Lower Bounds on Quantum Annealing Times

Luis Pedro García-Pintos,[1,2,3,*] Lucas T. Brady,[4,5] Jacob Bringewatt,[1,2] and Yi-Kai Liu[1,6]

[1]Joint Center for Quantum Information and Computer Science, University of Maryland, College Park, Maryland 20742, USA
[2]Joint Quantum Institute, University of Maryland, College Park, Maryland 20742, USA
[3]Theoretical Division (T4), Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA
[4]Quantum Artificial Intelligence Laboratory, NASA Ames Research Center, Moffett Field, California 94035, USA
[5]KBR, 601 Jefferson Street, Houston, Texas 77002, USA
[6]Applied and Computational Mathematics Division, National Institute of Standards and Technology, Gaithersburg, Maryland 20899, USA

(Received 9 November 2022; accepted 21 February 2023; published 5 April 2023)

The adiabatic theorem provides sufficient conditions for the time needed to prepare a target ground state. While it is possible to prepare a target state much faster with more general quantum annealing protocols, rigorous results beyond the adiabatic regime are rare. Here, we provide such a result, deriving lower bounds on the time needed to successfully perform quantum annealing. The bounds are asymptotically saturated by three toy models where fast annealing schedules are known: the Roland and Cerf unstructured search model, the Hamming spike problem, and the ferromagnetic p-spin model. Our bounds demonstrate that these schedules have optimal scaling. Our results also show that rapid annealing requires coherent superpositions of energy eigenstates, singling out quantum coherence as a computational resource.

DOI: 10.1103/PhysRevLett.130.140601

Introduction.—Generic computational tasks can be mapped to finding the ground state of a Hamiltonian. This is the basis for quantum annealing and adiabatic quantum computing [1–4]. In these approaches, a computational protocol consists of initializing a system in an easy-to-prepare ground state of a Hamiltonian H_0. Thereafter, a time-dependent evolution is performed where the Hamiltonian transitions from H_0 to a Hamiltonian H_1, whose ground state provides the solution to the desired problem. That is, the system is driven by

$$H(t) = (1 - g_t)H_0 + g_t H_1, \tag{1}$$

where the "annealing schedule" g satisfies $g_0 = 0$ and $g_{t_f} = 1$ and t_f is the total duration of the process.

If the transition $H_0 \to H_1$ is slow enough, the adiabatic theorem [5] ensures that the final state is close to the ground state $|E_0^{t_f}\rangle$ of $H(t_f) = H_1$, in which case the protocol performs the desired computation. More precisely, the system remains close to the ground state at all times if $t_f \geq T_{\text{adiab}}$, for $T_{\text{adiab}} \sim \theta/\Delta^2$, where $\theta = \max_t \|d/d(t/t_f)H(t)\|$, $\|\cdot\|$ denotes the spectral norm, and Δ is the minimum energy gap between the instantaneous ground state and first excited state of H_t over the whole schedule (tighter bounds can also be found in Refs. [3,6]). This mechanism is as powerful as standard quantum computation [7]. Throughout this work, we use $\sim$ to denote leading-order terms, up to multiplicative constants.

The condition $t_f \geq T_{\text{adiab}}$ is a sufficient one to perform adiabatic computation. Necessary conditions were also derived in Refs. [8,9]. In Ref. [8], it was shown that an adiabatic annealing process requires at least a time $\tau_{\text{adiab}} \sim L/\Delta$, where $L := \int_0^{t_f} \|d|\psi_t\rangle/dt\| dt$ is the length of the adiabatic path that the state $|\psi_t\rangle$ of the system takes in state space. An algorithm for achieving such $\sim 1/\Delta$ scaling by making use of an oracle for the gap is given in Ref. [10]. Reference [9] used bounds on the speed of adiabatic evolution to derive necessary conditions $t_f \geq \tau_{\text{adiab}}$ on adiabatic annealing times. Throughout this Letter, we denote lower bounds on the time of an annealing process by τ_* (for some descriptive $*$); we use T_{adiab} to denote timescales that ensure that the process is adiabatic.

However, adiabaticity is not a requirement for annealing—it is simply a (powerful yet demanding) condition that guarantees the successful preparation of the desired ground state with a bounded error. This does not exclude the existence of nonadiabatic annealing schedules that take the system to the desired target state more quickly. This is the motivation behind a plethora of popular (and somewhat overlapping) approaches including the quantum approximate optimization algorithm (QAOA) [11,12], diabatic quantum annealing [13], counterdiabatic driving [14–16], and optimal control [16]. Unfortunately, these approaches

are often heuristic with limited theoretical guarantees on performance.

In this Letter, we derive saturable lower bounds on the time necessary for the system to approach a desired target state $|E_0^{t_f}\rangle$ for quantum annealing protocols. In this way, we find general conditions that constrain how fast annealing can be successfully performed, including beyond the adiabatic regime.

