

Tighter Lower Bounds on Quantum Annealing Times

Luis Pedro García-Pintos,¹ Mrunmay Sahasrabudhe,² and Christian Arenz²

¹*Theoretical Division (T-4), Los Alamos National Laboratory, New Mexico, 87545, USA*

²*School of Electrical, Computer, and Energy Engineering,
Arizona State University, Tempe, Arizona 85287, USA*

We derive lower bounds on the time needed for a quantum annealer to prepare the ground state of a target Hamiltonian. These bounds do not depend on the annealing schedule and can take the local structure of the Hamiltonian into account. Consequently, the bounds are computable without knowledge of the annealer’s dynamics and, in certain cases, scale with the size of the system. We discuss spin systems where the bounds are polynomially tighter than existing bounds, qualitatively capturing the scaling of the exact annealing times as a function of the number of spins.

I. INTRODUCTION

Quantum annealing is a powerful framework for solving optimization problems on quantum computers [1]. In quantum annealing, one aims to steer the state of a quantum system by properly tailored control functions, from a simple-to-prepare initial state to a target state that encodes the solution to the problem of interest. Typically, the initial and final states are the ground states of an initial and final Hamiltonian H_i and H_f . The goal is to design annealing schedules that dynamically transition between both Hamiltonians such that, at some final time T , the system’s state is close to the desired target ground state. A prominent example of quantum annealing is adiabatic quantum computation. There, the annealing schedule consists of slowly changing the control functions to evolve into the ground state of H_f according to the adiabatic theorem [2]. The scaling of the annealing time T with respect to the system size depends on the complexity of the computational problem to be solved.

Unfortunately, general results that characterize the behavior of T are scarce. The most studied of such results rely on the adiabatic theorem to derive sufficient conditions for successful annealing. The adiabatic theorem shows that the system’s final state will be close to the desired ground state of H_f if the protocol is slow enough. Typically, T must scale as $1/\Delta^2$ for adiabaticity to hold, where Δ is the minimum spectral gap given by the energy difference between the ground and first excited energy eigenstate of the time-dependent Hamiltonian describing the annealer [3]. The spectral gap is typically exponentially small in system size, which leads to exponentially long computing times by adiabatic protocols, noting that necessary conditions for adiabaticity were also studied in Refs. [4] and [5].

However, adiabaticity can be an overly restrictive condition. Often, optimized schedules can be much faster than those guided by the adiabatic theorem (i.e., much faster than $T \sim 1/\Delta^2$). This motivates studying the smallest times T needed in quantum annealing protocols. General lower bounds on quantum annealing times, which hold for non-adiabatic annealing schedules, were derived in Ref. [6]. Leveraging the techniques used in [6], Ref. [7] focuses on the minimum number of rounds needed

to reach a target ground state via an analog optimization algorithm, which is inspired by quantum annealing. The results in Refs. [4–7] give bounds on quantum annealing times given an algorithm’s physical implementation (see also Ref. [8], which considers phenomenological limits to quantum annealing times).

The bounds in Refs. [5–7] were derived using techniques usually known as quantum speed limits [9]. A speed limit is a bound on the rate at which a physical quantity can change. Mandelstam and Tamm proved the first speed limit in Ref. [10]. They showed that an observable’s expectation value $\langle A \rangle$ satisfies $|d\langle A \rangle/dt|^2 \leq 4\text{Var}(A)\text{Var}(H)$ for any isolated quantum system, where $\text{Var}(A) := \langle A^2 \rangle - \langle A \rangle^2$ and $\text{Var}(H)$ are the variances of an observable A and of the system’s Hamiltonian H , respectively. Mandelstam and Tamm’s result sets a trade-off relation between speed and uncertainty. Reference [5] shows that a similar trade-off relation constrains adiabatic quantum annealing. Meanwhile, Ref. [6] shows a trade-off relation between runtime and quantum coherence in arbitrary (possibly diabatic) annealing.

One shortcoming of recent results that aim to describe the minimum runtime T is that they do not accurately reflect the annealing times in systems where locality matters. For instance, the bounds in Refs. [5–7] are loose for quantum annealing or the quantum approximate optimization algorithm (QAOA) that aims to prepare the ground state of an Ising Hamiltonian [11]; the bounds are independent of the system size, but the actual times typically grow with the system size.

A second shortcoming stems from an inherent property of most speed limits: evaluating the bounds often requires having information on the dynamics, e.g., the path taken in Hilbert space. For instance, evaluating Mandelstam and Tamm’s bound $|d\langle A \rangle/dt|^2 \leq 4\text{Var}(A)\text{Var}(H)$ requires knowing $\text{Var}(A)$, which changes as the system evolves. This problem is prevalent in most quantum [12–16] and classical [17–19] speed limits considered in the literature. As such, most quantum speed limits are of limited practical use for accurately capturing the behavior of T in explicitly time-dependent systems, which generally occurs in quantum control and quantum annealing.

In this work, we take steps to address the shortcomings described in the previous two paragraphs by leverag-

ing techniques introduced in [20, 21]. Namely, we derive lower bounds on the annealing time T , which in turn implies lower bounds on the circuit depth needed in QAOA. The bounds (i) can be evaluated without information of annealing schedules or state dynamics, and (ii) are tighter than previously considered bounds [5–7] for certain spin models, where the new bounds can scale with the system size.

The manuscript is structured as follows. We begin in Sec. II by deriving simple lower bounds on annealing times that do not depend on the annealing schedule or the system’s trajectory through Hilbert space. In Sec. III, we leverage the results from Sec. II to derive speed limits that are tighter than previously considered bounds. We illustrate the scaling with the system size of the new bounds on the analog version of Grover’s search algorithm and on a perturbed p -spin model. For each example, we compare to existing bounds in the literature. We show that for optimization problems characterized by a well-connected graph with a low-connected vertex, the new bounds display a polynomial improvement over previous results.

