Spectral Gap Amplification

R. D. Somma[1] and S. Boixo[2]

[1]*Los Alamos National Laboratory, Los Alamos, NM 87545, USA*[*]
[2] *Information Sciences Institute, USC, Marina del Rey, CA 90292,*
USA. Harvard University, Cambridge, MA 02138, USA.[†]
(Dated: October 15, 2018)

Many problems can be solved by preparing a specific eigenstate of some Hamiltonian H. The generic cost of quantum algorithms for these problems is determined by the inverse spectral gap of H for that eigenstate and the cost of evolving with H for some fixed time. The goal of spectral gap amplification is to construct a Hamiltonian H' with the same eigenstate as H but a bigger spectral gap, requiring that constant-time evolutions with H' and H are implemented with nearly the same cost. We show that a quadratic spectral gap amplification is possible when H satisfies a frustration-free property and give H' for these cases. This results in quantum speedups for optimization problems. It also yields improved constructions for adiabatic simulations of quantum circuits and for the preparation of projected entangled pair states (PEPS), which play an important role in quantum many-body physics. Defining a suitable black-box model, we establish that the quadratic amplification is optimal for frustration-free Hamiltonians and that no spectral gap amplification is possible, in general, if the frustration-free property is removed. A corollary is that finding a similarity transformation between a stoquastic Hamiltonian and the corresponding stochastic matrix is hard in the black-box model, setting limits to the power of some classical methods that simulate quantum adiabatic evolutions.

Keywords: Quantum Algorithms, Adiabatic Quantum Computing, Quantum Monte-Carlo

I. INTRODUCTION AND SUMMARY OF RESULTS

Many problems in physics and optimization, such as describing quantum phases of matter or solving satisfiability, can be reduced to the computation of low-energy states of Hamiltonians (see, e.g.[19, 31]). Methods to compute such states exist, with adiabatic state transformation (AST) being one of the most acknowledged and powerful heuristics for that goal [4, 18–20, 22, 24]. The problem in AST involves transforming a pure quantum state $|\psi_0\rangle$ into $|\psi_1\rangle$; these being the endpoints of a state path $|\psi_s\rangle$, $0 \le s \le 1$, where each $|\psi_s\rangle$ is an eigenstate of a Hamiltonian $H(s)$. Quantum strategies to solve the AST problem typically require evolving with the Hamiltonians for time $T > 0$, which determines the cost of the method. A well-known example is given by the quantum adiabatic theorem and requires changing the interaction parameter s sufficiently slowly in time [23, 27], but more efficient strategies exist [9, 10, 43]. When direct access to Hamiltonian evolutions is not possible, such evolutions can be simulated in the standard quantum circuit model by means of Trotter-Suzuki-like formulas [8, 29, 42]. For these cases, the cost of the method does not only depend on T, but also depends on the additional cost of approximating fixed-time Hamiltonian evolutions by quantum circuits. For this reason, it is standard to restrict to those $H(s)$ whose fixed-time evolutions can be approximated by quantum circuits of small size (e.g., size bounded by constant-degree polynomials).

Let L be the angular length determined by $\{|\psi_s\rangle\}_{0 \le s \le 1}$ and $\Delta_s \ge \Delta$ the spectral gap to the other nearest eigenvalue of $H(s)$. Recently, we provided fast quantum methods for eigenpath traversal that do not exploit the structure of the Hamiltonians and solve the AST problem in time $T \in \Omega(L/\Delta)$ [10]. Such methods were shown to be optimal in a black-box model [11]. However, the results in Ref. [11] do not forbid the existence of faster methods for solving the AST problem if additional knowledge about the structure of $H(s)$ is available. Our paper focuses on this observation and introduces the spectral gap amplification problem or GAP. Basically, we intend to answer: Can we construct a Hamiltonian $H'(s)$ that also has $|\psi_s\rangle$ as eigenstate, but spectral gap $\Delta'_s \gg \Delta_s$?

When spectral gap amplification is possible, it can be used as a technique to find new quantum speedups in both, the Hamiltonian-based model and the quantum circuit model. In Refs. [34, 36], for example, we constructed the so-called quantum simulated annealing algorithm that provided speedups of the well-known simulated annealing method implemented using Monte-Carlo. The reason for such speedups is a gap amplification step: if Δ is the spectral gap of the stochastic matrix S used in Monte-Carlo, there is a Hamiltonian H' with spectral gap $\Delta' \ge \sqrt{\Delta}$ and eigenstate that allows to sample from the fixed point of S. To build H' we used previously-developed tools for quantum walks [5, 38].

Motivated by the results in Refs. [34, 36] we study the GAP in different scenarios. We first show that, for Hamiltonians that satisfy a frustration-free property [13, 17, 30, 44], a quadratic spectral gap amplification is possible and *optimal* in some suitable black-box model. Quadratic spectral gap amplification was previously known only for Hamiltonians that result from a similarity transformation of stochastic matrices [34, 36, 38]. These Hamiltonians, which are the so-called discriminants, are also stoquastic, i.e. the off-diagonal entries are non-positive. A direct consequence of our new construction is that spectral gap amplification can now be achieved for any frustration-free Hamiltonian as those used in quantum adiabatic simulations of quantum circuits, improving upon the results in Refs. [1, 3, 28]. The eigenstate in our construction is not the lowest-energy state of the Hamiltonian induced by the quantum circuit (which was the case in Refs. [1, 3, 28]). Nevertheless, techniques for AST work for any eigenstate and still apply to our case, giving more efficient quantum adiabatic simulations to prepare the circuit's output state. In addition, some low-energy eigenstates of general frustration-free Hamiltonians (termed PEPS) were shown to play an important role in physics, renormalization, and optimization [30, 40, 44]: PEPS are useful variational states reproducing the local physics with high accuracy [39, 40]. Our method for spectral gap amplification can be used to speedup the preparation of PEPS on a quantum computer, with an expected improvement of other methods for this goal (e.g., [32]).

Interestingly, it is hard to achieve spectral gap amplification for general Hamiltonians that do not satisfy the frustration-free property. It would otherwise imply fast quantum methods for some problems, contradicting known complexity bounds in the oracle setting [7]. In addition, a natural generalization of our technique for spectral gap amplification would output a Hermitian operator only if the input Hamiltonian H is semidefinite positive (as is the case for frustration-free Hamiltonians). Such a generalization could then be used to solve the decision problem of whether a given H is semidefinite positive or not (with some small error bound), and provide an estimate to the lowest eigenvalue of H. This is a complete problem for the class QMA [26], and finding its solution may require a quantum circuit of undesirably large size.

The result on the impossibility of general spectral gap amplification has additional implications on the power of classical methods that solve the AST problem. For example, some stoquastic Hamiltonians can be mapped to stochastic matrices through a similarity transformation [12, 13, 33]. The fixed point of the stochastic matrix coincides with the distribution induced by the ground state of the Hamiltonian. The stochastic matrix would allow us to build a Monte-Carlo method to sample from the fixed point and classically solve the AST problem in this case. The same similarity transformation could then be used to give a frustration-free representation of the Hamiltonian [33, 40], allowing us to use the results for frustration-free Hamiltonians to amplify the gap. This contradicts the impossibility result, implying that finding the transformation is hard (in the black-box setting).

We organize the manuscript as follows. First, in Sec. II we define the GAP in more detail and present some methods for the proofs of the quadratic amplification in the frustration-free case, and for the impossibility of amplification in general. In Sec. III we show the construction of H' for frustration-free Hamiltonians. In Sec. III A we prove that the cost of implementing evolutions under H' is nearly the same as the cost of implementing evolutions under H in the black-box model, and present a simulation method. In Sec. III B we prove that the quadratic amplification is optimal for frustration-free Hamiltonians in this black-box model. In Sec. IV we comment on the advantages of our gap amplification construction for the adiabatic simulation of quantum circuits and give a local Hamiltonian. In Sec. V we prove that no gap amplification is possible for general Hamiltonians in the black-box model. This results in limits on the power of classical simulations for quantum systems, that are discussed in Sec. VI. We conclude in Sec. VII.

