ . This claim is easily seen to hold true for the term multiplying $(1-s)$ in Eq. (D10):

$$\begin{aligned} & \frac{\partial}{\partial \theta_i} \sum_{i=1}^n \frac{1}{2} [1 - \cos \varphi_i(s) \sin \theta_i(s)]|_{\theta_j=\theta, \varphi_j=\varphi \forall j} \\ &= \frac{\partial}{\partial \theta} \frac{1}{2} [1 - \cos \varphi(s) \sin \theta(s)] \\ &= \frac{1}{n} \frac{\partial}{\partial \theta} V_{\text{SC}}^{\text{sym}}(\theta, \phi, s = 0), \end{aligned} \quad (\text{D18})$$

where in the last line we use Eq. (D14). Next, we focus on the term multiplying s in Eq. (D10). This term has no φ dependence and thus we consider only the θ derivatives. First note that

$$\begin{aligned} & \frac{\partial}{\partial \theta_i} V_{\text{SC}}(\{\theta_i\}, \{\varphi_i\}, s = 1) \\ &= \sum_{x \in \{0,1\}^n} f(x) \prod_{j: x_j=0} \cos^2\left(\frac{\theta_j}{2}\right) \prod_{j: x_j=1} \sin^2\left(\frac{\theta_j}{2}\right) \\ & \quad \times \left[-\delta_{x_i,0} \sec^2\left(\frac{\theta_i}{2}\right) + \delta_{x_i,1} \csc^2\left(\frac{\theta_i}{2}\right) \right] \frac{\sin \theta_i}{2}. \end{aligned} \quad (\text{D19})$$

Now, we set all the θ_i 's equal. Let us define $p(\theta) \equiv \sin^2(\theta/2)$. Using this and the fact that f is only a function of the Hamming weight (which is equivalent to the qubit permutation symmetry), we can rewrite the last expression, after a few steps of algebra, as

$$\begin{aligned} & \sum_{w=0}^n f(w) p^{w-1} (1-p)^{n-w-1} \partial_{\theta} p \\ & \quad \times \left[(1-p) \binom{n-1}{w-1} - p \binom{n-1}{w} \right] \\ &= \sum_{w=0}^n f(w) p^{w-1} (1-p)^{n-w-1} \partial_{\theta} p \left[\frac{1}{n} \binom{n}{w} (w-np) \right] \\ &= \frac{1}{n} \frac{\partial}{\partial \theta} V_{\text{SC}}^{\text{sym}}(\theta, \varphi, s = 1). \end{aligned} \quad (\text{D20})$$

Similar to the quantum case, we can perform SVD without explicitly imposing the permutation symmetry and obtain the same results. Here, too, we are forced to explicitly exploit the symmetry due to the nonlocal nature of the problem under consideration, which makes directly implementing the SVD oracle (without the symmetry) exponentially hard. For local problems, we can efficiently implement the SVD oracle.

In the results presented in the main text, it is the implementation of SA that does not share this symmetry. However, while the quantum algorithms and SVD can be implemented without knowledge of the symmetry and still retain their advantage, an implementation of SA that uses the symmetry would require intimate knowledge of the problem. This would be an unfair advantage for SA, not for the quantum evolution.

4. Simulated quantum annealing

An alternative method to simulated annealing, simulated quantum annealing (or path-integral Monte Carlo along the quantum annealing schedule) [28,29], is an annealing algorithm based on discrete-time path-integral quantum Monte Carlo simulations of the transverse field Ising model using Monte Carlo dynamics. At a given time t along the anneal, the Monte Carlo dynamics samples from the Gibbs distribution defined by the action

$$S[\mu] = \Delta(t) \sum_{\tau} H_P(\mu_{\cdot, \tau}) - J_{\perp}(t) \sum_{i, \tau} \mu_{i, \tau} \mu_{i, \tau+1}, \quad (\text{D21})$$

where $\Delta(t) = \beta B(t)/N_{\tau}$ is the spacing along the timelike direction, $J_{\perp} = -\ln\{\tanh[A(t)/2]\}/2$ is the ferromagnetic spin-spin coupling along the timelike direction, and μ denotes a spin configuration with a spacelike direction (the original problem direction, indexed by i) and a timelike direction (indexed by τ). $\mu_{\cdot, \tau}$ denotes all spins along the τ th slice. For our spin updates, we perform Wolff cluster updates [57] along the imaginary-time direction only. For each spacelike slice, a random spin along the timelike direction is picked. The neighbors of this spin are added to the cluster (assuming they are parallel) with probability