II. FROM SCHEDULE-DEPENDENT TO SCHEDULE-INDEPENDENT SPEED LIMITS

We start by deriving a speed limit based on bounding the Bures distance [22]

$$D_B(|\psi(t)\rangle, |\tilde{\psi}(t)\rangle) := \sqrt{2 \left(1 - \left| \langle \psi(t) | \tilde{\psi}(t) \rangle \right| \right)} \quad (1)$$

between two quantum states $|\psi(t)\rangle = U(t)|\psi_0\rangle$ and $|\tilde{\psi}(t)\rangle = \tilde{U}(t)|\psi_0\rangle$. The unitary transformations $U(t)$ and $\tilde{U}(t)$ are generated by the Hamiltonians $H(t)$ and $\tilde{H}(t)$, respectively and $|\psi_0\rangle$ is the initial state of the system. **While one could consider other metrics, we adopt the Bures distance given its operational meaning: the Bures distance upper and lower bounds the likelihood to distinguish between two states given any possible measurement on the system [22, 23].**

Along the lines of Ref. [20], in Appendix A we show that Bures distance can be upper bounded by

$$D_B(|\psi(t)\rangle, |\tilde{\psi}(t)\rangle) \leq \int_0^t dt' \left\| (H(t') - \tilde{H}(t')) \tilde{U}(t') |\psi_0\rangle \right\|, \quad (2)$$

where $\|\cdot\|$ denotes the Euclidean vector norm. Thus, if we choose $\tilde{H}(t) = \langle \psi_0 | H(t) | \psi_0 \rangle \mathbb{1}$ so that $\tilde{U}(t)$ becomes a global phase, we obtain a lower bound

$$T \geq \frac{D_B(|\psi_T\rangle, |\psi_0\rangle)}{\frac{1}{T} \int_0^T \sqrt{\text{Var}_{|\psi_0\rangle}(H(t))} dt}, \quad (3)$$

on the time T to prepare a target state $|\psi_T\rangle \equiv |\psi(T)\rangle$. The bound depends on the variance $\text{Var}_{|\psi_0\rangle}(H(t))$ of $H(t)$

with respect to the initial state $|\psi_0\rangle$. The speed limit in Eq. (3) resembles the one derived by Mandelstam-Tamm in Ref. [10]. As alluded to in the introduction, such bounds are of limited use for time-dependent Hamiltonians since they depend on the control or annealing schedule that prepares the desired target state.

A. Schedule-independent speed limits on quantum annealing

To derive schedule-independent speed limits that apply to quantum annealing problems, we leverage a technique introduced in [20]. Consider a time-dependent Hamiltonian of the form

$$H(t) = f(t)H_i + g(t)H_f. \quad (4)$$

The goal in quantum annealing is to design control schedules described by the functions $f(t)$ and $g(t)$ that prepare the ground state $|\psi_T\rangle = |E_{\min}^f\rangle$ of the Hamiltonian H_f at some time T . The system starts in an initial state $|\psi_0\rangle$, which is typically assumed to be an easy-to-prepare ground state of the Hamiltonian H_i . Leveraging the bound (2), we find that the time T to prepare the ground state is lower bounded by

$$T \geq \max \left\{ \frac{D_B(|\psi_T\rangle, |\psi_0\rangle)}{g_{\max} \sqrt{\text{Var}_{|\psi_0\rangle}(H_f)}}, \frac{D_B(|\psi_T\rangle, |\psi_0\rangle)}{f_{\max} \sqrt{\text{Var}_{|\psi_T\rangle}(H_i)}} \right\}. \quad (5)$$

Here, $g_{\max} = \max_{t \in [0, T]} |g(t)|$ and $f_{\max} = \max_{t \in [0, T]} |f(t)|$ are the largest amplitudes of the control functions $g(t)$ and $f(t)$, respectively. The first expression is proved by picking $\tilde{H}(t) = f(t)H_i + g(t)\langle \psi_0 | H_f | \psi_0 \rangle \mathbb{1}$ in (2), in which case the unitary $\tilde{U}(t)$ applied to an initial eigenstate of H_i gives a phase that does not affect Bures distance $D_B(|\psi_T\rangle, |\psi_0\rangle)$. The second bound can be proven analogously by considering the reverse-annealing process. **(Note that we will improve upon these bounds in the next section.)**

As suggested by the bound (5), if the control functions are unconstrained, preparing the ground state of H_f can be achieved instantaneously [24], i.e., the ground state can be prepared to arbitrary precision in an arbitrarily short time. Physically, however, energy constraints limit how large $g(t)$ and $f(t)$ can be. Interestingly, even if $f(t)$ is unconstrained but $g(t)$ is limited in amplitude by g_{\max} , the ground state cannot be prepared faster than the lower bound (5) implies, and vice-versa. We further note that, unlike Eq. (3), the lower bounds in Eq. (5) are independent of the driving protocol [i.e., independent of the control functions $f(t)$ and $g(t)$].

The lower bound on T scales with the variances of H_f or H_i with respect to the initial or final states. Thus, if one of the variances decreases with system size, the algorithm’s runtime must increase accordingly. For instance,

if we assume that the initial state is an equal superposition $|\psi_0\rangle = \frac{1}{\sqrt{d}} \sum_{j=1}^d |E_j^f\rangle$ of all eigenstates $|E_j^f\rangle$ of the Hamiltonian H_f , then the bound (5) implies

$$T \geq \frac{\sqrt{2d(1 - \frac{1}{\sqrt{d}})}}{g_{\max} \sqrt{\text{Tr}\{H_f^2\}}} \quad (6)$$

where d is the dimension of the Hilbert space. Thus, the time T to prepare the ground state must at least scale as $\mathcal{O}(\sqrt{d})$ when $\text{Tr}\{H_f^2\} = \mathcal{O}(1)$ remains constant when the system is scaled. This is, for example, the case when $\text{Tr}\{H_f^2\} = 1$ is used as a normalization condition [25, 26]. The $\mathcal{O}(\sqrt{d})$ scaling obtained here is similar to the one found in Ref. [27], where system size-dependent speed limits were obtained for quantum control by estimating the diameter of the quotient spaces that exclude operations that can be achieved instantaneously through control. This scaling was also obtained in Refs. [5–7] for the analog version of Grover’s search algorithm, which we consider next to probe the new bound (5).