Asymptotically saturable bounds on annealing times.—We consider

$$C_1(\rho_t) \coloneqq \min_{\sigma_t} \|\rho_t - \sigma_t\|_1 \tag{2}$$

as a measure of energy coherence of the system's state $\rho_t = |\psi_t\rangle\langle\psi_t|$ [17–19], where σ_t is diagonal in the eigenbasis of $H(t) = \sum_j E_j^t |E_j^t\rangle\langle E_j^t|$ and $\|A\|_1 \coloneqq \mathrm{Tr}(\sqrt{AA^\dagger})$ denotes the trace norm. It holds that $C_1(\rho_t) \leq C_{l_1}(\rho_t)$, where $C_{l_1}(\rho_t) \coloneqq \sum_{j\neq k} |\langle E_k^t|\rho_t|E_j^t\rangle|$ is another popular measure of coherence [18].

Without loss of generality, we take the ground state energies of H_0 and H_1 to be zero, and we denote their time evolving energy expectation values as $\langle H_0\rangle_t \coloneqq \mathrm{Tr}(\rho_t H_0)$ and $\langle H_1\rangle_t \coloneqq \mathrm{Tr}(\rho_t H_1)$. The aim of a successful annealing schedule is to maximize the probability p_{0,t_f} to end in the ground state (or ground subspace in a degenerate spectrum) of $H(t_f) = H_1$ in the shortest t_f possible, where $p_{j,t} \coloneqq \langle E_j^t|\rho_t|E_j^t\rangle$.

With this setup, we derive a hierarchy of lower bounds on the time t_f needed to perform annealing [20]:

$$t_f \geq \tau_{\mathrm{anneal1}} \geq \tau_{\mathrm{anneal2}} \geq \tau_{\mathrm{anneal3}}, \tag{3}$$

with

$$\tau_{\mathrm{anneal1}} \coloneqq 2\frac{\langle H_0\rangle_{t_f} + \langle H_1\rangle_0 - \langle H_1\rangle_{t_f}}{\|[H_1, H_0]\| \frac{1}{t_f}\int_0^{t_f} C_1(\rho_t)dt}, \tag{4a}$$

$$\tau_{\mathrm{anneal2}} \coloneqq \frac{\langle H_0\rangle_{t_f} + \langle H_1\rangle_0 - \langle H_1\rangle_{t_f}}{\|[H_1, H_0]\| \frac{1}{t_f}\int_0^{t_f} \sqrt{1-\sum_j p_{j,t}^2}dt}, \tag{4b}$$

$$\tau_{\mathrm{anneal3}} \coloneqq \frac{\langle H_0\rangle_{t_f} + \langle H_1\rangle_0 - \langle H_1\rangle_{t_f}}{\|[H_1, H_0]\|}. \tag{4c}$$

These limits on the time to reach a solution through annealing processes constitute the main result of this Letter.

While the first two bounds depend on the trajectory of the state of the system through the annealing process, the loosest of the bounds, τ_{anneal3}, depends only on properties of the Hamiltonians H_0 and H_1 and on how close the final state is to the desired ground state. The error term $\langle H_1\rangle_{t_f}$ describes how far the final state is from the desired solution. For perfect annealing, $\langle H_1\rangle_{t_f} = 0$. Note, too, that the second bound implies that an annealing process where the system remains in the instantaneous ground state at all times, $p_{0,t} = 1$, requires an infinite time, since τ_{anneal2} diverges. This is consistent with truly adiabatic evolution.

Alternatively, these bounds set constraints on the minimum coherence and the minimum excitations needed to anneal a system. For concreteness, assume one desires to perfectly anneal a system within a time t_f much shorter than the adiabatic timescale. Equations (3), (4a), and (4b) then imply that

$$\int_0^{t_f} \sqrt{1-\sum_j p_{j,t}^2}dt \geq \frac{1}{2}\int_0^{t_f} C_1(\rho_t)dt \geq \frac{\langle H_0\rangle_{t_f} + \langle H_1\rangle_0}{\|[H_1, H_0]\|}. \tag{5}$$

While the rightmost term depends only on the Hamiltonians that define the problem, the other two terms are path dependent. The leftmost term is an entropic quantity that characterizes energy excitations. The middle term depends on the energy coherence of the system. Fast annealing, thus, requires populating many energy levels. This serves as a sort of converse for the adiabatic theorem, which states that no excitations occur as long as the process is sufficiently slow. Fast annealing also requires coherence $C_1 > 0$ in the energy basis. This cements the role of coherence as a resource in quantum computations [21–23].

Next, we prove that these bounds are asymptotically saturable in the size of the system, correctly capturing the optimal annealing timescales of certain toy models.