$$P = 1 - \exp(-2J_{\perp}). \quad (\text{D22})$$




FIG. 12. Scaling of the p_{GS} peak of QA for the fixed plateau problem with $l = 0$ and $u = 6$. (a) p_{GS} versus t_f curves for several problem sizes n . The peak at $t_f \approx 10$ is the cause of the optimal annealing time. (b) For each p_{GS} versus t_f curve, we fit a Gaussian to the peak. Plotted is the standard deviation of the fitted Gaussians as a function of n . An inverse polylogarithmic function provides an excellent fit to this data.

When all neighbors of the spin have been checked, the newly added spins are checked. When all spins in the cluster have had their neighbors along the timelike direction tested, the cluster is flipped according to the Metropolis probability using the spacelike change in energy associated with flipping the cluster. A single sweep involves attempting to update a single cluster on each spacelike slice.

As in SA, we can use SQA both as an *annealer* and as a *solver* [52]. In the former, we randomly pick one of the states on the Trotter slices at the end of the evolution as the output of the algorithm, while for the latter, we pick the state with the lowest energy found along the entire anneal as the output of the algorithm. We use the latter to maximize the performance of the algorithm.

APPENDIX E: BEHAVIOR OF p_{GS} VERSUS t_f CURVES

We find that for many of the PHWO problems studied, the optimal t_f lies around $t_f = 10$. This is because there is a peak in the probability of finding the ground state p_{GS} at this t_f . Moreover, we find that this peak becomes increasingly higher as the problem size n grows. This is what allows the problem to have an $\mathcal{O}(1)$ scaling. Since this peak becomes increasingly sharper with growing n , there may be the worry that one might need an arbitrarily high precision in setting $t_f \approx t_f^{\text{opt}}$. We address this concern by showing that, in fact, the width of the p_{GS} versus t_f curve decreases as $\mathcal{O}[1/\text{polylog}(n)]$ for the fixed plateau. This shows that we require only a polylogarithmically increasing precision in our ability to set t_f at the optimal value in order to obtain the speedup.

The evidence is summarized in Fig. 12. The first plot, Fig. 12(a), shows p_{GS} versus t_f curves for several values

of n . The second plot, Fig. 12(b), shows the scaling of the standard deviation of Gaussian fit to the peak at t_f^{opt} . This scaling is well matched by polylogarithmic fit.

-
- [1] H. Nishimori, *Statistical Physics of Spin Glasses and Information Processing: An Introduction* (Oxford University Press, Oxford, 2001).
 - [2] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, *Optimization by Simulated Annealing*, *Science* **220**, 671 (1983).
 - [3] S. Suzuki and A. Das, *Discussion and Debate—Quantum Annealing: The Fastest Route to Quantum Computation?*, *Eur. Phys. J. Spec. Top.* **224**, 1 (2015).
 - [4] B. Heim, T. F. Rønnow, S. V. Isakov, and M. Troyer, *Quantum versus Classical Annealing of Ising Spin Glasses*, *Science* **348**, 215 (2015).
 - [5] P. Ray, B. K. Chakrabarti, and A. Chakrabarti, *Sherrington-Kirkpatrick Model in a Transverse Field: Absence of Replica Symmetry Breaking due to Quantum Fluctuations*, *Phys. Rev. B* **39**, 11828 (1989).
 - [6] 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).
 - [7] T. Kadowaki and H. Nishimori, *Quantum Annealing in the Transverse Ising Model*, *Phys. Rev. E* **58**, 5355 (1998).
 - [8] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, *A Quantum Adiabatic Evolution Algorithm Applied to Random Instances of an NP-Complete Problem*, *Science* **292**, 472 (2001).
 - [9] A. Das and B. K. Chakrabarti, *Colloquium: Quantum Annealing and Analog Quantum Computation*, *Rev. Mod. Phys.* **80**, 1061 (2008).
 - [10] S. Suzuki, J. Inoue, and B. K. Chakrabarti, *Quantum Ising Phases and Transitions in Transverse Ising Models*,