B. Case study: analog Grover search

The aim of Grover’s search algorithm is to prepare the state $|m\rangle$ where $m \in \{0, 1\}^n$ is a bitstring of size n that corresponds to the solution of an (unstructured) search problem. The n -qubit system is initialized in the equal-superposition state $|\psi_0\rangle = \frac{1}{\sqrt{d}} \sum_{x \in \{0, 1\}^n} |x\rangle$, where $d = 2^n$. The initial and target states are minimum energy eigenstates of the Hamiltonians $H_i = \mathbb{1} - |\psi_0\rangle\langle\psi_0|$ and $H_f = \mathbb{1} - |m\rangle\langle m|$, respectively. Roland and Cerf proved that an optimized annealing schedule gives a quadratic improvement in runtime [28], i.e., $T = \mathcal{O}(\sqrt{d})$, matching the quantum improvement of Grover’s search algorithm over classical algorithms [29]. The optimized annealing schedule derived by Roland and Cerf saturates the lower bound derived in Ref. [6], which shows that going beyond the adiabatic schedule cannot improve the \sqrt{d} scaling.

Assuming ideal annealing where the target state is given by $|\psi_T\rangle = |m\rangle$, and using that $\langle m|\psi_0\rangle = \frac{1}{\sqrt{d}}$, Eq. (5) gives the lower bound

$$T \geq \max \left\{ \frac{1}{g_{\max}}, \frac{1}{f_{\max}} \right\} \frac{\sqrt{2(1 - \frac{1}{\sqrt{d}})}}{\sqrt{\frac{1}{d} - \frac{1}{d^2}}} = \mathcal{O}(\sqrt{d}), \quad (7)$$

on the time needed to prepare the target state $|m\rangle$ that yields the solution to the search problem. This proves that even if one of the control fields $f(t)$ or $g(t)$ could be made arbitrarily strong, the $\mathcal{O}(\sqrt{d})$ scaling cannot be beaten. The minimum computation time is bounded by the weakest of the control fields. In the case where M bitstrings are solutions to the search problem, i.e., $H_i = -\sum_{j=1}^M |j\rangle\langle j|$, $j \in \{0, 1\}^n$, for $d \gg M$ the bound becomes $T \gtrsim \max \left\{ \frac{1}{g_{\max}}, \frac{1}{f_{\max}} \right\} \sqrt{\frac{d}{M}}$.




FIG. 1. Numerical optimization using BFGS to determine the optimal annealing schedule and corresponding annealing times T for the analog version of Grover search. The simulations were performed using 100 piecewise constant control field values averaged over 100 randomly chosen initial control values. The colormap in (a) shows the fidelity error ϵ as a function of the system size d and T for a fixed maximum control amplitude $g_{\max} = 1$, while the control field $f(t)$ was assumed to be unconstrained. The colormap in (b) shows ϵ as function of g_{\max} and T for fixed $d = 10$. In both plots, the bound in Eq. (7) is shown as a black curve.

In Fig. 1, we compare the exact annealing times with bound (7). The exact annealing times were obtained by numerically optimizing the fidelity error for preparing the ground state using the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm. In the numerical simulations, we assumed that $f(t)$ is unconstrained and used 100 piecewise constant control amplitudes over which we optimize, reporting the average over 100 randomly chosen initial control field values. The results shown in Fig. 1 illustrate that the bound (7) shown as a solid black curve is remarkably tight. While asymptotically tight bounds on search by annealing were also derived in Refs. [5, 6], Eq. (7) is slightly more general as it does not restrict the

magnitude of the schedule functions.

While the bound (5) captures the annealing time in search problems, it does not accurately reflect the times needed to perform annealing on systems with geometric locality (e.g., spin chains). This is because the Hamiltonian's variances typically grow with system size [30]. This shortcoming is shared with lower bounds on annealing times previously considered in the literature [5–7]. Next, we partially address this shortcoming by deriving bounds on the annealing time that, in certain instances, grow with system size.

III. TIGHTER BOUNDS ON ANNEALING TIMES FOR SPIN MODELS

The schedule- and path-independent bound (5) is a special case of a broader class of results that aim to bound quantum control times (i.e., the time to prepare a desired target state or unitary transformation through control fields) by considering choices of $\tilde{H}(t)$ that make the system the most uncontrollable [31]. In fact, the derivation of the bound (5) rests on comparing the exact evolution $U(t)$ with an evolution $\tilde{U}(t)$ where the control term $g(t)H_f$ is removed. Motivated by the results in [21], we now consider other choices of $\tilde{H}(t)$ to derive tighter bounds on annealing times.