Example of optimally fast annealing: Unstructured search.—Consider the standard model for unstructured search over d elements on an analog quantum computer [24]. Let the system be initialized in a state $|E_0(0)\rangle = |\psi_0\rangle = (1/\sqrt{d})\sum_{j=1}^d |j\rangle$ [25], with

$$H_0 = \mathbb{1} - |\psi_0\rangle\langle\psi_0|, \qquad H_1 = \mathbb{1} - |m\rangle\langle m|. \tag{6}$$

The aim is to find the eigenstate $|E_0^{t_f}\rangle \equiv |m\rangle$ among the d possible states. In the limit $d \gg 1$, Roland and Cerf proved that an optimized adiabatic schedule drives the system to a state that is close to the desired state, with $|\langle E_0^{t_f}|\psi_{t_f}\rangle|^2 \geq 1-\epsilon^2$, in an adiabatic annealing time $T_{\mathrm{adiab}} = (\pi/2\epsilon)\sqrt{d}$ [24]. That is, whereas classically it takes $\sim d$ trials to find an item from an unstructured list, quantum mechanical protocols can do this in a time $\sim\sqrt{d}$, recovering the $1/\sqrt{d}$ speedup from Grover's algorithm in the digital case.

Using that

$$\|[H_1, H_0]\| = \frac{1}{\sqrt{d}}, \qquad \langle H_1\rangle_0 = 1 - \frac{1}{d}, \tag{7}$$

and

$$\langle H_0\rangle_{t_f} - \langle H_1\rangle_{t_f} = |\langle\psi_{t_f}|m\rangle|^2 - |\langle\psi_{t_f}|\psi_0\rangle|^2 \geq p_{0,t_f}, \tag{8}$$

we obtain that any annealing protocol requires a time $t_f \geq \tau_{\text{anneal3}}$, with

$$\tau_{\text{anneal3}} \geq \frac{1 - \frac{1}{d} + p_{0,t_f}}{\frac{1}{\sqrt{d}}} \approx 2\sqrt{d}. \tag{9}$$

That is, the scaling with system size of Roland and Cerf's optimal adiabatic protocol cannot be beaten by diabatic protocols. This also shows that the lower bound Eq. (4) on annealing times is (asymptotically) saturable.

If we further impose, as Roland and Cerf do, that $p_{0,t} \geq 1 - \epsilon^2$ with $\epsilon \ll 1$, we get that $1 - \sum_j p_{j,t}^2 \leq 1 - p_{0,t}^2 \lesssim 2\epsilon^2$ [26]. Then, we find that adiabatic annealing requires a time $t_f \geq \tau_{\text{anneal2}}$, where

$$\tau_{\text{anneal2}} \gtrsim \frac{1 - \frac{1}{d} + p_{0,t_f}}{\frac{1}{\sqrt{d}}\sqrt{2}\epsilon} \approx \frac{\sqrt{2d}}{\epsilon}. \tag{10}$$

Both the scaling with system size d and target distance ϵ are saturated by Roland and Cerf's optimal adiabatic protocol, which requires a time $T_{\text{adiab}} = (\pi/2\epsilon)\sqrt{d}$.

The gaps between τ_{anneal} and T_{adiab}.—However, as we argued in the introduction, adiabatic schedules can be far from optimal. In certain models, the gap between the timescales in Eq. (4) and the ones implied by the adiabatic theorem can be large. In order to explore such a gap, we consider two toy models where free parameters govern the adiabatic timescale T_{adiab}.

An example where this is the case is the Hamming spike problem, defined by the Hamiltonians

$$H_0 = \frac{1}{2}(N\mathbb{1} - M_x), \tag{11a}$$

$$H_1 = \frac{1}{2}(N\mathbb{1} - M_z) + b(W), \tag{11b}$$

where $M_\xi := \sum_{\nu=1}^N \sigma_\nu^\xi$ is the magnetization of the N qubits along direction $\xi \in \{x, y, z\}$ and b is a function of the so-called Hamming weight operator $W := (N\mathbb{1} - M_z)/2$ which, when acting on a computational basis state, returns its Hamming weight w, defined as the number of ones in the bit string. We assume b is localized around the region $w = (N/4)$ and models a "spike" or "barrier" of some form [27–32]. The barrier is assumed large enough to hinder tunneling of the quantum state during the annealing process but small enough to act perturbatively. In particular, assume the barrier has height $\sim N^\alpha$ and width $\sim N^\beta$ with $\alpha < 1$ and $\beta < \frac{1}{2}$ [28]. The size of the barrier dictates the timescales derived from the adiabatic theorem. It holds that $T_{\text{adiab}} \sim \text{poly}(N)$ when $2\alpha + \beta < 1$ and that $T_{\text{adiab}} \sim \exp(N)$ for $2\alpha + \beta > 1$ [28].