II. DEFINITIONS AND METHODS

We first comment on the implementation cost of Hamiltonian evolutions. Ignoring precision, we assume Hamiltonians H given as black boxes that implement the evolution $\exp\{-iHt\}$ on ancilliary input $|t\rangle$, satisfying $|t| \leq \pi$. When we write $H = \sum_k \Lambda_k$ for the Hamiltonian, we assume a modified black-box that implements $\exp\{-i\Lambda_k t\}$ on ancilliary input $|k\rangle|t\rangle$. Such a black box is typically an efficient quantum circuit in applications. We write O_H for the black-box. The implementation cost of $\exp\{-iH's\}$ in the black-box model, where H' is the Hamiltonian resulting from the gap amplification construction applied to H, will be the number of times O_H is called to implement the operation in either case.

Definition 1 *Let $H = \sum_k \Lambda_k \in \mathbb{C}^N \times \mathbb{C}^N$ be a finite-dimensional Hamiltonian, $|\psi\rangle$ a (unique) eigenstate of H with eigenvalue λ, and Δ the gap to the other nearest eigenvalue of λ; i.e. Δ is the spectral gap. The goal of the GAP is to find a new Hamiltonian H' that has $|\psi\rangle \otimes |\mathsf{o}\rangle$ as (unique) eigenstate and spectral gap $\Delta' \in \Omega(\Delta^{1-\epsilon})$, for $\epsilon > 0$. $|\mathsf{o}\rangle$ is any subsystem's state for which an efficient quantum circuit is known. The implementation cost of $\exp\{-iH't\}$ in the black-box model must be of the same order as the implementation cost of $\exp\{-iHt\}$.*

The last requirement forbids naive constructions such as $H' = cH$ for $c \gg 1$ and is needed to carry quantum speedups from the Hamiltonian-based model, in which the cost is the evolution time, to the black-box model, in which the cost is the number of black boxes. If both H and H' are sparse with a bounded number of (efficiently computable) non-zero entries per row and matrix oracles for the Λ_k are provided, the last requirement can be satisfied [2, 8, 15, 35].

Definition 2 *A Hamiltonian $H \in \mathbb{C}^N \times \mathbb{C}^N$ is frustration free if it can be written as $H = \sum_{k=1}^{L} a_k \Pi_k$, with (known) $0 \leq a_k \leq 1$, $(\Pi_k)^2 = \Pi_k$ projectors, and $L \in \mathcal{O}[polylog(N)]$. Further, if $|\psi\rangle$ is a lowest-eigenvalue eigenstate (ground state) of H, then $\Pi_k |\psi\rangle = 0 \ \forall \ k$.*

$|\psi\rangle$ is then a ground state of every term in the decomposition of H. We will assume that $a_k = 1 \ \forall k$ when there is no loss of generality. Then $H = \sum_k \Pi_k$ is still frustration free, it has $|\psi\rangle$ as ground state, it has at least the spectral gap of $\sum_k a_k \Pi_k$, and $\|H\| \leq L$. While the Π_k's may be local operators, we do not make that assumption here. It will suffice to have access to the black box O_H.

We consider the unstructured search problem or SEARCH to prove the optimality of quadratic gap amplification for frustration-free Hamiltonians and also the impossibility of gap amplification for general Hamiltonians. We introduce SEARCH and describe a simple variant of Grover's algorithm to solve it.

Definition 3 *Consider a family of Boolean functions or* oracles *$\{f_x\}$, with domain $[0, .., M-1]$, that satisfy $f_x(y) = \delta_{xy}$. The goal of SEARCH is to find the unknown word x, with high probability, by querying f_x the least possible number of times.*

For quantum algorithms, we use a reversible version of f_x defined as

$$Q_x |y\rangle |b\rangle = |y\rangle |b \oplus f_x(y)\rangle \ , \tag{1}$$

with $b \in \{0, 1\}$. We refer to Q_x as the search oracle. Quantum algorithms that solve SEARCH with $\Theta(\sqrt{M})$ uses of Q_x, and other x-independent operations, exist [7, 21]. Classical algorithms require $\Theta(M)$ queries to f_x.

From Ref. [16] we can give a quantum algorithm to solve SEARCH using projective measurements.

Lemma 1 *Let $\{H_x\}$ be a family of Hamiltonians with unique eigenstates $|\psi_x\rangle$ satisfying $p_x = |\langle x|\psi_x\rangle|^2 \in \mathcal{O}(1)$ and $p_s = |\langle s|\psi_x\rangle|^2 \in \mathcal{O}(1)$, where $|s\rangle = \frac{1}{\sqrt{M}} \sum_{y=0}^{M-1} |y\rangle$ is the equal superposition state. A quantum algorithm can find x with cost $\mathcal{O}(1/\Delta)$, where Δ is a bound on the eigenvalue gap of H_x.*

Proof 1 *The probabilities p_x and p_s are bounded independently of M. If we initialize a quantum computer in $|s\rangle$ and next we perform two projective measurements of $|\psi_x\rangle$ and $|x\rangle$, respectively, we can find x with probability no smaller than $p_x.p_s$. The cost of this method is measured by T, the time of evolving with H_x. This is needed to simulate the measurement of $|\psi_x\rangle$, with high accuracy, using the well-known phase estimation algorithm [16, 25] or by phase randomization [9]. The cost T is dominated by the inverse gap of H_x. In optimal constructions, $\Delta \in \mathcal{O}(1/\sqrt{M})$ so that $T \in \mathcal{O}(\sqrt{M})$.*

III. GAP: FRUSTRATION-FREE HAMILTONIANS

When a Hamiltonian $H = \sum_k \Pi_k$ is frustration free, local, and stoquastic [12], we could use the results in Ref. [13] and Refs. [34, 36] to build a Hamiltonian H' with a quadratically bigger spectral gap. In that case, H is the so-called discriminant of a Markov process or stochastic matrix. Such a construction uses tools from quantum walks [38] to speed up Markov chain based algorithms and requires full knowledge of each Π_k. Here, we present a novel and more efficient technique that allows for quadratic gap amplification of any frustration-free Hamiltonian, without the requirement of H being local and stoquastic. Later we show that this technique provides the optimal amplification and present a black-box simulation method for the evolution with H' that does not require full knowledge of each Π_k.

Theorem 1 *Let $H = \sum_{k=1}^{L} \Pi_k \in \mathbb{C}^N \times \mathbb{C}^N$ be a frustration-free Hamiltonian, $|\psi\rangle$ a ground state, and Δ the spectral gap or smallest nonzero eigenvalue. Then, there exists H' satisfying $H'|\psi\rangle|\mathfrak{o}\rangle = 0$ and eigenvalue gap $\Delta' \in \Omega(\sqrt{\Delta L})$. The dimension of the null space of H' is that of the null space of H.*

Proof 2 *The proof is constructive. We let*

$$X = \sum_{k=1}^{L} \Pi_k \otimes \Upsilon_k \tag{2}$$

be a Hamiltonian, and $\Upsilon_k = |k\rangle\langle k|$ a rank-one projector acting on $\mathbb{C}^L$. It is easy to show that $X^n = X$ for $n \geq 1$. We define the unitary operator

$$U = e^{-i\pi X} = \mathbb{1} - 2\sum_{k=1}^{L} \Pi_k \otimes \Upsilon_k , \tag{3}$$

with $\mathbb{1}$ the identity matrix whose dimension will be clear from context (here it is $NL \times NL$). The implementation of U requires a single black box O_H and $U^2 = \mathbb{1}$. In addition, $|\mathbf{o}\rangle = \frac{1}{\sqrt{L}}\sum_{k=1}^{L}|k\rangle$ is an ancilliary state and $P = \mathbb{1} \otimes |\mathbf{o}\rangle\langle\mathbf{o}|$ its projector. Because $L \in \mathcal{O}[polylog(N)]$, quantum circuits that prepare $|\mathbf{o}\rangle$ efficiently exist (e.g. [35]).