We consider a rotation of $H(t)$ by a unitary transformation W [21], such that $[W, H_i] = 0$. That is, $\tilde{H}(t) = W^\dagger H(t) W = g(t)W^\dagger H_f W + f(t)H_i$, and $W|\psi_0\rangle = e^{i\phi}|\psi_0\rangle$. By relying on the bound (2), we prove in Appendix B that

$$T \geq \frac{D_B(|\psi_T\rangle, W|\psi_0\rangle)}{g_{\max}\|H_f - WH_fW^\dagger\|} = \frac{\sqrt{2}\sqrt{1 - |\langle\psi_T|W|\psi_0\rangle|}}{g_{\max}\|[H_f, W]\|}, \quad (8)$$

where here $\|A\|$ denotes the spectral norm of an operator A , and we used that $\|H_f - WH_fW^\dagger\| = \|[H_f, W]\|$ since the spectral norm is unitarily invariant. An analogous result holds by choosing $\tilde{H}(t) = V^\dagger H(t)V$, with $[V, H_f] = 0$ and $V|\psi_T\rangle = e^{i\varphi}|\psi_T\rangle$. We find, $T \geq \frac{D_B(|\psi_0\rangle, V|\psi_T\rangle)}{f_{\max}\|H_i - V^\dagger H_i V\|}$ where $f_{\max} = \max_{t \in [0, T]} |f(t)|$. Combining both bounds yields

$$T \geq \sqrt{2} \max \left\{ \frac{\sqrt{1 - |\langle\psi_0|V|\psi_T\rangle|}}{f_{\max}\|[H_i, V]\|}, \frac{\sqrt{1 - |\langle\psi_T|W|\psi_0\rangle|}}{g_{\max}\|[H_f, W]\|} \right\}. \quad (9)$$

In Appendix C, we leverage these results to bound the circuit depth in QAOA.

The bounds in (9) do not depend on the state's path or the annealing schedule. Evaluating them only requires knowledge of (i) the Hamiltonians H_i and H_f that drive the quantum annealing algorithm, (ii) the initial or final states $|\psi_0\rangle$ or $|\psi_T\rangle$, and (iii) the free-to-choose unitaries V or W . Next, we show that properly chosen W or V yields speed limits for spin networks that are tighter than previously known bounds [6].

A. Bounds on annealing times for spin networks

We consider an n -spin 1/2 network on a graph $G(V, E)$ with edges E and vertices V . We denote by $|V| = n$ and $|E|$ the number of vertices and edges in the graph, respectively. For concreteness, we consider 2-local systems, where the spin network is described by the Hamiltonian,

$$H_f = \mathcal{N} \sum_{\alpha, \beta \in \{x, y, z\}} \sum_{(i, j) \in E} h_{\alpha, \beta}^{(i, j)} \sigma_\alpha^{(i)} \sigma_\beta^{(j)}. \quad (10)$$

Here, $\sigma_\alpha^{(i)}$ denote Pauli operators $\alpha \in \{x, y, z\}$ acting on the i th spin, $h_{\alpha, \beta}^{(i, j)}$ are coupling constants between spins (i) and (j) , and \mathcal{N} is a normalization constant. We denote the largest coupling constant in the system by $h_{\max} = \max\{|h_{\alpha, \beta}^{(i, j)}|\}$. The solutions to optimization problems described by the graph $G(V, E)$ can be encoded in the ground states of Hamiltonians like (10) [32].

For the initial Hamiltonian, we take

$$H_i = - \sum_{j \in V} \sigma_x^{(j)}, \quad (11)$$

so that the initial state is $|\psi_0\rangle = |+\rangle^{\otimes n}$, the ground state of H_i , where $|+\rangle$ is an eigenstate of σ_x . We aim to leverage the bound in Eq. (8) to lower bound the time T needed to prepare the ground state of H_f via the time-dependent Hamiltonian $H(t) = f(t)H_i + g(t)H_f$. For simplicity we assume $f_{\max} = g_{\max} = 1$.

In general, the time to prepare the ground state of H_f depends on the magnitudes of the Hamiltonians used to drive the dynamics. For instance, rescaling a Hamiltonian to $H' = aH$ with $a > 1$ trivially decreases time to $T' = T/a$. As the bounds (8) and (9) hold regardless of the rescaling of the Hamiltonian, it is convenient to fix a magnitude of the H s when comparing the scaling of T with the system size. For spin networks, we adopt the convention that all Hamiltonians are extensive [33]. That is, $\|H_i\| = n$ and $\|H_f\| = \mathcal{O}(n)$. For geometrically local Hamiltonians, this is ensured by picking the normalization constant \mathcal{N} in H_f to be $\mathcal{N} = \frac{|V|}{|E|} = \frac{n}{|E|}$. This choice ensures that $\|H(t)\| \sim \|H_i\| \sim \|H_f\| = \mathcal{O}(n)$, and that doubling the number of spins in the network approximately doubles the system's energy.

Taking $W = \sigma_\alpha^{(i)}$ in Eq. (8), where i identifies the vertex with the smallest degree δ , it holds that $\|[H_f, W]\| \leq 6 h_{\max} \delta \mathcal{N} = 6 h_{\max} \delta \frac{|V|}{|E|}$, which yields

$$T \geq \frac{\sqrt{2}\sqrt{1 - |\langle\psi_T|W|\psi_0\rangle|} |E|}{6 h_{\max} \delta |V|}. \quad (12)$$

To illustrate the bound's scaling, consider an almost-complete graph where $|E| = \mathcal{O}(n^2)$ and $|V| = n$, where one node is minimally connected, e.g., $\delta = 1$. Then, the time to prepare the ground state must at least increase linearly with the number of spins n [34]. That

is, the lower bound on T grows with n for sufficiently connected graphs where $|E|/|V|$ grows, as long as the smallest vertex degree δ is small and $|\langle\psi_T|W|\psi_T\rangle|$ is sufficiently small (as we show below, one can often choose W such that $|\langle\psi_T|W|\psi_T\rangle| = 0$). The bound (12) may thus be ideally suited to study long-range systems [35, 36], such as the SYK model [37, 38], where the corresponding graphs have high connectivity.

For instance, consider the minimum time to find the ground state of an Ising Hamiltonian $H_f = |V|/|E| \sum_{(i,j) \in E} h_{z,z}^{(i,j)} \sigma_z^{(i)} \sigma_z^{(j)}$ where the smallest vertex degree is $\delta = 1$. The ground state is given by some bit string $|\psi_T\rangle = |m\rangle$ with $m \in \{0,1\}^n$, so $\langle\psi_T|W|\psi_T\rangle = 0$ for $W = \sigma_x^{(i)}$. Thus, Eq. (12) becomes $T \geq \frac{\sqrt{2}}{h_{\max}} \frac{|E|}{|V|}$, which grows with the graph's connectivity.