As $|0\rangle$ and $|+\rangle$ are the eigenstates corresponding to the minimum eigenvalues of $-\sigma^z$ and $-\sigma^x$, respectively, the ground states of H_0 and H_1 are $|+\rangle^{\otimes N}$ and $|0\rangle^{\otimes N}$, respectively. Then, assuming ideal annealing gives

$$\langle H_0\rangle_{t_f} = \frac{N}{2}, \qquad \langle H_1\rangle_0 = \frac{N}{2} + \mathcal{O}(N^{\alpha+\beta-1/2}), \tag{12}$$

and it holds that [20]

$$\|[H_0, H_1]\| \leq \frac{N}{2} + \mathcal{O}(N^{\alpha+\beta}). \tag{13}$$

The important thing to note is that these correction terms depend on the area under the barrier curve b and that in most parameter regimes considered (including some with exponentially small spectral gaps [28]) they will scale linearly or sublinearly in N.

Therefore, successful annealing in the Hamming spike problem requires at least a time $t_f \geq \tau_{\text{anneal3}}$, where $\tau_{\text{anneal3}} \gtrsim 1$. Remarkably, this scaling matches that of the numerically optimized annealing schedules [33,34] and the QAOA schedule for this problem [35]. This is in stark contrast to the size-dependent timescales obtained from the adiabatic theorem. This shows a second toy model where the scaling of the new bounds (4) is saturated and that the annealing times for the Hamming spike problem previously found numerically in the literature are, in fact, optimal.

Another toy model with a large gap between adiabatic and nonadiabatic timescales is the p-spin model [36,37]. In the ferromagnetic p-spin model,

$$H_0 = \frac{N}{2}\left(\mathbb{1} - \frac{M_x}{N}\right), \qquad H_1 = \frac{N}{2}\left(\mathbb{1} - \frac{M_z^p}{N^p}\right). \tag{14}$$

The integer $p \geq 1$ governs the timescales in the adiabatic theorem via the minimum gap, which scales as [38]

$$\Delta \sim 1, \qquad p = 1, \tag{15a}$$

$$\Delta \sim N^{-1/3}, \qquad p = 2, \tag{15b}$$

$$\Delta \sim \exp(-N), \qquad p \geq 3, \tag{15c}$$

yielding adiabatic timescales of $T_{\text{adiab}} \sim \{1, N^{2/3}, \exp(2N)\}$, respectively.

In contrast, our bound in Eq. (4) yields

$$\tau_{\text{anneal3}} \geq 2 \quad \forall \; p \geq 1, \tag{16}$$

where we used the fact that the ground states of H_0 and H_1 are $|+\rangle^{\otimes N}$ and $|0\rangle^{\otimes N}$ (for odd p) or $\{|0\rangle^{\otimes N}, |1\rangle^{\otimes N}\}$ (for even p)—implying that $\langle H_0\rangle_{t_f} = \langle H_1\rangle_0 = N/2$—and the fact that $\|[H_1, H_0]\| \leq N/2$ [20].

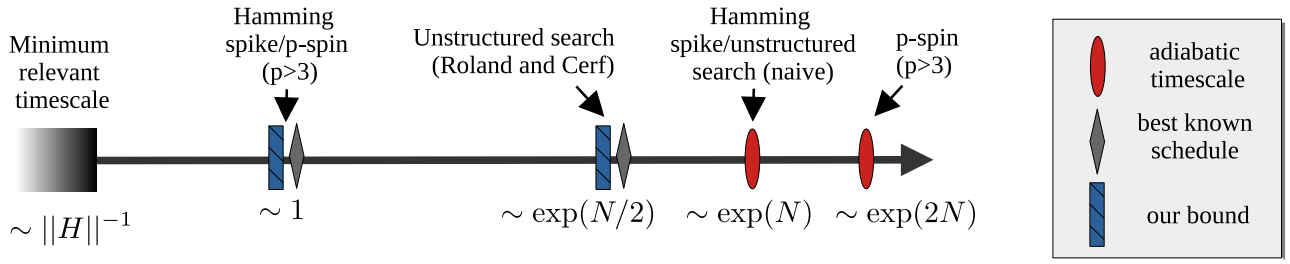

FIG. 1. Annealing timescales. An illustration of the range of possible timescales in annealing problems and how our bounds and the adiabatic timescales fit in the picture.

The outstanding question is, which of these widely different timescales better characterizes the performance of an optimal schedule? For odd N, it is known analytically that a constant time, single-round QAOA-style, or bang-bang, annealing schedule (with $g_t = 1$ for an initial interval of time and $g_t = 0$ for the rest) allows one to exactly reach the target ground state [39]. Equation (16) demonstrates that this scaling is, in fact, optimal. We show a simple proof of this in Supplemental Material [20]. While analytically less straightforward, numerics for even N also indicate $t_f \sim 1$ scaling to reach the target state with high fidelity.