We define the Hamiltonian

$$G = UPU - P . \tag{4}$$

If $\{|\psi_j\rangle\}_{1 \leq j \leq N}$ are the eigenstates of H with eigenvalues $\lambda_1 = 0 \leq \lambda_2 \leq \ldots \leq \lambda_N \leq L$, G leaves invariant the at most two-dimensional subspace $\mathcal{V}_j$ spanned by $\{|\psi_j\rangle|\mathbf{o}\rangle, U|\psi_j\rangle|\mathbf{o}\rangle\}$, for all j. To show this, we need an important property that is a consequence of our construction:

$$PUP = \left(\mathbb{1} - \frac{2}{L}H\right) \otimes |\mathbf{o}\rangle\langle\mathbf{o}| = A \otimes |\mathbf{o}\rangle\langle\mathbf{o}| . \tag{5}$$

The eigenvalues of A are $\gamma_j = 1 - 2\lambda_j/L$ and the eigenstates are also the $|\psi_j\rangle$. To avoid technical issues we assume that $\gamma_N \geq 0$. If this were not the case, the projector P could be easily modified so that $\gamma_j = 1 - \lambda_j/L$ instead, and the assumption is fulfilled. We can think of γ_j as being $\cos\alpha_j = \langle\psi_j|\langle\mathbf{o}|PUP|\psi_j\rangle|\mathbf{o}\rangle$, so that α_j is the angular distance between the states $|\psi_j\rangle|\mathbf{o}\rangle$ and $U|\psi_j\rangle|\mathbf{o}\rangle$ (Fig. 1). When j is such that $\lambda_j = 0$, we have $U|\psi_j\rangle|\mathbf{o}\rangle = |\psi_j\rangle|\mathbf{o}\rangle$ and $\mathcal{V}_j$ is one dimensional. This implies that $|\psi_j\rangle|\mathbf{o}\rangle$ is an eigenstate of G of eigenvalue 0. When $\lambda_j \neq 0$, Eq. (5) gives $G|\psi_j\rangle|\mathbf{o}\rangle = \gamma_j U|\psi_j\rangle|\mathbf{o}\rangle - |\psi_j\rangle|\mathbf{o}\rangle$, and G is invariant in $\mathcal{V}_j$.


FIG. 1:   Geometric representation of the action of G on $|\psi_j\rangle|\mathbf{o}\rangle$.

When $\lambda_j \neq 0$ (or $\alpha_j \neq 0$), we define $|\perp_j\rangle$ as the unit state orthogonal to $|\psi_j\rangle|\mathbf{o}\rangle$ in the subspace $\mathcal{V}_j$. We can use simple geometry arguments (Fig. 1) to show that the two-dimensional block of G in the basis $\{|\psi_j\rangle|\mathbf{o}\rangle, |\perp_j\rangle\}$ is

$$G_j = \begin{pmatrix} -\sin^2\alpha_j & \sin\alpha_j\cos\alpha_j \\ \sin\alpha_j\cos\alpha_j & \sin^2\alpha_j \end{pmatrix} . \tag{6}$$

The eigenvalues of G in this subspace are $\pm\sin\alpha_j$. Because $\sin\alpha_j \geq \sqrt{2\lambda_j/L}$, a quadratic amplification in the spectral gap follows.

We need to account for a small technical issue to obtain H'. G acts trivially on any other eigenstate $|\phi\rangle$ orthogonal to $\oplus_j\mathcal{V}_j$. The dimension of the null space of G is much larger than the dimension of the null space of H. Usually,

this poses no difficulty in the AST problem because G does not generate transitions between the states $|\psi_j\rangle|\mathbf{o}\rangle$ and any other state in the null space of G. In those cases we can assume $H' = LG$ so that $\|H'\| \leq L$. Otherwise, since $P|\phi\rangle = (\sum_j |\psi_j\rangle\langle\psi_j| \otimes |\mathbf{o}\rangle\langle\mathbf{o}|)|\phi\rangle = 0$, we can easily remove the unwanted degeneracy of G by adding a "penalty term" to those states whose ancilliary register is not in $|\mathbf{o}\rangle$. That is,

$$H' = LG + \delta(\mathbb{1} - P) \tag{7}$$

for some $\delta > 0$ determined below. Note that, if $\lambda_j = 0$, $|\psi_j\rangle|\mathbf{o}\rangle$ is still in the null space of H'. For $\lambda_j \neq 0$, the matrix of H' in $\mathcal{V}_j$ is

$$H_j' = L \begin{pmatrix} -\sin^2\alpha_j & \sin\alpha_j\cos\alpha_j \\ \sin\alpha_j\cos\alpha_j & \delta/L + \sin^2\alpha_j \end{pmatrix} . \tag{8}$$

It suffices to choose $\delta = \sqrt{\Delta L}$ to guarantee that $\xi_j^{\pm}$, the eigenvalues of H' in this subspace, satisfy $|\xi_j^{\pm}| \geq (1/6)\sqrt{\lambda_j L}$. The gap of H' is then $\Delta' \in \mathcal{O}(\sqrt{\Delta L})$ and the dimension of the null space of H' coincides with that of the null space of H.

Our choice of δ uses Δ, which may be unknown. We can then safely use a lower bound on the spectral gap of H, $\bar{\Delta} \leq \Delta$, and build H' using $\bar{\Delta}$ instead. The loss is that the resulting gap of H' may not be the largest possible. However, this is not a drawback if we use H' (or LG) for the AST problem. As mentioned, the relevant spectral gap is of order $\sqrt{\Delta L}$ even if Δ is unknown, and the role of $\bar{\Delta}$ is unimportant in this case.

A. Black-box simulation

In a Hamiltonian-based model of quantum computation, a gain on the spectral gap readily provides a quantum speedup for the AST problem. This model assumes the ability of evolving under H', i.e. applying $V(t) = \exp\{-iH't\}$ directly. To carry those speedups to the quantum circuit model, we need a way of approximating a fixed-time evolution by a short sequence of gates. Usually, such a sequence is obtained by using Trotter-Suzuki-like approximations [8, 42]. In those approximations, the evolution of a Hamiltonian $H = \sum_k \Lambda_k$ is built upon short-time evolutions under each Λ_k, and each such evolution is written as a constant-size sequence of gates [35]. In other words, what determines the size of the circuit that approximates $V(t)$ is the implementation cost as given by the number of uses of O_H, the black box described in Sec. II. Here we describe a simulation method for $V(t)$ whose overall cost is almost linear in $L|t|$. In most cases this result implies a quantum speed-up, in the circuit model, for those methods that solve the AST problem using paths of frustration-free Hamiltonians.

Lemma 2 *For any κ and ϵ, there is a black-box simulation $W(t)$ of $V(t)$ that requires using the black box $\mathcal{O}\left(|Lt|^{1+1/(2\kappa)}\right)$ times, $\kappa \in \mathbb{N}^*$, and satisfies $\|W(t) - V(t)\| \leq \epsilon$.*

Proof 3 *Up to a trivial offset, $H' = A_1 - (L+\delta)P$, with $A_1 = LUPU$, $\|A_1\| = L$, and $\|P\| = 1$. We can then use known results on Hamiltonian simulations [2, 8, 15, 29] to approximate $\exp\{-iH't\}$ at precision ϵ by a sequence of $c(\epsilon,\kappa)|Lt|^{1+1/\kappa}$ concatenated evolutions under A_1 and P. For fixed ϵ and κ, $c(\epsilon,\kappa)$ is constant. The proof follows by noticing that*

$$e^{-iA_1 s} = U e^{-iLPs} U , \tag{9}$$

i.e. it requires two uses of U, and each U can be implemented using O_H once on input $|\pi\rangle$. Evolutions under P for any time do not require the black box and can be implemented with a quantum circuit of size L or smaller.

B. Optimal amplification

We are interested in finding the biggest spectral gap amplification possible for frustration-free H. To this end, we consider the set $\mathcal{F}$ of all those H' that have $|\psi\rangle|\mathbf{o}\rangle$ as eigenstate and such that evolving with H' for time t can be done using $\mathcal{O}(|t|)$ black boxes O_H; i.e. L is constant. Under these assumptions, the quadratic spectral gap amplification is optimal in the black-box model:

Theorem 2 $\Delta' \in \Theta(\sqrt{\Delta})$.