We can use the previous example to show that the bound in Eq. (9) is typically tighter than the only schedule-independent bound derived in Ref. [6]. (Note that the tightest bound derived in Ref. [6], which involves the state's coherence along the annealing schedule, depends on the time-evolved state of the system.) The schedule-independent bound in Ref. [6] is

$$T \geq \tau_{\text{anneal3}} = \frac{\langle H_i \rangle_T - \langle H_i \rangle_0 + \langle H_f \rangle_0 - \langle H_f \rangle_T}{\| [H_f, H_i] \|}, \quad (13)$$

where $\langle H_i \rangle_T = \langle \psi_T | H_i | \psi_T \rangle$, $\langle H_f \rangle_T = \langle \psi_T | H_f | \psi_T \rangle$, $\langle H_i \rangle_0 = \langle \psi_0 | H_i | \psi_0 \rangle$ and $\langle H_f \rangle_0 = \langle \psi_0 | H_f | \psi_0 \rangle$. For the example considered in the previous paragraph, $\| [H_f, H_i] \| \approx h_{\max} |V| = h_{\max} n$. Moreover, $\langle H_i \rangle_T - \langle H_i \rangle_0 = \mathcal{O}(n)$ and $\langle H_f \rangle_0 - \langle H_f \rangle_T = \mathcal{O}(n)$. Thus, the bound (13) yields $T \geq \tau_{\text{anneal3}} = \mathcal{O}(1)$. Next, we show a second example where the new bound (9) is tighter than the ones previously considered in the literature.

B. Case study: a perturbed p -spin model

The p -spin model is described by the Hamiltonians

$$H_i = n \left(\mathbb{1} - \frac{M_x}{n} \right) \quad \text{and} \quad H_f = n \left(\mathbb{1} - \frac{M_z^p}{n^p} \right), \quad (14)$$

where $M_z = \sum_{j=1}^n \sigma_z^{(j)}$ and $M_x = \sum_{j=1}^n \sigma_x^{(j)}$ are the spins' magnetizations along z and x , respectively [39, 40].

The free parameter p determines the energy gap between the ground and first excited states which, in turn, determines the timescales to anneal adiabatically [41]: $\tau_{\text{adiabatic}} \sim \text{poly}(n)$ for $p = 2$ and $\tau_{\text{adiabatic}} \sim \exp(n)$ for $p \geq 3$ [41]. Optimized schedules significantly beat the times to anneal adiabatically; a depth-2 QAOA can prepare the target state on a timescale $T = \mathcal{O}(n^{p-1})$ [42].

Consider a perturbed p -spin model where an extra, low-connectivity vertex is added to the final Hamiltonian,

$$H'_f = n \left(\mathbb{1} - \frac{M_z^p + \lambda \sigma_z^{(n)} \sigma_z^{(n+1)}}{n^p} \right), \quad (15)$$

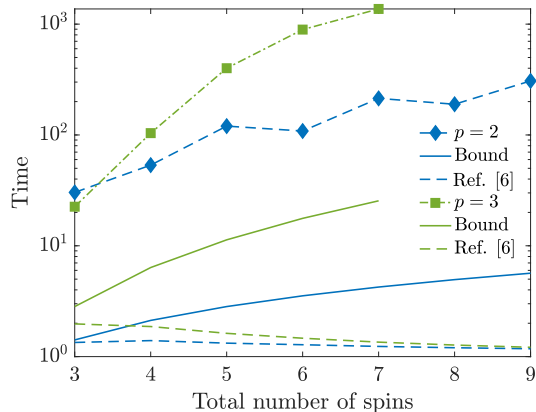


FIG. 2. Annealing times to prepare the ground state of the p -spin model with an extra node, as described by Eq. (15) for $\lambda = 1$, as a function of the number of total spins. We consider $p = 2$ (blue diamonds) and $p = 3$ (green squares). The annealing times were obtained by numerically optimizing the fidelity to prepare the ground state using BFGS, reporting the smallest time for which the ground state is prepared with a fidelity $F > 0.99$. The corresponding bounds given in Eq. (16) are shown as blue and green lines, with $g_{\max} = 1$. The schedule-independent bound in Ref. [6] is shown for this example as a dashed line.

where $\lambda \in \mathbb{R}$. The extra vertex acts like a perturbation $\propto 1/n^{p-1}$ that does not significantly change the Hamiltonian's norm. We assume the initial Hamiltonian is extended to act on the appended vertex.

Choosing $W = \sigma_x^{(n+1)}$ in (8) yields the lower bound

$$T \geq \frac{n^{p-1}}{\sqrt{2} g_{\max} \lambda}, \quad (16)$$

for the annealing time T to prepare the ground state of the modified p -spin Hamiltonian (15). For the p -spin model, the bound (13) gives a constant scaling $\tau_{\text{anneal3}} = \mathcal{O}(1)$ [6] (the perturbation to the p -spin Hamiltonian does not affect the bound's scaling). Eq. (16) thus shows a polynomial improvement.

In Fig. 2 we compare the bound in Eq. (16) to the times corresponding to an optimized QAOA schedule for the perturbed p -spin model [Eq. (15)]. QAOA consists of sequentially applying gates $e^{-i\beta_j H_i}$ and $e^{-i\gamma_j H_f}$ for $j = \{1, \dots, L\}$, where the angles β_j and γ_j are optimized to approach the target ground state of H_f [11]. We refer to Appendix C for further details. The runtimes $T = \sum_j |\beta_j| + |\gamma_j|$ to prepare the ground state were obtained by numerically maximizing the fidelity $F = \langle \psi(T) | P | \psi(T) \rangle$ using the BFGS algorithms where P is the projector onto the ground state for different annealing times T . We additionally optimized over 100 randomly chosen initial angles and report the smallest T for which the ground state is prepared with fidelity $F > 0.99$ (green squares and blue diamonds). From the results shown in Fig. 2, we see that the bound (16) cap-

tures the scaling in system size of the annealing time.