Therefore, we have a third toy model where the optimal schedule saturates the lower bounds τ_{anneal} and where the gap to the adiabatic timescale T_{adiab} is large (exponential for $p \geq 3$). A comparison of annealing timescales in the three toy models considered in this Letter is illustrated in Fig. 1.

Lower bounds for k-local Hamiltonians.—Consider the N-particle Hamiltonians

$$H_0 = \sum_{\nu=1}^{N} h_\nu^0, \qquad H_1 = \sum_{\nu=1}^{N} h_\nu^1, \tag{17}$$

where h_ν^0 and h_ν^1 are k-local Hamiltonians with support on at most k subsystems [40], where $\|h_\nu^0\| = \|h_\nu^1\| = 1$.

The scaling with N of the bounds on annealing times [Eq. (4)] intricately depends on the constituent Hamiltonians. However, it holds that $\|[H_0, H_1]\| \leq kN$, and one can typically expect that $\langle H_0\rangle_{t_f} \sim \langle H_1\rangle_0 \sim N$. This gives that any annealing protocol that aims to connect ground states of k-local Hamiltonians requires a time $t_f \geq \tau_{\text{anneal3}}$ with

$$\tau_{\text{anneal3}} \gtrsim \frac{2}{k}. \tag{18}$$

This scaling is in stark contrast to the one obtained from the adiabatic theorem. For many-body systems, the minimum energy gap between the ground and first excited state typically scale as $\Delta \sim 1/\text{poly}(N)$ or as $\Delta \sim \exp(-N)$ [41]. In the latter case, for example, the adiabatic theorem ensures a schedule that anneals the system if $t_f > T_{\text{adiab}}$ with

$$T_{\text{adiab}} \sim \frac{\theta}{\Delta^2} \sim \exp(2N). \tag{19}$$

This gives the same scaling as, for instance, the p-ferromagnetic spin model for $p = 3$, which is 3-local. However, in that case we found that the scaling of lower bound $\tau_{\text{anneal3}} \geq 2$ was saturated by a single-round QAOA schedule. This, thus, shows that the minimum annealing time [Eq. (18)] for k-local systems is indeed saturable.

Annealing times with extra control Hamiltonians.—So far, we adopted the standard quantum annealing scenario where one carefully tailors a schedule that combines H_0 and H_1 to reach the desired state. However, including extra control Hamiltonians H_C adds freedom to the dynamics that can, in principle, speed up an annealing process [42–46]. One extreme example of this is that of a counterdiabatic Hamiltonian H_{CD} that implements a shortcut to adiabaticity dynamics by inhibiting excitations out of the instantaneous ground state [14,47–50].

How much can extra physical control Hamiltonians speed up an annealing process? Let us assume access to a set of N_C control Hamiltonians $\{H_C^a\}$ with schedules $\{f_t^a\} \geq 0$ such that $f_0^a = f_{t_f}^a = 0$. Their aim is to speed up the transition to the eigenstate $|E_0(t_f)\rangle$ of H_1. The total Hamiltonian is

$$\tilde{H}(t) = H(t) + \sum_{a=1}^{N_C} f_t^a H_C^a. \tag{20}$$

Then, we prove a constraint on the annealing times $t_f \geq \tau_{\text{anneal}}$ under dynamics with the extra control knobs provided by $\{H_C^a\}$, where [20]

$$\tau_{\text{anneal}} := \frac{\langle H_0\rangle_{t_f} + \langle H_1\rangle_0 - \langle H_1\rangle_{t_f}}{\|[H_1, H_0]\| + \sum_{a=1}^{N_C} \|[H_1 - H_0, H_C^a]\|}. \tag{21}$$

Consequently, while a control Hamiltonian $H_C \propto (H_1 - H_0)^p$ may improve the performance of some schedules, it cannot improve the performance of an optimal schedule that saturates the lower bounds (4). Some interesting connections can also be made between Eq. (21) and

counterdiabatic Hamiltonians. As noted in Ref. [51], a counterdiabatic Hamiltonian designed to prepare a target state in time t_f carries that t_f-dependence only in its norm, which scales like $1/t_f$. As counterdiabatic driving allows one to prepare a target state with perfect fidelity for any t_f, the bound in Eq. (21) must go to zero as $t_f \to 0$. This is only possible if $\|[H_1 - H_0, H_{CD}]\| \to \infty$ for $t_f \to \infty$, implying that this term is non-zero for any t_f. That is, one can catalyze the evolution with an additional control Hamiltonian that is equivalent to the counterdiabatic Hamiltonian that will lead to a sizable increase in the denominator of Eq. (21) and, therefore, a reduction in the lower bound on the annealing time.