Proof 4 *We sketch the proof here and leave the full version in Appendix A. We consider instances* $\{\tilde{H}_x\}$, $x \in \{0, 1, \ldots M-1\}$ *unknown, that can be used to solve SEARCH as described in Lemma 1. Each* $\tilde{H}_x$ *is frustration free and the spectral gap is of order* $1/M$. *The corresponding black box* $O_{\tilde{H}_x}$ *can be implemented with a single call to the search oracle of Eq. (1). If* $\tilde{H}'_x$ *is the Hamiltonian with the amplified gap* Δ', *SEARCH can be solved by evolving with* $\tilde{H}'_x$ *for time* $\Omega(1/\Delta')$. *By assumption, this would require using* $\Omega(1/\Delta')$ *black boxes* O_{H_x} *and thus* $\Omega(1/\Delta')$ *search oracles. The lower bound on SEARCH implies* $\Delta' \in \Theta(\sqrt{\Delta})$, *with* $\sqrt{\Delta} \in \mathcal{O}(1/\sqrt{M})$.

IV. GAP AMPLIFICATION AND ADIABATIC QUANTUM COMPUTATION

A property of our spectral gap amplification method is that it yields quantum speed-ups for the AST problem. It is also important to describe additional advantages of the method, including the simulation of quantum circuits. It is known that any quantum circuit specified by quantum gates $U_1, \ldots, U_Q$ can be simulated, in a Hamiltonian-based model, by an adiabatic quantum evolution involving frustration-free Hamiltonians $H(s) = \sum_{k=1}^{L} a_k(s)\Pi_k(s)$. For $s = 1$, the ground state of $H(1)$ has large overlap with the state output by the quantum circuit. In some constructions, $L \in \mathcal{O}[\mathrm{poly}(Q)]$ and Π_k denotes nearest-neighbor (two-body) interactions between spins of corresponding many-body systems in one or two-dimensional lattices [1, 3].

The cost of the adiabatic simulation depends on the inverse minimum gap of $H(s)$, Δ. We can therefore use our gap amplification construction for frustration-free Hamiltonians to speed-up the adiabatic simulation. However, in doing so, we encounter difficulties regarding the locality of H'. For example, the projector P can be interpreted as an interaction term involving $\log_2(L)$ qubits and this is undesired if we want to find a physical H' whose interactions address a limited number of subsystems. A simple procedure to avoid many-body terms is to consider each state $|k\rangle$ in $|\mathsf{o}\rangle$ as a state of L qubits of the form

$$|k\rangle \rightarrow |0\ldots0 \underbrace{1}_{k\text{th position}} 0\ldots0\rangle . \tag{10}$$

The subspace defined by Eq. (10) is typically termed the *single-particle* subspace. The terms $|k\rangle\langle k'|$ appearing in P map to $\sigma_k^+ \sigma_{k'}^-$, with $\sigma^\pm = (\sigma^x \pm i\sigma^y)/2$, and σ^α the corresponding Pauli operator. The terms $|k\rangle\langle k|$ map to $(\mathbb{1} - \sigma_k^z)/2$. In this subspace,

$$G = \frac{1}{L}\left[\sum_{k,k'} e^{i\pi\Pi_k(\mathbb{1}-\sigma_k^z)/2}\sigma_k^+\sigma_{k'}^- e^{-i\pi\Pi_{k'}(\mathbb{1}-\sigma_{k'}^z)/2} - \sigma_k^+\sigma_{k'}^-\right] . \tag{11}$$

If the Π_k are two-subsystem interaction terms, $H' = LG + \delta(\mathbb{1} - P)$ will map to a combination of four-subsystem terms. This makes H' local but not spatially local. We also note that, while the overall spectrum of H' may have changed because of the mapping, the spectrum of H' in the single-particle subspace corresponds to the one analyzed in Sec. III.

We note that our construction to amplify the gap is not unique. Indeed, as seen in Appendix B, the Hamiltonian

$$\tilde{G} = \sum_{k=1}^{L} \Pi_k \otimes (|k\rangle\langle 0| + |0\rangle\langle k|) \tag{12}$$

has a similar property: its eigenvalues are $\pm\sqrt{\lambda_j}$. Note that if $a_k \neq 1$ we can modify it as $\tilde{G} = \sum_{k=1}^{L} \sqrt{a_k}\,\Pi_k \otimes (|k\rangle\langle 0| + |0\rangle\langle k|)$. The action of $\tilde{G}$ in the single-particle subspace is equivalent to that of [45]

$$\tilde{G} = \sum_{k=1}^{L} \Pi_k \otimes (\sigma_k^+\sigma_0^- + \sigma_k^-\sigma_0^+) . \tag{13}$$

To eliminate the degeneracy of the null eigenvalue in this subspace, we can proceed as in Sec. III and add a term proportional to $(\mathbb{1} - |0\rangle\langle 0|) \rightarrow (1 + \sigma_0^z)/2$, so that

$$\tilde{G} \rightarrow \bar{G} = \sum_{k=1}^{L} \Pi_k \otimes (\sigma_k^+\sigma_0^- + \sigma_k^-\sigma_0^+) + \frac{\sqrt{\Delta}}{2}\frac{(1+\sigma_0^z)}{2} . \tag{14}$$

The relevant spectral gap of $\bar{G}$ in the single-particle subspace is still of order $\sqrt{\Delta}$.

Nevertheless, an unwanted degeneracy of the null eigenvalue from other (many-particle) states is possible. (Note that any eigenstate in the 0-particle subspace has eigenvalue exactly $\sqrt{\Delta}/2$). To remove such a degeneracy we set "penalties" to states that belong to different particle subspaces. To this end, we note that

$$0 \le \tilde{G}^2 = \sum_{k,k'} \Pi_k \Pi_{k'} (\sigma_k^+ \sigma_{k'}^- \sigma_0^- \sigma_0^+ + H.c.)$$

$$\le S^+ S^- \sigma_0^- \sigma_0^+ + H.c. , \tag{15}$$

where $S^+ = \sum_{k=1}^L \sigma_k^+$ is proportional to the so-called $SU(2)$ spin *raising* operator. Let $|S, m\rangle$ be the eigenstates of $S^z = \sum_{k=1}^L \sigma_k^z$, with eigenvalue m, and total $SU(2)$ spin proportional to S. Then, $m \in \{-L, -L+2, \ldots, L\}$, $0 \le S \le L$, and

$$S^+ |S, m-2\rangle = N(S, m)|S, m\rangle ,$$

$$N(S, m) = \frac{1}{2}\sqrt{(S+m)(S-m+2)} . \tag{16}$$

It implies

$$\left(S^+ S^- \sigma_0^- \sigma_0^+ + H.c.\right)|b\rangle_0 \otimes |S, m\rangle = \tag{17}$$

$$= \left(\delta_{b,0} N(S, m+2)^2 + \delta_{b,1} N(S, m)^2\right)|b\rangle_0 \otimes |S, m\rangle ,$$

where $b \in \{0, 1\}$ denotes the state of the qubit in the 0-th position. The operator $\tilde{G}^2$ leaves the a-particle subspace invariant, and we write $a = (L-m)/2$. From Eqs. (16) and (17), the eigenvalues of $\tilde{G}^2$ in that subspace are bounded from above by $(L+m+2)(L-m+2)/4 = (L+1-a)(1+a)$. The upper bound to the absolute value of the eigenvalues of $\tilde{G}$, in that subspace, is $\sqrt{(L+1-a)(1+a)}$.