- 2nd ed., Lecture Notes in Physics, Vol. 862 (Springer-Verlag, Berlin, 2013).
- [11] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, *Quantum Computation by Adiabatic Evolution*, arXiv:quant-ph/0001106.
- [12] E. Farhi, J. Goldstone, and S. Gutmann, *Quantum Adiabatic Evolution Algorithms versus Simulated Annealing*, arXiv:quant-ph/0201031.
- [13] E. Farhi, J. Goldstone, and S. Gutmann, *Quantum Adiabatic Evolution Algorithms with Different Paths*, arXiv:quant-ph/0208135.
- [14] G. Schaller and R. Schützhold, *The Role of Symmetries in Adiabatic Quantum Algorithms*, *Quantum Inf. Comput.* **10**, 0109 (2010).
- [15] S. Boixo, V. N. Smelyanskiy, A. Shabani, S. V. Isakov, M. Dykman, V. S. Denchev, M. Amin, A. Smirnov, M. Mohseni, and H. Neven, *Computational Role of Collective Tunneling in a Quantum Annealer*, arXiv:1411.4036.
- [16] B. W. Reichardt, in *Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing, STOC '04*, (ACM, New York, 2004), pp. 502–510.
- [17] W. van Dam, M. Mosca, and U. Vazirani, in *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science, 2001* (IEEE, New York, 2001), pp. 279–287.
- [18] Hamming weight can also be understood as the magnetization of a system of spins.
- [19] L. T. Brady and W. van Dam, *Quantum Monte Carlo Simulations of Tunneling in Quantum Adiabatic Optimization*, *Phys. Rev. A* **93**, 032304 (2016).
- [20] By an entropic barrier, we mean that the hardness for SA is predicated purely on the fact that there are many more ways to move away from the ground state than move towards the ground state.
- [21] J. Roland and N. J. Cerf, *Quantum Search by Local Adiabatic Evolution*, *Phys. Rev. A* **65**, 042308 (2002).
- [22] A. T. Rezakhani, A. K. Pimachev, and D. A. Lidar, *Accuracy versus Run Time in an Adiabatic Quantum Search*, *Phys. Rev. A* **82**, 052305 (2010).
- [23] John R. Klauder, *Path Integrals and Stationary-Phase Approximations*, *Phys. Rev. D* **19**, 2349 (1979).
- [24] T. F. Rønnow, Z. Wang, J. Job, S. Boixo, S. V. Isakov, D. Wecker, J. M. Martinis, D. A. Lidar, and M. Troyer, *Defining and Detecting Quantum Speedup*, *Science* **345**, 420 (2014).
- [25] J. A. Smolin and G. Smith, *Classical Signature of Quantum Annealing*, *Front. Phys.*, DOI: (2014).
- [26] T. Albash, T. F. Rønnow, M. Troyer, and D. A. Lidar, *Reexamining Classical and Quantum Models for the D-Wave One Processor*, *Eur. Phys. J. Spec. Top.* **224**, 111 (2015).
- [27] S. A. Owerre and M. B. Paranjape, *Macroscopic Quantum Tunneling and Quantum-Classical Phase Transitions of the Escape Rate in Large Spin Systems*, *Phys. Rep.* **546**, 1 (2015).
- [28] R. Martoňák, G. E. Santoro, and E. Tosatti, *Quantum Annealing by the Path-Integral Monte Carlo Method: The Two-Dimensional Random Ising Model*, *Phys. Rev. B* **66**, 094203 (2002).
- [29] G. E. Santoro, R. Martoňák, E. Tosatti, and R. Car, *Theory of Quantum Annealing of an Ising Spin Glass*, *Science* **295**, 2427 (2002).
- [30] A. Messiah, *Quantum Mechanics* (Dover Publication, New York, 1999).
- [31] We note that while this adiabatic criterion matches the numerical scaling we observe for the quantum evolution, it is well known to be neither exact nor general; see, e.g., Refs. [32–34].
- [32] S. Jansen, M.-B. Ruskai, and R. Seiler, *Bounds for the Adiabatic Approximation with Applications to Quantum Computation*, *J. Math. Phys. (N.Y.)* **48**, 102111 (2007).
- [33] M. H. S. Amin, *Consistency of the Adiabatic Theorem*, *Phys. Rev. Lett.* **102**, 220401 (2009).
- [34] D. A. Lidar, A. T. Rezakhani, and A. Hamma, *Adiabatic Approximation with Exponential Accuracy for Many-Body Systems and Quantum Computation*, *J. Math. Phys. (N.Y.)* **50**, 102106 (2009).
- [35] V. T. Stefanov, *Mean Passage Times for Tridiagonal Transition Matrices*, *J. Appl. Probab.* **32**, 846 (1995).
- [36] O. Krafft and M. Schaefer, *Mean Passage Times for Tridiagonal Transition Matrices and a Two-Parameter Ehrenfest Urn Model*, *J. Appl. Probab.* **30**, 964 (1993).
- [37] This may be an artifact of our implementation of SQA, whereby we only include cluster updates along the imaginary-time direction and not along the spatial (problem) direction. An implementation with spacelike cluster updates may allow SQA via its thermal relaxation to mimic the tunneling of AQA more efficiently. Whether this is the case will be addressed in future work.
- [38] R. D. Somma, D. Nagaj, and M. Kieferová, *Quantum Speedup by Quantum Annealing*, *Phys. Rev. Lett.* **109**, 050501 (2012).
- [39] E. Crosson, E. Farhi, C. Y.-Y. Lin, H.-H. Lin, and P. Shor, *Different Strategies for Optimization Using the Quantum Adiabatic Algorithm*, arXiv:1401.7320.
- [40] I. Hen, *Continuous-Time Quantum Algorithms for Unstructured Problems*, *J. Phys. A* **47**, 045305 (2014).
- [41] D. S. Steiger, T. F. Rønnow, and M. Troyer, *Heavy Tails in the Distribution of Time to Solution for Classical and Quantum Annealing*, *Phys. Rev. Lett.* **115**, 230501 (2015).
- [42] M. Jarret and S. P. Jordan, *The Fundamental Gap for a Class of Schrödinger Operators on Path and Hypercube Graphs*, *J. Math. Phys. (N.Y.)* **55**, 052104 (2014).
- [43] W. M. Kaminsky, in *Quantum Optimization Workshop—The Fields Institute for Research in Mathematical Sciences*, 2014.
- [44] W. M. Kaminsky, *The USC Quantum Information Seminar*, 2014.
- [45] N. V. Vitanov and B. M. Garraway, *Landau-Zener Model: Effects of Finite Coupling Duration*, *Phys. Rev. A* **53**, 4288 (1996).
- [46] N. V. Vitanov and B. M. Garraway, *Erratum: Landau-Zener Model: Effects of Finite Coupling Duration*, *Phys. Rev. A* **54**, 5458 (1996).
- [47] S. Boixo and R. D. Somma, *Necessary Condition for the Quantum Adiabatic Approximation*, *Phys. Rev. A* **81**, 032308 (2010).
- [48] R. A. Horn and C. R. Johnson, *Matrix Analysis* (Cambridge University Press, Cambridge, England, 2012).
- [49] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, *Equation of State Calculations by Fast Computing Machines*, **21**, 1087 (1953).