Finally, Eq. (21) implies a lower limit on the number of control Hamiltonians N_C needed to perform annealing within a short time t_f. For instance, assume that one desires to anneal the system in the maximum timescale $1/\|H(t)\|$ possible given the original Hamiltonian $H(t)$ in Eq. (1). For concreteness, let us assume, as in the k-local Hamiltonians above, that $\langle H_0 \rangle_{t_f} \sim \langle H_1 \rangle_0 \sim N$ and that the control Hamiltonians are also k-local, so that $\|[H_1 - H_0, H_C^a]\| \leq kN$. Then,

$$\|H(t)\|\tau_{\text{anneal}} \geq N \frac{2N}{kN(1+N_C)} \sim \frac{N}{kN_C}, \tag{22}$$

and at least $N_C \sim N/k$ control Hamiltonians are needed to perform annealing at the maximum rate $\|H(t)\|$ defined by the original Hamiltonian. Similarly, Eq. (21) implies that at least $N_C \sim (N/kt_f)$ control Hamiltonians are needed to implement a counterdiabatic Hamiltonian that enforces adiabatic evolution in t_f for a many-body system.

Discussion.—Extensive work has been devoted to understanding the timescales T_{adiab} that ensure that a process is adiabatic. However, less is rigorously known about diabatic schedules that can anneal a system faster. In fact, to our knowledge, the scaling of the optimal annealing time is known only in a few toy models, which include the unstructured search model, the Hamming spike problem, and the p-ferromagnetic spin model.

In this Letter, we derived easy-to-evaluate lower bounds on the times necessary for annealing to occur which are saturated by the best known annealing schedules for all of these toy models. While the Roland and Cerf schedule appears to have optimal scaling even without confirmation from our bounds due to the fact it recovers the Grover-type speedup for unstructured search [52], studies of the Hamming spike problem were numerical in nature. Moreover, we found that previously considered QAOA schedules for the p-ferromagnetic spin model also saturate the lower bounds, proving those schedules to be optimal.

Note that all models considered here are Hamming symmetric. That is, the Hamiltonians are invariant under permutations of basis elements with the same Hamming weight in the computational basis, or, equivalently, they conserve the total spin along a given (z) direction. This high degree of symmetry could conceivably be responsible for the saturation of our bounds in these models. Strikingly, however, the collection of models for which we can show our bounds are saturable exhibit vastly different optimal annealing schedules, ranging from optimized adiabatic schedules to "bang-bang" controls. This means that, if symmetry is indeed responsible for the tightness of the bounds, the direct means by which it causes this tightness is not obvious. We leave exploring this as an open question while observing that the variety of different schedules which saturate bounds provides compelling evidence for their usefulness, especially given the dearth of rigorous results on the timescales needed to perform quantum annealing beyond the adiabatic regime.

Unlike the timescales obtained from the adiabatic theorem, our bounds do not depend on the spectral gap of the system, which makes it easier to evaluate the latter. While we found our bounds to better reflect the timescales of optimized annealing in all toy models considered, this highlights the importance of understanding the role of the spectral gap in the performance of optimal schedules in more physically realistic scenarios. In addition, because our bounds involve the quantum coherence of the system, this suggests an approach to understanding when the system can escape from local minima [55] in the diabatic regime. Finally, the role that geometric locality plays in the lower bounds on annealing times remains a problem to be explored.

We thank Michael Gullans for discussions related to this work. We acknowledge funding by the U.S. Department of Energy (DOE) Advanced Scientific Computing Research (ASCR) Accelerated Research in Quantum Computing program (Grant No. DE-SC0020312), DOE Quantum Systems Accelerator (QSA), National Science Foundation (NSF) Quantum Leap Challenge Institutes (QLCI) (Grant No. OMA-2120757), DOE ASCR Quantum Testbed Pathfinder program (Grant No. DE-SC0019040), NSF PFCQC program, Air Force Office of Scientific Research (AFOSR), Army Research Laboratory (ARO) Multidisciplinary University Research Initiatives (MURI), AFOSR MURI, and Defense Advanced Research Projects Agency (DARPA) Science of Atomic Vapors for New Technologies (SAVaNT) ADVENT. This research was supported in part by the National Science Foundation under Grant No. NSF PHY-1748958, the Heising-Simons Foundation, and the Simons Foundation (216179, LB). The work at Los Alamos National Laboratory was carried out under the auspices of the U.S. DOE and National Nuclear Security Administration (NNSA) under Contract No. DEAC52-06NA25396. We also acknowledge support by the DOE Office of Science, Office of Advanced Scientific Computing Research, Accelerated Research for Quantum Computing program, Fundamental Algorithmic Research

for Quantum Computing (FAR-QC) project. L. T. B. is a KBR employee working under the Prime Contract No. 80ARC020D0010 with the National Aeronautics and Space Administration (NASA) Ames Research Center and is grateful for support from the DARPA RQMLS program under IAA 8839, Annex 128. The United States Government retains, and by accepting the article for publication, the publisher acknowledges that the United States Government retains, a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for United States Government purposes.