We define $Z = L - 2 - \sum_{k=0}^L \sigma_k^z$ and note that Z acts trivially in the single-particle space and nontrivially elsewhere. We also define

$$H' = \frac{1}{L} \bar{G} + 2Z , \tag{18}$$

$d > 0$. H' acts invariantly in the a-particle subspaces. The eigenvalues of H' are then bounded from below by

$$-\frac{1}{L^{1/d}}\sqrt{(L+1-a)(1+a)} + 4(a-1) . \tag{19}$$

For the specific case of $d = 2$, $a = 1$, we already know that the eigenvalues of H' are in the range $[-1, 1]$; this case corresponds to the single-particle one whose spectrum was analyzed previously. Also, for $d \le 2$, $a \ge 2$, all the eigenvalues coming from the a-particle subspaces do not *mix* with the eigenvalues of the single-particle subspace, and the potentially high degeneracy of the null space is avoided. For $a = 0$ the exact eigenvalues of H' are $\sqrt{\Delta}/(2L^{1/d}) - 4 \le -7/2$ if $d \le 2$.

Summarizing, the Hamiltonian

$$H' = \frac{1}{L^{1/d}}\left[\sum_{k=1}^L \Pi_k(\sigma_k^+ \sigma_0^- + \sigma_k^- \sigma_0^+) + \frac{\sqrt{\Delta}}{2}\frac{(1+\sigma_z^0)}{2}\right] + 2Z \tag{20}$$

has $|\psi_0\rangle|10\ldots0\rangle$ as unique eigenstate of eigenvalue 0, and all other eigenvalues are at distance at least $\sqrt{\Delta}/L^{1/d}$ if $d \le 2$. H' can be regarded as a physical Hamiltonian. In Fig. 2 we represent H' for the case in which H is the frustration-free Hamiltonian of a system in one spatial dimension, and $d = 1$.

The construction in Ref. [1] leads indeed a family of frustration-free Hamiltonians in one spatial dimension. For a quantum circuit with Q gates, the spectral gap of H in Ref. [1] can be made of order $1/Q^2$ [46]. H' in Eq. (20) allows us to amplify the gap up to order $1/Q^{1+1/d}$, with an important gain if $d = 2$ and $Q \gg 1$.

V. GAP: GENERAL HAMILTONIANS

We show that spectral gap amplification is not possible for general Hamiltonians in the black-box model. Again, we consider instances H_x whose unique ground states $|\psi_x\rangle$ can be used to solve SEARCH as in Lemma 1. As before, the cost of an evolution with H_x is the number of black boxes O_{H_x} required to implement it. We obtain:

FIG. 2: A physically local representation of H'. Each additional qubit $1, \ldots, L$ has assigned a color and interacts with a particular term Π_k in the Hamiltonian and the centered qubit. In a one dimensional construction, Π_k represents a nearest-neighbor interaction. Therefore, Π_k and Π_{k+1} have support overlapping on a single body that is represented by lines of color that match that of the additional qubit.

Theorem 3 *There are stoquastic (sparse) Hamiltonians for which spectral gap amplification is not possible.*

Proof 5 *The result will follow again from the known lower bound on SEARCH. We consider the family $\{H_x\}_{0 \le x \le M-1}$ with*

$$H_x = -|x\rangle\langle x| - cA , \tag{21}$$

and x unknown. The constant $c > 0$ will be specified later and A is a $M \times M$ symmetric matrix proportional to the adjacency matrix of a square lattice in five spatial dimensions. The matrix entries of A are $a_{y'y} = 1$ if (y, y') is an edge in the lattice and $a_{y'y} = 0$ otherwise. A is sparse and independent of the unknown x. In this case we assume a matrix oracle W_x for H_x. This is stronger than O_{H_x} in the sense that O_{H_x} can be implemented with a single call to W_x. Further, by construction, W_x requires calling the search oracle of Eq. (1) only once.

We present some useful properties of H_x and refer to Ref. [14] for details. There is a phase transition point at $c = c^* > 0$ where the spectral gap is

$$\Delta \to \frac{I_1}{2\sqrt{I_2 M}} \in \mathcal{O}(1/\sqrt{M}) , \tag{22}$$

and

$$p_x = |\langle x|\psi_x\rangle|^2 \to 0.37 , \tag{23}$$

$$p_s = |\langle s|\psi_x\rangle|^2 \to \frac{1}{2} . \tag{24}$$

Then, the ground state $|\psi_x\rangle$ has good overlap with both $|x\rangle$ and $|s\rangle$ as in Lemma 1. SEARCH can be solved with cost $\mathcal{O}(\sqrt{M})$ using the black boxes O_{H_x} or, similarly, the search oracles. If gap amplification for H_x were possible, this would yield a faster method to solve SEARCH. The proof follows by contradiction.

VI. STOQUASTIC VS. STOCHASTIC MATRICES

Because stoquastic Hamiltonians do not suffer from the so-called sign problem [37], an important task is to devise classical probabilistic methods to solve the AST problem by sampling from the distribution as determined by the ground state. A well-known method for this purpose is quantum Monte-Carlo (QMC) [6]. Another method may be obtained by finding a diagonal similarity transformation that maps stoquastic Hamiltonians to stochastic matrices [13]. The fixed point of the resulting Markov chain is the desired distribution. Contrary to QMC, this second method requires knowledge of the ratio of amplitudes of the ground state for the mapping. Our result on no gap amplification for all stoquastic Hamiltonians implies that finding the transformation to stochastic matrices should be hard in the black-box model. Otherwise it would imply a method to amplify the gap by following the results on frustration-free Hamiltonians.

We can go further and give a family of stoquastic Hamiltonians for which classical probabilistic methods, that solve the AST by sampling from the ground state, has a large black-box implementation cost. For some QMC algorithms, our result implies that the spectral gap of the constructed stochastic matrix rapidly decreases with the problem size. Once more, our proof considers a family of Hamiltonians $\{H_x\}_{0 \leq x \leq M-1}$ such that sampling from the ground state of H_x outputs the unknown x with large probability.

Lemma 3 *Consider the family of Hamiltonians*

$$H_x = -|x\rangle\langle x| + F/4 \ , \tag{25}$$

where $F \in \mathbb{C}^M \times \mathbb{C}^M$ is independent of x and $\|F\| \leq 1$. Then, the lowest eigenvalue is bounded as $E_0 \leq -3/4$, the spectral gap satisfies $\Delta \geq 1/2$, and the ground state has large support on $|x\rangle$, i.e. $|\langle x|\psi_x\rangle|^2 \in \mathcal{O}(1)$.

Proof 6 *(Sketched) All results can be simply proved by considering $-|x\rangle\langle x| + gF/4$ and using perturbation theory in the range $[0,1]$ for g.*

If F is stoquastic and irreducible, the ground state of H_x is $|\psi_x\rangle = \sum_y \sqrt{\pi_y}|y\rangle$, with $\pi_y > 0$. In this case we let S_x be a stochastic matrix whose fixed point is almost the distribution π. We assume that a matrix oracle for S_x follows from a matrix oracle for H_x. The latter requires using the (classical) search oracle only once: for input y, the matrix elements of H_x in the yth row depend on whether $x = y$ or not. In path integral QMC, for example, S_x can be constructed efficiently as follows. For sufficiently large β the probability π_y is (almost) proportional to

$$\langle y|e^{-\beta H_x}|y\rangle \ . \tag{26}$$

Because $\Delta \geq 1/2$, it suffices to choose $\beta \in \mathcal{O}(\log M)$ to bound the approximation error. Next, because the exponentiation of H_x might be hard to compute, we write $e^{-\beta H_x} = \prod_{l=1}^{p} e^{-\eta H_x}$, with $\eta \in \mathcal{O}(1/\beta)$, and approximate $e^{-\eta H_x} = \mathbb{1} - \eta H_x + \mathcal{O}(\eta^2)$. We can relate $\prod_{l=1}^{p}(\mathbb{1} - \eta H_x)$ with a "classical system" H_c (diagonal Hamiltonian) that acts on a larger Hilbert space of dimension M^L. More precisely, we define

$$H_c = -\sum \log(f_{r-1,r})|z^{(r-1)}\rangle|z^{(r)}\rangle\langle z^{(r-1)}|\langle z^{(r)}| \ , \tag{27}$$

with $f_{r-1,r} = \langle z^{(r-1)}|\mathbb{1} - \eta H_x|z^{(r)}\rangle$. Note that $\mathrm{tr}\,(e^{-\beta H_x}) \approx \mathrm{tr}\,(e^{-H_c})$. The computation of any diagonal entry of H_c requires evaluating $p \in \mathcal{O}[\mathrm{polylog}(M)]$ matrix elements of H_x and thus the same number of search oracles. Sampling from the distribution determined by H_c requires choosing an appropriate Markov chain, like the one in the Metropolis algorithm. From known complexity bounds on the search problem, any choice of Markov chain will lead, in the black-box model, to a mixing time

$$\tau_{\mathrm{mix}} \in \mathcal{O}\left[\frac{M}{\mathrm{polylog}(M)}\right] \ . \tag{28}$$