- [50] W. K. Hastings, *Monte Carlo Sampling Methods Using Markov Chains and Their Applications*, *Biometrika* **57**, 97 (1970).
- [51] V. I. Manousiouthakis and M. W. Deem, *Strict Detailed Balance is Unnecessary in Monte Carlo Simulation*, *J. Chem. Phys.* **110**, 2753 (1999).
- [52] I. Hen, J. Job, T. Albash, T. F. Rønnow, M. Troyer, and D. Lidar, *Probing for Quantum Speedup in Spin Glass Problems with Planted Solutions*, *Phys. Rev. A* **92**, 042325 (2015).
- [53] J. R. Cash and A. H. Karp, *A Variable Order Runge-Kutta Method for Initial Value Problems with Rapidly Varying Right-Hand Sides*, *ACM Trans. Math. Softw.* **16**, 201 (1990).
- [54] J. R. Dormand and P. J. Prince, *A Family of Embedded Runge-Kutta Formulae*, *J. Comput. Appl. Math.* **6**, 19 (1980).
- [55] A. Messiah, *Quantum Mechanics* (North-Holland, Amsterdam, 1962), Vol. 2.
- [56] F. T. Arecchi, E. Courtens, R. Gilmore, and H. Thomas, *Atomic Coherent States in Quantum Optics*, *Phys. Rev. A* **6**, 2211 (1972).
- [57] U. Wolff, *Collective Monte Carlo Updating for Spin Systems*, *Phys. Rev. Lett.* **62**, 361 (1989).