The scaling in (16) is a result of the perturbation added to the p -spin model. The perturbation adds a minimally connected vertex to a well-connected graph, the scenario that we identified after (12) as the most favorable for the bound derived in this work. It is not surprising that the minimum timescale may be determined by the weakest link since, in this case, the state of the appended spin can only evolve due to the perturbation. But, we highlight that previous bounds on annealing times did not capture such behaviour, as their scaling was determined by more global properties of the defining Hamiltonians (e.g., as in Eq. (13)).

IV. ON THE TIGHTNESS OF THE BOUNDS

This article’s main goal is to derive bounds on the minimum annealing times that

- (a) can be evaluated without knowledge of the annealers’ schedule or its state along a computation, and
- (b) are tighter than those previously known bounds.

The bounds in Eq. (5) of Sec. II focus on addressing goal (a). Evaluating Eq. (5) requires no information of an annealing algorithm’s schedule or the evolution of the annealer’s state. Moreover, the bounds are saturated for search algorithms. However, such bounds are loose for spin chain models. To see this, note that the denominators *typically* scale as $\text{var}_{|\psi_0\rangle}(H_f) \sim \text{var}_{|\psi_T\rangle}(H_i) \sim n$ for an n -body spin chain with local interactions and decaying correlations (this can be evaluated directly for a the former variance in a known initial state). However, one expects that the annealing time grows with system size. Thus, Eq. (5) can be very loose, e.g., for optimization problems encoded in spin chain Hamiltonians.

In response to this shortcoming, the main results of this article [Eq. (9) in Sec. III] take steps to address goal (b), too. Specifically, the bounds in Eq. (9) can be tighter than those in Eq. (5) in situations when the latter are very loose, such as optimization problems encoded in spin chain Hamiltonians. As Fig. 2 illustrates, our new bounds in Eq. (9) are polynomially tighter than previously known ones [and, for this model, also significantly tighter than Eq. (5)].

Admittedly, the bounds in Eq. (9) are tighter when one considers a model that includes a perturbative term. In such a case, the bounds are able to capture the scaling of the runtime with the weakest term in the Hamiltonian (previous bounds in the literature were not able to

capture such behavior). In other cases, the new bounds are loose. Nevertheless, our results add to a growing set of tools to analyze annealing times without the need to solve rarely tractable dynamics for large-scale problems.

V. DISCUSSION

In this work, we have taken steps to address shortcomings of previous bounds on annealing times considered in the literature. Our main result, bound (9), can take into consideration a problem’s locality structure and, for optimization problems characterized by a well-connected graph with a low-connected vertex, displays a polynomial improvement over previous results. To illustrate regimes where the bounds obtained here are tighter than previous ones, we considered a perturbed annealing model where the bound (9)’s scaling depends on the perturbation. Along similar lines, it has been noted in Ref. [43] that, if uncorrected, perturbations can strongly influence annealing processes.

Our results thus add to a series of recent works deriving runtime bounds on computing times from the physical models used to compute. When sufficiently tight, these sort of bounds can be used to benchmark concrete algorithms; an algorithm that computes saturating a bound is optimal.

In the literature of speed limits, it has been noted that the bounds often depend on the metrics chosen [15, 44]. It would be interesting to explore if the bounds derived in this letter can be improved upon by exploring other metrics in state space or observable-based.

ACKNOWLEDGMENTS

L.P.G.P. acknowledges support from the Beyond Moore’s Law project of the Advanced Simulation and Computing Program at Los Alamos National Laboratory (LANL) managed by Triad National Security, LLC, for the National Nuclear Security Administration of the U.S. DOE under contract 89233218CNA000001, the Laboratory Directed Research and Development (LDRD) program of LANL under project number 20230049DR, and the U.S. Department of Energy, Office of Advanced Scientific Computing Research, Accelerated Research for Quantum Computing program, Fundamental Algorithmic Research for Quantum Computing (FAR-QC) project. C.A. acknowledges support from the National Science Foundation (Grant No. 2231328).

[1] 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).

[2] T. Albash and D. A. Lidar, Adiabatic quantum computation, *Rev. Mod. Phys.* **90**, 015002 (2018).