VII.   CONCLUSIONS

We introduced and studied the spectral gap amplification problem or GAP. The goal is to efficiently construct a Hamiltonian that has the ground state of a given Hamiltonian as eigenstate, but a bigger spectral gap. This problem is motivated by the results in Ref. [36], where we showed quantum speedups of Monte-Carlo methods by giving a quadratic gap amplification for a specific family of Hamiltonians. The GAP has important applications in adiabatic quantum computing and other methods for adiabatic state transformations, where the implementation cost is dominated by the inverse gap: If gap amplification is possible, this results in a quantum speedup.

We showed that a quadratic spectral gap amplification is indeed possible for frustration-free Hamiltonians. This generalizes the results in Ref. [36] and gives a more efficient construction to speedup classical Markov-chain based methods. In a suitable black-box model, we proved that our method provides the biggest amplification possible. The result for frustration-free Hamiltonians also provides more efficient quantum adiabatic methods to simulate quantum circuits and for the preparation of projected entangled pair states or PEPS.

We presented a family of Hamiltonians for which spectral gap amplification is not possible under some assumptions. The Hamiltonians in this case can be sparse and stoquastic. They are constructed so that by preparing their ground states one can solve the unstructured search problem, and then use known lower complexity bounds to prove the impossibility result by contradiction. A corollary is that classical probabilistic methods that sample from the ground state distribution of some Hamiltonians may require a large implementation cost.

Appendix A: Proof of optimal amplification

Each H_x is a symmetric matrix related to the adjacency matrix A of a sparse expander graph $G = (M, d, \lambda)$. M is the number of vertexes, $d \in \mathcal{O}(1)$ is the degree of G, and λd is the second largest eigenvalue of A. WLOG, $\lambda \leq 1/2$ (e.g., a Ramanujan graph). Before specifying H_x in a frustration-free form, we define H_x by its matrix elements as follows:

$$\langle y|H_x|z\rangle = \begin{cases} \frac{1}{M-1} & \text{if } y = z = x \\ \frac{-1}{d\sqrt{M-1}} & \text{if } \{y, z\} \in E \text{ and either } y = x \text{ or } z = x \\ \frac{-1}{d} & \text{if } \{y, z\} \in E \text{ and } x \neq y \neq z \neq x \\ 1 & \text{if } y = z \neq x \\ 0 & \text{otherwise.} \end{cases} \tag{A1}$$

E is the set of edges of the graph G. H_x has the following unique ground state of eigenvalue 0:

$$|\psi_x\rangle = \frac{1}{\sqrt{2}}|x\rangle + \frac{1}{\sqrt{2(M-1)}}\sum_{y \neq x}|y\rangle \,. \tag{A2}$$

This state satisfies the requirements of the algorithm in Lemma 1: $|\psi_x\rangle$ has constant overlap with $|x\rangle$ and the equal superposition state. Since x is unknown, the preparation of $|\psi_x\rangle$ will allow us to solve SEARCH.

We can bound the spectral gap Δ of H_x in terms of M. We define $Z_x = \mathbb{1} - H_x - |\psi_x\rangle\langle\psi_x|$. Because $\|H_x\| \leq 1$, the norm of Z_x is $\|Z_x\| = \sup_{|\phi\rangle}\langle\phi|Z_x|\phi\rangle = 1 - \Delta$. We can write $|\phi\rangle = \beta|x\rangle + \sqrt{1 - \beta^2}|x^\perp\rangle$, with $\beta > 0$ and $|x^\perp\rangle$ a unit state orthogonal to $|x\rangle$, to obtain

$$|\langle\phi|Z_x|\phi\rangle| \leq 2\beta\sqrt{1 - \beta^2}|\langle x|Z_x|x^\perp\rangle| + $$
$$+ \beta^2|\langle x|Z_x|x\rangle| + (1 - \beta^2)|\langle x^\perp|Z_x|x^\perp\rangle| \,. \tag{A3}$$

An upper bound on the *rhs* of Eq. (A3) implies a lower bound on the gap. We write $|x^\perp\rangle = \sum_{y \neq x} l_y|y\rangle$. The first term on the *rhs* of Eq. (A3) is

$$2\beta\sqrt{1 - \beta^2}\sum_{y \neq x}|l_y|\left|\left[\langle x|H_x|y\rangle + \frac{1}{2\sqrt{M-1}}\right]\right| \,. \tag{A4}$$

We use

$$\left|\langle x|H_x|y\rangle + \frac{1}{2\sqrt{M-1}}\right| \leq \frac{1}{2\sqrt{M-1}} \; \forall \; y \tag{A5}$$

and

$$\sum_{y \neq x}|l_y| = \||x^\perp\rangle\|_1 \leq \sqrt{M-1}\||x^\perp\rangle\|_2 = \sqrt{M-1} \tag{A6}$$

to show

$$|2\beta\sqrt{1 - \beta^2}\langle x|Z_x|x^\perp\rangle| \leq |2\beta\sqrt{1 - \beta^2}|\frac{1}{2} \,. \tag{A7}$$

Further, from the diagonal matrix elements of Z_x, we have ($M \geq 3$):

$$\beta^2 |\langle x|Z_x|x\rangle| \leq \beta^2 \left[\frac{1}{2} - \frac{1}{M-1}\right] . \tag{A8}$$

Finally, if $P_{\bar{x}} = \sum_{y \neq x} |y\rangle\langle y|$ is the projector onto the orthogonal subspace of $|x\rangle$, then the third term in the *rhs* of Eq. (A3) is bounded by $(1 - \beta^2)\|P_{\bar{x}} Z_x P_{\bar{x}}\|$. In addition,

$$P_{\bar{x}} H_x P_{\bar{x}} = P_{\bar{x}}(\mathbb{1} - A/d)P_{\bar{x}} , \tag{A9}$$

where A is the adjacency matrix of $G(M, d, \lambda)$. If $|s\rangle$ is the equal superposition state, $A|s\rangle = d|s\rangle$, and this is the only eigenstate of A with largest eigenvalue d. We can write

$$P_{\bar{x}}|\psi_x\rangle = \frac{\sqrt{M}}{\sqrt{2(M-1)}} P_{\bar{x}}|s\rangle , \tag{A10}$$

and consequently,

$$\|P_{\bar{x}} Z_x P_{\bar{x}}\| = \left\|P_{\bar{x}}\left(A/d - \frac{M}{2(M-1)}|s\rangle\langle s|\right) P_{\bar{x}}\right\|$$
$$\leq \left\|\left(A/d - \frac{M}{2(M-1)}|s\rangle\langle s|\right)\right\| .$$

Because the second largest eigenvalue of A/d is $\lambda \leq 1/2$, we obtain

$$\|P_{\bar{x}} Z_x P_{\bar{x}}\| \leq 1 - \frac{M}{2(M-1)} \leq \frac{1}{2} . \tag{A11}$$

Using Eqs. (A7), (A8), and (A11), we get

$$\langle \phi|Z_x|\phi\rangle \leq \frac{1}{2} + \beta\sqrt{1 - \beta^2} - \frac{\beta^2}{M-1} . \tag{A12}$$

The maximum is found when β satisfies

$$-\frac{\beta^2}{M-1} = \frac{\beta\sqrt{1 - \beta^2}}{(M-1)^2} - \frac{1}{2(M-1)} . \tag{A13}$$

