Dequantization Barriers for Guided Stoquastic Hamiltonians

Yassine Hamoudi* Yvan Le Borgne† Shrinidhi Teganahally Sridhara‡

Université de Bordeaux, CNRS, LaBRI, France

Last update: November 5, 2026

Abstract

We construct a probability distribution, induced by the Perron–Frobenius eigenvector of an exponentially large graph, that cannot be efficiently sampled by any classical algorithm, even when given the *best-possible* warm-start distribution. In the quantum setting, this problem can be viewed as preparing the ground state of a stoquastic Hamiltonian given a guiding state as input, and is known to be efficiently solvable on a quantum computer. We further show that even estimating the ground-state energy of that Hamiltonian remains classically hard. Our results suggest that no efficient classical algorithm can solve a broad class of stoquastic ground-state problems.

Our graph is constructed from a class of *high-degree, high-girth spectral expanders* (such as the LPS Ramanujan graphs) to which self-similar trees are attached. This builds on and extends prior work of Gilyén, Hastings, and Vazirani [Has21; GHV21], which ruled out dequantization for a specific stoquastic adiabatic path algorithm. We strengthen their result by ruling out any classical algorithm for guided ground-state preparation or energy estimation, while also providing a derandomized construction.

1 Introduction

Understanding and analyzing the ground state $|\psi_H\rangle$ of a many-body system represented by a Hamiltonian H acting on m qubits is of fundamental interest in quantum computation with critical applications in quantum chemistry (molecular modeling), materials science (high-temperature superconductivity), and quantum machine learning, where ground states are used to represent optimal model parameters. One of the main obstructions to efficient preparation of the ground state or estimation of the ground-state energy is the exponential size of the matrix H. In fact, in the absence of other information apart from the Hamiltonian, the problem of estimating the ground-state energy is computationally hard even for quantum computers; the problem is famously QMA-complete for local Hamiltonians [KSV02]. In practice, we often assume access to *guiding states* $|\psi_{\text{in}}\rangle$, which are states close to the actual ground state (measured by the overlap $\langle\psi_{\text{in}}\,|\,\psi_H\rangle$), as a computational resource for the ground-state preparation/energy estimation problems. We refer to this setting as the *Guided Ground-State Preparation* (GGSP) problem and to the related energy estimation problem as *Guided Ground-State Energy Estimation* (GGSE).

Guiding states can be seen as analogous to *warm starts* in Markov Chain Monte Carlo algorithms [AF02], ansätze in variational quantum algorithms [TCC+22], or the intermediate states along a path of adiabatic quantum computation [AL18]. Guiding states can also be classified based on their closeness to the actual ground state and the existence of efficient algorithms to prepare them or query their amplitudes. The well-studied types include guiding states with

*ys.hamoudi@gmail.com

†borgne@labri.fr

‡shrinidhi.teganahally-sridhara@u-bordeaux.fr

constant or $1/\operatorname{poly}(m)$ *global* overlap, *strong* guiding states with constant or $1/\operatorname{poly}(m)$ *pointwise* overlap [Bra15], *succinct* guiding states with efficient classical circuits for amplitude evaluation [Liu21; GG22; Jia25; Gal25], etc. In the presence of these guiding states, the GGSP/GGSE problems have been proved to be efficiently solvable by quantum algorithms and, in some cases, even by classical algorithms under additional assumptions on H and $|\psi_{\text{in}}\rangle$.

Dequantization. Recently, there have been several results on *dequantized* algorithms [GG22; Gal25; CFG+23] for the GGSE problem. These *classical* algorithms usually assume sampling access to the guiding state, along with amplitude access to the guiding state via an evaluation oracle. For instance, Le Gall [Gal25] showed that when $|\langle\psi_{\text{in}}\,|\,\psi_H\rangle| = \Omega(1)$ and the required precision is $1/\log(m)$, there exists a classical algorithm for energy estimation that runs in $\operatorname{poly}(m)$ time. In fact, when the exact ground state is given as input, $|\psi_{\text{in}}\rangle = |\psi_H\rangle$, a straightforward algorithm computes the energy *exactly*, provided amplitude access is available (see, e.g., the comment in [GG22]). Our result shows that sampling access alone is insufficient, even when the guiding state is very close to the exact ground state, thereby demonstrating the necessity of amplitude access for efficient classical energy estimation in this setting.