*lpgp@lanl.gov

[1] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse Ising model, Phys. Rev. E **58**, 5355 (1998).
[2] P. Hauke, H. G. Katzgraber, W. Lechner, H. Nishimori, and W. D. Oliver, Perspectives of quantum annealing: Methods and implementations, Rep. Prog. Phys. **83**, 054401 (2020).
[3] T. Albash and D. A. Lidar, Adiabatic quantum computation, Rev. Mod. Phys. **90**, 015002 (2018).
[4] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, Quantum computation by adiabatic evolution, arXiv:quant-ph/0001106.
[5] M. Born and V. Fock, Beweis des adiabatensatzes, Z. Phys. **51**, 165 (1928).
[6] 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).
[7] D. Aharonov, W. van Dam, J. Kempe, Z. Landau, S. Lloyd, and O. Regev, Adiabatic quantum computation is equivalent to standard quantum computation, SIAM Rev. **50**, 755 (2008).
[8] S. Boixo and R. D. Somma, Necessary condition for the quantum adiabatic approximation, Phys. Rev. A **81**, 032308 (2010).
[9] J.-H. Chen, Lower bounds for adiabatic quantum algorithms by quantum speed limit, arXiv:2207.01604.
[10] M. Jarret, B. Lackey, A. Liu, and K. Wan, Quantum adiabatic optimization without heuristics, arXiv:1810.04686.
[11] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, arXiv:1411.4028.
[12] N. Barraza, G. A. Barrios, J. Peng, L. Lamata, E. Solano, and F. Albarrán-Arriagada, Analog quantum approximate optimization algorithm, Quantum Sci. Technol. **7**, 045035 (2022).
[13] E. Crosson and D. Lidar, Prospects for quantum enhancement with diabatic quantum annealing, Nat. Rev. Phys. **3**, 466 (2021).
[14] M. V. Berry, Transitionless quantum driving, J. Phys. A **42**, 365303 (2009).
[15] M. Demirplak and S. A. Rice, Adiabatic population transfer with control fields, J. Phys. Chem. A **107**, 9937 (2003).
[16] D. Guéry-Odelin, A. Ruschhaupt, A. Kiely, E. Torrontegui, S. Martínez-Garaot, and J. G. Muga, Shortcuts to adiabaticity: Concepts, methods, and applications, Rev. Mod. Phys. **91**, 045001 (2019).
[17] T. Baumgratz, M. Cramer, and M. B. Plenio, Quantifying Coherence, Phys. Rev. Lett. **113**, 140401 (2014).
[18] S. Rana, P. Parashar, and M. Lewenstein, Trace-distance measure of coherence, Phys. Rev. A **93**, 012110 (2016).
[19] A. Streltsov, G. Adesso, and M. B. Plenio, Colloquium: Quantum coherence as a resource, Rev. Mod. Phys. **89**, 041003 (2017).
[20] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.130.140601 for detailed proofs of some of the statements in the main text.
[21] M. Hillery, Coherence as a resource in decision problems: The Deutsch-Jozsa algorithm and a variation, Phys. Rev. A **93**, 012111 (2016).
[22] H.-L. Shi, S.-Y. Liu, X.-H. Wang, W.-L. Yang, Z.-Y. Yang, and H. Fan, Coherence depletion in the grover quantum search algorithm, Phys. Rev. A **95**, 032307 (2017).
[23] F. Ahnefeld, T. Theurer, D. Egloff, J. M. Matera, and M. B. Plenio, Coherence as a Resource for Shor's Algorithm, Phys. Rev. Lett. **129**, 120501 (2022).
[24] J. Roland and N. J. Cerf, Quantum search by local adiabatic evolution, Phys. Rev. A **65**, 042308 (2002).
[25] Interestingly, the starting state is maximally coherent with respect to the target Hamiltonian H_1 [17].
[26] We adopt a notation where $f \gtrsim a$ ($f \lesssim a$) means $f > c$ ($f < c$) with $c \sim a$.
[27] E. Farhi, J. Goldstone, and S. Gutmann, Quantum adiabatic evolution algorithms versus simulated annealing, arXiv:quant-ph/0201031.
[28] L. T. Brady and W. van Dam, Spectral-gap analysis for efficient tunneling in quantum adiabatic optimization, Phys. Rev. A **94**, 032309 (2016).
[29] B. W. Reichardt, The quantum adiabatic optimization algorithm and local minima, in *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC)* (Association for Computing Machinery, New York, 2004), pp. 502–510.
[30] E. Crosson and M. Deng, Tunneling through high energy barriers in simulated quantum annealing, arXiv:1410.8484.
[31] S. Muthukrishnan, T. Albash, and D. A. Lidar, Tunneling and Speedup in Quantum Optimization for Permutation-Symmetric Problems, Phys. Rev. X **6**, 031010 (2016).
[32] J. Bringewatt, W. Dorland, S. P. Jordan, and A. Mink, Diffusion Monte Carlo approach versus adiabatic computation for local Hamiltonians, Phys. Rev. A **97**, 022323 (2018).
[33] S. Muthukrishnan, T. Albash, and D. A. Lidar, When diabatic trumps adiabatic in quantum optimization, arXiv:1505.01249.
[34] L. T. Brady and W. van Dam, Necessary adiabatic run times in quantum optimization, Phys. Rev. A **95**, 032335 (2017).
[35] A. Bapat and S. Jordan, Bang-bang control as a design principle for classical and quantum optimization algorithms, Quantum Inf. Comput. **19**, 424 (2019).
[36] Y. Seki and H. Nishimori, Quantum annealing with antiferromagnetic fluctuations, Phys. Rev. E **85**, 051112 (2012).
[37] B. Seoane and H. Nishimori, Many-body transverse interactions in the quantum annealing of the p-spin ferromagnet, J. Phys. A **45**, 435301 (2012).
[38] V. Bapst and G. Semerjian, On quantum mean-field models and their quantum annealing, J. Stat. Mech. (2012) P06007.