Inserting Eq. (A13) in Eq. (A12), and since $\beta\sqrt{1 - \beta^2} \leq 1/2$, yields

$$\langle \phi|Z_x|\phi\rangle \leq 1 - \frac{1}{2(M-1)} + \frac{1}{2(M-1)^2} \tag{A14}$$

$$\leq 1 - \frac{1}{4(M-1)} . \tag{A15}$$

for $M \geq 3$. This implies

$$\Delta \geq \frac{1}{4(M-1)} . \tag{A16}$$

We now focus on the specification of H_x (or a modified version $\tilde{H}_x$) as a frustration-free Hamiltonian. To this end, we first describe a particular choice for the projectors appearing in H_x. We define the function

$$c(y) = \begin{cases} \frac{1}{\sqrt{d(M-1)}} & \text{if } y = x \\ \frac{1}{\sqrt{d}} & \text{e.o.c.} \end{cases} \tag{A17}$$

We also define E_o as the set of ordered pairs (y, z) of edges in G, where each edge appears only once. That is, if $\{y, z\} \in E$, then either (y, z) or (z, y) appears in E_o. For each edge $(y, z) \in E_0$ we define the (unnormalized) state

$$|\phi_{(y,z)}\rangle = c(y)|y\rangle - c(z)|z\rangle . \tag{A18}$$

Then,

$$H_x = \sum_{(y,z) \in E_o} |\phi_{(y,z)}\rangle\langle\phi_{(y,z)}| \, , \tag{A19}$$

as the matrix entries coincide with those in Eq. (A1). Because $\langle\phi_{(y,z)}|\psi_x\rangle = 0$ for all edges, H_x is frustration free.

For consistency with the definitions in previous sections, we define the projectors

$$\Pi_{(y,z)} = \frac{|\phi_{(y,z)}\rangle\langle\phi_{(y,z)}|}{\langle\phi_{(y,z)}|\phi_{(y,z)}\rangle} \tag{A20}$$

and the corresponding modified Hamiltonian

$$\tilde{H}_x = \sum_{(y,z) \in E_o} \Pi_{(y,z)} \, . \tag{A21}$$

$|\psi_x\rangle$ is still the unique ground state of $\tilde{H}_x$. Since

$$\langle\varphi|\tilde{H}_x|\varphi\rangle = \sum_{(y,z) \in E_o} \langle\varphi|\Pi_{(y,z)}|\varphi\rangle > \sum_{(y,z) \in E_o} \langle\varphi|\phi_{(y,z)}\rangle\langle\phi_{(y,z)}|\varphi\rangle = \langle\varphi|H_x|\varphi\rangle \, , \tag{A22}$$

the spectral gap of $\tilde{H}_x$ is also bounded from below as $\Delta \geq \frac{1}{4(M-1)}$.

The number of terms in Eq. (A21) grows with the number of edges in the graph, $dM/2$. It is convenient to reduce the number of projectors so that $L \in \mathcal{O}(1)$, and L playing no significant role in the scaling of the amplified gap. We then assume a given edge coloring of the graph G with edge chromatic number χ of order d [41] [i.e., $\chi \in \mathcal{O}(1)$]. Let $c_1, \ldots, c_\chi$ be the different colors. All the projectors $\Pi_{(y,z)}$ belonging to one of the colors are, by construction, orthogonal to each other as they don't share a vertex. For each $k \in \{1, \ldots, \chi\}$, we define the projectors

$$\Pi_k = \sum_{(y,z) \in c_k} \Pi_{(y,z)} \, , \tag{A23}$$

Then, $\tilde{H}_x = \sum_{k=1}^{\chi} \Pi_k$.

The specification of $\tilde{H}_x$ in this proof is as follows. For each k we assume the existence of a reversible version of the matrix oracle W that, on input (k, y), it outputs $(z, \langle z|\Pi_k|y\rangle)$, where z is the only vertex that shares an edge of coloring c_k with y. Note that all the matrix elements $\langle z|\Pi_k|y\rangle$ are equal unless $y = x$ or $z = x$. The implementation of W requires then deciding whether its input y is the marked vertex x or not. It implies that the reversible version of W can be implemented using a single call to the search oracle of Eq. (1). Further, the black box $O_{\tilde{H}_x}$, which implements $\exp\{-i\Pi_k t\}$ on ancilliary input $|k\rangle|t\rangle$, can be constructed using $\mathcal{O}(1)$ oracles W. This follows from the simple observation that $\exp\{-i\Pi_k t\}$ induces rotations in the two-dimensional subspaces spanned by $\{|y\rangle, |z\rangle\}$, where z is the only vertex that shares an edge of c_k with y. The implementation of a reversible $O_{\tilde{H}_x}$ requires then using the search oracle at most two times ($y \overset{?}{=} x$ or $z \overset{?}{=} x$).

Let $\tilde{H}_x'$ be the Hamiltonian that has $|\psi_x\rangle|\mathbf{o}\rangle$ as a non-degenerate eigenstate and spectral gap Δ'. SEARCH can then be solved using the technique in Sec. II which requires evolving with $\tilde{H}_x'$ for time $\mathcal{O}(1/\Delta')$. By assumption, it implies that SEARCH can be solved using $\mathcal{O}(1/\Delta')$ black boxes $O_{\tilde{H}_x}$ or, similarly, $\mathcal{O}(1/\Delta')$ search oracles. Therefore, the known complexity bound on SEARCH implies that $\Delta' \leq 1/\sqrt{M} \in \mathcal{O}(\sqrt{\Delta})$. In this black-box model, our construction in Sec. III gives the biggest gap amplification possible for frustration-free Hamiltonians.

Appendix B: Spectral properties of $\tilde{G}$

The first property to notice is that if $|\psi_0\rangle$ is in the null space of H, then $\Pi_k|\psi_0\rangle = 0$. Therefore $\tilde{G}|\psi_0\rangle|0\rangle = 0$ and $|\psi_0\rangle|0\rangle$ is an eigenstate in the null space of $\tilde{G}$.

As in Sec. III, we label the eigenvalues of H by λ_j. Assume now $\lambda_j \neq 0$, and consider the action of $\tilde{G}$ on the state $|\psi_j 0\rangle$

$$\tilde{G}|\psi_j 0\rangle = \sum_k \Pi_k |\psi_j k\rangle \, . \tag{B1}$$

Notice that

$$\langle\psi_j 0|\tilde{G}|\psi_j 0\rangle = \sum_k \langle\psi_j k|\Pi_k|\psi_j 0\rangle = 0 \tag{B2}$$

and that

$$\|\tilde{G}|\psi_j 0\rangle\|^2 = \sum_k \langle\psi_j k|\Pi_k|\psi_j k\rangle = \lambda_j \ . \tag{B3}$$

We denote by

$$|\perp_j\rangle = \frac{1}{\sqrt{\lambda_j}}\tilde{G}|\psi_j 0\rangle \tag{B4}$$

the corresponding normalized state. Next note that

$$\tilde{G}|\perp_j\rangle = \frac{1}{\sqrt{\lambda_j}}\sum_k \Pi_k|\psi_j 0\rangle = \sqrt{\lambda_j}|\psi_j 0\rangle \ . \tag{B5}$$

Therefore, the Hamiltonian $\tilde{G}$ is invariant in the subspace $\mathcal{V}_j = \{|\psi_j 0\rangle, |\perp_j\rangle\}$. Define $\tilde{G}_j = \tilde{G}_{|\mathcal{V}_j}$ the projection of $\tilde{G}$ in this invariant subspace. In the basis $\{|\psi_j 0\rangle, |\perp_j\rangle\}$ we can write

$$\tilde{G}_j = \begin{pmatrix} 0 & \sqrt{\lambda_j} \\ \sqrt{\lambda_j} & 0 \end{pmatrix} \ , \tag{B6}$$

with eigenvalues $\pm\sqrt{\lambda_j}$.