Another line of classical algorithms for the GGSP problems is the Quantum Monte Carlo (QMC) method, a widely used heuristic in computational condensed matter physics. Bravyi and Terhal [BT10] showed that when H is stoquastic and frustration-free, then the ground-state preparation is classically tractable. Later, in [BCGL23], it was shown that there exists a rapidly mixing Markov Chain that can sample from the ground state of any gapped, local H when given access to ratios of amplitudes of the ground state via an oracle, along with a starting state satisfying a mild technical condition. However, implementing the amplitude-ratio oracle is unknown for general Hamiltonians, including stoquastic Hamiltonians that are not frustration-free.

General barriers. Dequantization results have led to the question of *to what extent classical algorithms can suffice to solve the GGSP/GGSE problems*. This question has received most of its attention when the guiding state is an *untrusted* witness, i.e., claimed by the prover in prover-verifier protocols. While the general energy estimation problem in this setting is famously QMA-complete [KSV02; GHLS14], restricting to Hamiltonians that admit specific guiding states makes it QCMA-complete [WFC24] or MA-complete [Jia25]. Turning to the *trusted* case, where a guiding state is given as input (as in the present work), GGSE is BQP-complete when the required precision is $1/\operatorname{poly}(m)$ and the guiding state is chosen adversarially [CFG+23; GG22]. For weaker precision, the problem becomes solvable in BPP [GG22; Gal25].

Stoquastic barriers. A further restriction on Hamiltonians is to assume them to be *stoquastic*, namely, that all their off-diagonal entries are real and non-positive in some given basis. Such Hamiltonians are particularly interesting because of their classical flavor. Indeed, the ground state of a stoquastic Hamiltonian can be chosen to have non-negative real entries, which allows it to be interpreted as a probability distribution. Most Hamiltonians arising in quantum architectures like D-Wave come from Transverse Field Ising Models (TIM), which are stoquastic in the standard basis. Other examples of stoquastic Hamiltonians include the Heisenberg model on bipartite graphs and the bosonic Hubbard model. Since a classical objective function can be seen as a diagonal Hamiltonian and hence stoquastic, the ground-state properties of stoquastic Hamiltonians are of considerable importance in optimization. The general energy estimation problem becomes StoqMA-complete when H is stoquastic, with $\text{MA} \subseteq \text{StoqMA} \subseteq \text{SBP} \cap \text{QMA}$ [BDOT08; BBT06]. In the untrusted setting, the problem collapses to MA-complete under various assumptions about the guiding states [Bra15; Liu21]. In the trusted setting, the problem is only known to be BPP-hard [Wai25]. We note that some of these results, such as the AM and SBP containments [BDOT08; BBT06], or the recent StoqMA-completeness result of [GR26],

do not require the Hamiltonian to be local, but only sparse (as is the case for the Hamiltonian constructed in our work).

Connection to Adiabatic Quantum Computation. Adiabatic quantum computation is a leading framework for solving ground-state preparation problems on a quantum computer when a target Hamiltonian has no available guiding state initially [AT07]. The paradigm relies on designing a sequence (or *path*) of slowly varying Hamiltonians whose spectral gap remains open, ensuring – by the adiabatic theorem – that the system stays close to its instantaneous ground state throughout the evolution. From a computational perspective, this process can be viewed as a *bootstrapped* guided ground-state preparation algorithm, where access to intermediate ground states enables the preparation of the final target state. The possibility of dequantizing adiabatic algorithms has been raised in several prior works. For instance, [BDOT08] asked whether *"every efficient adiabatic path using stoquastic Hamiltonians can be simulated by a polynomial-time probabilistic machine"*. Positive results in this direction include a BPP algorithm for simulating any path of stoquastic frustration-free Hamiltonians [BT10], and a PostBPP algorithm for simulating paths of general stoquastic Hamiltonians [BDOT08]. On the negative side, recent breakthrough results rule out the possibility of dequantizing a *specific* stoquastic path in BPP, relative to an oracle [Has21; GHV21]. Our results align with these latter works but have broader implications: we rule out any dequantization of an adiabatic path – stoquastic or not – that must pass through the stoquastic Hamiltonian constructed in our paper.