[39] M. M. Wauters, G. B. Mbeng, and G. E. Santoro, Polynomial scaling of the quantum approximate optimization algorithm for ground-state preparation of the fully connected p-spin ferromagnet, Phys. Rev. A **102**, 062404 (2020).
[40] We are adopting the naming convention typically used in computer science. Note, though, that these Hamiltonians need not be geometrically local. The role of space locality on the minimum annealing timescales remains an interesting problem to be explored.
[41] E. Farhi, J. Goldston, D. Gosset, S. Gutmann, H. B. Meyer, and P. Shor, Quantum adiabatic algorithms, small gaps, and different paths, Quantum Inf. Comput. **11**, 181 (2011).
[42] L. Hormozi, E. W. Brown, G. Carleo, and M. Troyer, Nonstoquastic Hamiltonians and quantum annealing of an ising spin glass, Phys. Rev. B **95**, 184416 (2017).
[43] Y. Susa, J. F. Jadebeck, and H. Nishimori, Relation between quantum fluctuations and the performance enhancement of quantum annealing in a nonstoquastic Hamiltonian, Phys. Rev. A **95**, 042321 (2017).
[44] W. Vinci and D. A. Lidar, Non-stoquastic Hamiltonians in quantum annealing via geometric phases, npj Quantum Inf. **3**, 38 (2017).
[45] E. Crosson, T. Albash, I. Hen, and A. P. Young, De-signing Hamiltonians for quantum adiabatic optimization, Quantum **4**, 334 (2020).
[46] P. R. Hegde, G. Passarelli, A. Scocco, and P. Lucignano, Genetic optimization of quantum annealing, Phys. Rev. A **105**, 012612 (2022).
[47] A. del Campo, Shortcuts to Adiabaticity by Counterdiabatic Driving, Phys. Rev. Lett. **111**, 100502 (2013).
[48] D. Guéry-Odelin, A. Ruschhaupt, A. Kiely, E. Torrontegui, S. Martínez-Garaot, and J. G. Muga, Shortcuts to adiabaticity: Concepts, methods, and applications, Rev. Mod. Phys. **91**, 045001 (2019).
[49] A. del Campo and K. Kim, Focus on shortcuts to adiabaticity, New J. Phys. **21**, 050201 (2019).
[50] G. Passarelli, V. Cataudella, R. Fazio, and P. Lucignano, Counterdiabatic driving in the quantum annealing of the p-spin model: A variational approach, Phys. Rev. Res. **2**, 013283 (2020).
[51] F. Petiziol, B. Dive, F. Mintert, and S. Wimberger, Fast adiabatic evolution by oscillating initial Hamiltonians, Phys. Rev. A **98**, 043436 (2018).
[52] Note that this argument is not rigorous, as it has been argued that analog classical models can also provide this same speedup, calling into question whether the speedup in the Roland and Cerf model is genuinely due to quantum effects or just due to finely tuned analog control [53,54]. Either way, our results show that the schedule achieves optimal scaling given access to H_0 and H_1.
[53] I. Hen, How quantum is the speedup in adiabatic unstructured search?, Quantum Inf. Process. **18**, 1 (2019).
[54] M. Slutskii, T. Albash, L. Barash, and I. Hen, Analog nature of quantum adiabatic unstructured search, New J. Phys. **21**, 113025 (2019).
[55] B. Altshuler, H. Krovi, and J. Roland, Anderson localization makes adiabatic quantum optimization fail, Proc. Natl. Acad. Sci. U.S.A. **107**, 12446 (2010).