Finally, we note that because the $|\psi_j\rangle$ form a complete basis, $\tilde{G}$ acts trivially on any other state orthogonal to $\oplus_j \tilde{\mathcal{V}}_j$. While the dimension of the null space of $\tilde{G}$ is larger than that of the null space of H, this has no effects when solving the AST problem: transitions between $|\psi_j\rangle|0\rangle$ and any other orthogonal state in the null space are forbidden. We can also use the trick in Sec. III to reduce such a degeneracy, if necessary.

---

[1] D. AHARONOV, D. GOTTESMAN, S. IRANI, AND J. KEMPE, *The power of quantum systems on a line*, Comm. Math. Phys., 287 (2009), p. 41.
[2] D. AHARONOV AND A. TA-SHMA, *Adiabatic quantum state generation and statistical zero knowledge*, Proceedings of the 35th ACM Symposium on Theory of Computing, (2003), pp. 20–29.
[3] D. AHARONOV, W. VAN DAM, J. KEMPE, Z. LANDAU, S. LLOYD, AND O. REGEV, *Adiabatic quantum computation is equivalent to standard quantum computation*, SIAM J. Comp., 37 (2007), pp. 166–194.
[4] P. AMARA, D. HSU, AND J. E. STRAUB, *Global energy minimum searches using an approximate solution of the imaginary time schroedinger equation*, The J. of Phys. Chem., 97 (1993), pp. 6715–6721.
[5] A. AMBAINIS, *Quantum walk algorithm for element distinctness*, Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science, (2004), pp. 22–31.
[6] J. B. ANDERSON, *Quantum Monte Carlo*, Oxford University Press, Oxford, UK, 2007.
[7] C. BENNETT, E. BERNSTEIN, G. BRASSARD, AND U. VAZIRANI, *Strengths and weaknesses of quantum computing*, SIAM J. Comput., 26 (1997), pp. 1510–1523.
[8] D. BERRY, G. AHOKAS, R. CLEVE, AND B. SANDERS, *Efficient quantum algorithms for simulating sparse hamiltonians*, Comm. Math. Phys., 270 (2007), p. 359.
[9] S. BOIXO, E. KNILL, AND R. D. SOMMA, *Eigenpath traversal by phase randomization*, Quantum Inf. and Comp., 9 (2009), pp. 833–855.
[10] S. BOIXO, E. KNILL, AND R. D. SOMMA, *Fast quantum algorithms for traversing paths of eigenstates*, arXiv:1005.3034, (2010).
[11] S. BOIXO AND R. D. SOMMA, *Necessary condition for the quantum adiabatic approximation*, Phys. Rev. A, 81 (2010), p. 032308.
[12] S. BRAVYI, D. P. DIVICENZO, R. OLIVEIRA, AND B. M. TERHAL, *The complexity of stoquastic local hamiltonian problems*, Quant. Inf. and Comp., 8 (2008), p. 0361.
[13] S. BRAVYI AND B. TERHAL, *Complexity of stoquastic frustration-free hamiltonians*, SIAM J. Comput., 39 (2009), p. 1462.
[14] A. CHILDS AND J. GOLDSTONE, *Spatial search by quantum walk*, Phys. Rev. A, 70 (2004), p. 022314.
[15] A. CHILDS AND R. KOTHARI, *Simulating sparse hamiltonians with star decompositions.* arXiv:1003.3683, 2010.

[16] A. M. Childs, E. Deotto, E. Farhi, J. Goldstone, S. Gutmann, and A. J. Landahl, *Quantum search by measurement*, Phys. Rev. A, 66 (2002), p. 032314.

[17] N. de Beaudrap, M. Ohliger, T. Osborne, and J. Eisert, *Solving frustration-free spin systems*, Phys. Rev. Lett., 105 (2010), p. 060504.

[18] E. Farhi, J. Goldstone, and S. Gutmann, *Quantum adiabatic evolution algorithms versus simulated annealing*, quant-ph/0201031, (2002).

[19] 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 (2001), pp. 472–476.

[20] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, *Quantum computation by adiabatic evolution*, quant-ph/0001106, (2000).

[21] E. Farhi and S. Gutmann, *Analog analogue of a digital quantum computation*, Physical Review A (Atomic, Molecular, and Optical Physics), 57 (1998), pp. 2403–2406.

[22] 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 (1994), pp. 343–348.

[23] S. Jansen, M. Ruskai, and R. Seiler, *Bounds for the adiabatic approximation with applications to quantum computation*, J. of Math. Phys., 48 (2007), pp. 102111–15.

[24] T. Kadowaki and H. Nishimori, *Quantum annealing in the transverse ising model*, Phys. Rev. E, 58 (1998), pp. 5355–5363.

[25] A. Y. Kitaev, *Quantum measurements and the abelian stabilizer problem*, arxiv:quant-ph/9511026, (1995).

[26] A. Y. Kitaev, A. Shen, and M. Vyalyi, *Classical and Quantum Computation*, American Mathematical Society, 2002.

[27] A. Messiah, *Quantum Mechanics*, Dover Publications, 1999.

[28] A. Mizel, D. A. Lidar, and M. Mitchell, *Simple proof of the equivalence between adiabatic quantum computation and the circuit model*, Phys. Rev. Lett., 99 (2007), p. 070502.

[29] A. Papageorgiou and C. Zhang, *On the efficiency of quantum algorithms for hamiltonian simulation*, arXiv:1005.1318, (2010).

[30] D. Perez-Garcia, F. Verstraete, J. I. Cirac, and M. Wolf, *Peps as unique ground states of local hamiltonians*, Quant. Inf. Comp., 8 (2008), p. 0650.

[31] S. Sachdev, *Quantum Phase Transitions*, Cambridge University Press, UK, 2001.

[32] M. Schwarz, K. Temme, and F. Verstraete, *Contracting tensor networks and preparing peps on a quantum computer*, arXiv:1104.1410, (2011).

[33] R. Somma, C. Batista, and G. Ortiz, *Quantum approach to classical statistical mechanics*, Phys. Rev. Lett., 99 (2007), p. 030603.

[34] R. Somma and G. Ortiz, *Quantum approach to classical thermodynamics and optimization*, Lect. Notes Phys., 802 (2010), p. 1.

[35] R. Somma, G. Ortiz, J. Gubernatis, E. Knill, and R. Laflamme, *Simulating physical phenomena by quantum networks*, Phys. Rev. A, 65 (2002), p. 042323.

[36] R. D. Somma, S. Boixo, H. Barnum, and E. Knill, *Quantum simulations of classical annealing processes*, Phys. Rev. Lett., 101 (2008), pp. 130504–4.

[37] M. Suzuki, *Quantum Monte Carlo Methods in Condensed Matter Physics*, World Scientific, Singapore, 1998.

[38] M. Szegedy, *Quantum speed-up of markov chain based algorithms*, Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science, (2004), pp. 32–41.

[39] F. Verstraete and J. Cirac, *Renormalization algorithms for quantum-many body systems in two and higher dimensions*, Phys. Rev. A, 70 (2004), p. 060302(R).

[40] F. Verstraete, M. Wolf, D. Perez-Garcia, and J. Cirac, *Criticality, the area law, and the computational power of projected entangled pair states*, Phys. Rev. Lett., 96 (2006), p. 220601.

[41] V. Vizing, *On an estimate of the chromatic class of a p-graph*, Diskret. Analiz., 3 (1964), p. 25.

[42] N. Wiebe, D. Berry, P. Hoyer, and B. C. Sanders, *Higher-order decompositions of ordered operator exponentials*, J. Phys. A: Math. Theor., 43 (2010), p. 065203.

[43] P. Wocjan and A. Abeyesinghe, *Speed-up via quantum sampling*, Phys. Rev. A, (2008), p. 042336.

[44] M. M. Wolf, G. Ortiz, F. Verstraete, and J. I. Cirac, *Quantum phase transitions in matrix product systems*, Phys. Rev. Lett., 97 (2006), p. 110403.

[45] The corresponding states are still those in Eq. (10), but each having $L + 1$ qubits labeled from 0 (left) to L (right).

[46] We acknowledge discussions with S. Irani.