- [3] S. Jansen, M.-B. Ruskai, and R. Seiler, Bounds for the adiabatic approximation with applications to quantum computation, *Journal of Mathematical Physics* **48**, [10.1063/1.2798382](#) (2007).
- [4] S. Boixo and R. D. Somma, Necessary condition for the quantum adiabatic approximation, *Phys. Rev. A* **81**, [032308](#) (2010).
- [5] J.-H. Chen, Lower bounds for adiabatic quantum algorithms by quantum speed limits, *Phys. Rev. Res.* **5**, [033175](#) (2023).
- [6] L. P. García-Pintos, L. T. Brady, J. Bringewatt, and Y.-K. Liu, Lower bounds on quantum annealing times, *Phys. Rev. Lett.* **130**, [140601](#) (2023).
- [7] N. Benchasattabuse, A. Bärtshi, L. P. García-Pintos, J. Golden, N. Lemons, and S. Eidenbenz, Lower bounds on number of QAOA rounds required for guaranteed approximation ratios, arXiv preprint [arXiv:2308.15442](#) [10.48550/arXiv.2308.15442](#) (2023).
- [8] S. Lloyd, Ultimate physical limits to computation, *Nature* **406**, [1047](#) (2000).
- [9] S. Deffner and S. Campbell, Quantum speed limits: from heisenberg's uncertainty principle to optimal quantum control, *J. Phys. A: Math. Theor.* **50**, [453001](#) (2017).
- [10] L. Mandelstam and I. Tamm, The uncertainty relation between energy and time in nonrelativistic quantum mechanics, *J. Phys.(USSR)* **9**, [1](#) (1945).
- [11] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, arXiv preprint [arXiv:1411.4028](#) (2014).
- [12] M. M. Taddei, B. M. Escher, L. Davidovich, and R. L. de Matos Filho, Quantum speed limit for physical processes, *Phys. Rev. Lett.* **110**, [050402](#) (2013).
- [13] A. del Campo, I. L. Egusquiza, M. B. Plenio, and S. F. Huelga, Quantum speed limits in open system dynamics, *Phys. Rev. Lett.* **110**, [050403](#) (2013).
- [14] S. Deffner and E. Lutz, Quantum speed limit for non-markovian dynamics, *Phys. Rev. Lett.* **111**, [010402](#) (2013).
- [15] D. P. Pires, M. Cianciaruso, L. C. Céleri, G. Adesso, and D. O. Soares-Pinto, Generalized geometric quantum speed limits, *Phys. Rev. X* **6**, [021031](#) (2016).
- [16] L. P. García-Pintos, S. B. Nicholson, J. R. Green, A. del Campo, and A. V. Gorshkov, Unifying quantum and classical speed limits on observables, *Phys. Rev. X* **12**, [011038](#) (2022).
- [17] S. B. Nicholson, L. P. García-Pintos, A. del Campo, and J. R. Green, Time-information uncertainty relations in thermodynamics, *Nat. Phys.* **16**, [1211](#) (2020).
- [18] V. T. Vo, T. Van Vu, and Y. Hasegawa, Unified approach to classical speed limit and thermodynamic uncertainty relation, *Phys. Rev. E* **102**, [062132](#) (2020).
- [19] E. Aghion and J. R. Green, Thermodynamic speed limits for mechanical work, *J. Phys. A: Math. Theor.* **56**, [05LT01](#) (2023).
- [20] C. Arenz, B. Russell, D. Burgarth, and H. Rabitz, The roles of drift and control field constraints upon quantum control speed limits, *New J. Phys.* **19**, [103015](#) (2017).
- [21] J. Lee, C. Arenz, H. Rabitz, and B. Russell, Dependence of the quantum speed limit on system size and control complexity, *New J. Phys.* **20**, [063002](#) (2018).
- [22] C. A. Fuchs, Distinguishability and accessible information in quantum theory, arXiv preprint [quant-ph/9601020](#) (1996).
- [23] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information* (Cambridge university press, 2010).
- [24] D. d'Alessandro, *Introduction to quantum control and dynamics* (CRC press, 2021).
- [25] B. Craps, M. D. Clerck, O. Evnin, P. Hacker, and M. Pavlov, Bounds on quantum evolution complexity via lattice cryptography, *SciPost Phys.* **13**, [090](#) (2022).
- [26] B. Craps, M. D. Clerck, O. Evnin, and P. Hacker, Integrability and complexity in quantum spin chains, *SciPost Phys.* **16**, [041](#) (2024).
- [27] M. T. Johnsson, L. van Luijk, and D. Burgarth, Exact and lower bounds for the quantum speed limit in finite-dimensional systems, *Phys. Rev. A* **108**, [052403](#) (2023).
- [28] J. Roland and N. J. Cerf, Quantum search by local adiabatic evolution, *Phys. Rev. A* **65**, [042308](#) (2002).
- [29] L. K. Grover, A fast quantum mechanical algorithm for database search, in *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96 (Association for Computing Machinery, New York, NY, USA, 1996) p. 212–219.
- [30] M. C. Bañuls, D. A. Huse, and J. I. Cirac, Entanglement and its relation to energy variance for local one-dimensional hamiltonians, *Phys. Rev. B* **101**, [144305](#) (2020).
- [31] D. Burgarth, J. Borggaard, and Z. Zimborás, Quantum distance to uncontrollability and quantum speed limits, *Phys. Rev. A* **105**, [042402](#) (2022).
- [32] A. Das and B. K. Chakrabarti, *Quantum annealing and related optimization methods*, Vol. 679 (Springer Science & Business Media, 2005).
- [33] R. C. Tolman, The measurable quantities of physics, *Physical Review* **9**, [237](#) (1917).
- [34] We highlight that the choice of normalization can change the bound on T , but the conclusions stand when one compares dimensionless quantities such as $T\|H\|$.
- [35] M. C. Tran, C.-F. Chen, A. Ehrenberg, A. Y. Guo, A. Deshpande, Y. Hong, Z.-X. Gong, A. V. Gorshkov, and A. Lucas, Hierarchy of linear light cones with long-range interactions, *Phys. Rev. X* **10**, [031009](#) (2020).
- [36] T. Kuwahara and K. Saito, Strictly linear light cones in long-range interacting systems of arbitrary dimensions, *Phys. Rev. X* **10**, [031010](#) (2020).
- [37] S. Sachdev and J. Ye, Gapless spin-fluid ground state in a random quantum heisenberg magnet, *Phys. Rev. Lett.* **70**, [3339](#) (1993).
- [38] H. Schlömer and S. Sachdev, [Quantum annealing with chaotic driver hamiltonians](#) (2024), [arXiv:2409.20538](#) [quant-ph].
- [39] Y. Seki and H. Nishimori, Quantum annealing with antiferromagnetic fluctuations, *Phys. Rev. E* **85**, [051112](#) (2012).
- [40] B. Seoane and H. Nishimori, Many-body transverse interactions in the quantum annealing of the p-spin ferromagnet, *J. Phys. A: Math. Theor.* **45**, [435301](#) (2012).
- [41] V. Bapst and G. Semerjian, On quantum mean-field models and their quantum annealing, *J. Stat. Mech-Theory E* **2012**, [P06007](#) (2012).
- [42] 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 in a transverse field, *Phys. Rev. A* **102**, [062404](#) (2020).

[43] A. Pearson, A. Mishra, I. Hen, and D. A. Lidar, Analog errors in quantum annealing: doom and hope, npj Quantum Information **6**, [10.1038/s41534-019-0210-7](https://doi.org/10.1038/s41534-019-0210-7) (2020).

[44] J. Bringewatt, Z. Steffen, M. A. Ritter, A. Ehrenberg, H. Wang, B. Palmer, A. J. Kollár, A. V. Gorshkov, and L. P. García-Pintos, Generalized geometric speed limits for quantum observables, arXiv preprint arXiv:2409.04544 [10.48550/arXiv.2409.04544](https://doi.org/10.48550/arXiv.2409.04544) (2024).

Appendix A: Upper bounding the Bures distance

Here, we establish the upper bound for Bures distance given in (2) in the main body of the manuscript.

We consider two quantum states $|\psi(t)\rangle = U(t)|\psi_0\rangle$ and $|\tilde{\psi}(t)\rangle = \tilde{U}(t)|\psi_0\rangle$ whose unitary evolutions $U(t)$ and $\tilde{U}(t)$ are generated by the Hamiltonians $H(t)$ and $\tilde{H}(t)$, respectively. Since [20],

$$U^\dagger(t)\tilde{U}(t) - \mathbb{1} = i \int_0^t U^\dagger(t')(H(t') - \tilde{H}(t'))\tilde{U}(t') dt', \quad (\text{A1})$$

the Euclidian distance between the two quantum states is given by

$$\| |\tilde{\psi}(t)\rangle - |\psi(t)\rangle \| = \|(U^\dagger(t)\tilde{U}(t) - \mathbb{1})|\psi_0\rangle\| = \left\| i \int_0^t U^\dagger(t')(H(t') - \tilde{H}(t'))\tilde{U}(t')|\psi_0\rangle dt' \right\|,$$

where we used unitary invariance of the Euclidian vector. As $\| |\tilde{\psi}(t)\rangle - |\psi(t)\rangle \| \geq D_B(|\psi(t)\rangle, |\tilde{\psi}(t)\rangle)$ is lower bounded by Bures distance D_B , upper bounding the right-hand side and using unitary invariance again establishes the desired result.

$$D_B(|\psi(t)\rangle, |\tilde{\psi}(t)\rangle) \leq \left\| i \int_0^t U^\dagger(t')(H(t') - \tilde{H}(t'))\tilde{U}(t')|\psi_0\rangle dt' \right\| \leq \int_0^t dt' \|(H(t') - \tilde{H}(t'))\tilde{U}(t')|\psi_0\rangle\| \quad (\text{A2})$$

Appendix B: Deriving the bound (8)

To derive the bound (8) in the main text, we further upper the right-hand side of (A2) to obtain

$$D_B(|\psi(t)\rangle, |\tilde{\psi}(t)\rangle) \leq \int_0^t \|(H(t') - \tilde{H}(t'))\| dt', \quad (\text{B1})$$

where here $\|\cdot\|$ denotes the spectral norm. Choosing $\tilde{H}(t) = W^\dagger H(t)W$ where the unitary W satisfies $[W, H_i] = 0$ and $W|\psi_0\rangle = e^{i\phi}|\psi_0\rangle$ gives

$$|\langle \tilde{\psi}(t) | \psi(t) \rangle| = |\langle \psi_0 | W^\dagger U^\dagger(t) W U(t) | \psi_0 \rangle| = |\langle \psi(t) | W | \psi(t) \rangle| \quad (\text{B2})$$

the desired result.

Appendix C: Bounding the number of layers in QAOA

QAOA consists of sequentially applying gates $e^{-i\beta_j H_i}$ and $e^{-i\gamma_j H_f}$ for $j = \{1, \dots, L\}$ [11]. The angles β_j and γ_j are usually optimized through a classical optimization routine, in tandem with a quantum computer to minimize the expectation value of the Hamiltonian H_f . The goal is to prepare a state close to the ground state of H_f with as few layers L as possible.

QAOA can be interpreted as an annealing problem where the annealing schedule is of the form $f(t) = 1 - g(t)$, and $g(t)$ switches between the values $g(t) \in \{0, 1\}$ to alternate between applying the unitaries $e^{-i\beta_j H_i}$ and $e^{-i\gamma_j H_f}$. In this setting, the angles in QAOA correspond to the time intervals each Hamiltonian acts for in the quantum annealing algorithm.

Following the steps taken in Ref. [7], the lower bounds (8) and (9) on annealing times imply lower bounds on the number of layers L needed for QAOA to prepare the desired ground state. We first note that a bound on annealing times implies a bound on the sum of QAOA angles. From (9) we obtain

$$\sum_{j=1}^L (|\beta_j| + |\gamma_j|) \geq \max \left\{ \frac{D_B(|\psi_0\rangle, V|\psi_0\rangle)}{\|[H_i, V]\|}, \frac{D_B(|\psi_T\rangle, W|\psi_T\rangle)}{\|[H_f, W]\|} \right\}. \quad (\text{C1})$$

For periodic Hamiltonians with $\{|\beta_j|, |\gamma_j|\} \in [0, 2\pi]$ [7], the previous expression yields a lower bound on the number of layers

$$L \geq \max \left\{ \frac{D_B(|\psi_0\rangle, V|\psi_0\rangle)}{4\pi\|[H_i, V]\|}, \frac{D_B(|\psi_T\rangle, W|\psi_T\rangle)}{4\pi\|[H_f, W]\|} \right\} \quad (\text{C2})$$

needed to prepare the ground state of H_f .

The latter expression lower bounds the circuit depth of any QAOA that uses periodic Hamiltonians. It can be used to benchmark the optimization procedure: i.e., there is no more need to continue optimizing the angles β_j and γ_j for QAOA protocols where the system's state is sufficiently close to the target state.