Connection to QMA vs QCMA separation oracles. A long-standing open question in quantum complexity theory asks whether quantum witnesses are more powerful than classical witnesses in quantum verification protocols. Several works have attacked this question via oracle separations [AK07], culminating recently in positive results relative to oracles inspired by Forrelation [BHNZ26] and the Yamakawa-Zhandry problem [BHV26]. Some prior candidates for such separations are based on expander mixing problems with graph oracles [Lut11; NN24; BFM23; LMY25]. A key difference with our construction is that the graphs are degree-regular, so the ground energy (and ground state) of their negative adjacency matrices are trivial to compute. These works instead rely on testing the second-lowest eigenvalue. In general, there is presumably a gap in hardness between computing the ground and excited-state energies of stoquastic Hamiltonians, with the latter being QMA-complete rather than StoqMA-complete [JGL10]. In contrast to these works, our construction poses a genuine ground-state problem, and may lead to a tighter oracle separation between StoqMA and QCMA (or MA).

Connection to other exponential speedups. Our contribution can also be viewed as part of a line of work showcasing exponential quantum speedups based on oracular graph constructions with particular spectral properties. A foundational result in this direction is due to Childs et al. [CCD+03] for the so-called *welded-tree problem*. These ideas have since been extended to demonstrate exponential quantum speedups in a variety of settings, including quantum-walk-based search algorithms [CCD+03; LLL24; BLH25; JZ23; LZ25], quantum property testing [BCG+24], quantum annealing [SNK12], adiabatic computing [Has21; GHV21], quantum distributed computing [GLMR25], and quantum optimization [LWWZ25].

1.1 Contributions

Our first main result is a no-go theorem for the efficient dequantization of GGSP for stoquastic Hamiltonians, even when the algorithm has access to an arbitrarily good guiding state. Below, we give an informal definition of the problem.

Definition 2.1 (Informal definition of GGSP). *We are given an m-qubit stoquastic Hamiltonian H with a unique* ground state $|\psi_H\rangle$ *and a spectral gap of at least* $1/\operatorname{poly}(m)$*. In addition,*

we are given independent copies of a guiding state $|\psi_{\text{in}}\rangle$ *satisfying* $|\langle\psi_{\text{in}}\,|\,\psi_H\rangle| \geq 1/\operatorname{poly}(m)$. *The goal is to output a* refined state $|\psi_{\text{out}}\rangle$ *such that* $|\langle\psi_{\text{out}}\,|\,\psi_H\rangle| \geq 3/4$.

For classical algorithms, the input states are replaced with i.i.d. samples from the distribution $(\langle x\,|\,\psi_{\text{in}}\rangle^2)_x$. *The output must be a sample from the distribution* $(\langle y\,|\,\psi_{\text{out}}\rangle^2)_y$ *and it must be approximately statistically independent of the input samples.*

Note that we allow guiding states that need not be easy to prepare, including guiding states that are exactly equal to the ground state. This generality enables us to understand the power of a wide range of guiding-state preparation procedures, such as more general adiabatic paths than those previously considered (e.g., in [Has21; GHV21]) and arbitrary ansätze used as guiding states.

We focus on GGSP because the ability to sample from or prepare the ground state is useful for a variety of tasks beyond energy estimation, such as estimating expectation values of general observables on the ground state, or simulating its real-time dynamics. It also helps decouple the complexity of these problems from the hardness of preparing good guiding states. For instance, the dequantized energy estimation algorithm of [Gal25] runs in $\operatorname{poly}\big(\frac{1}{|\langle\psi_{\text{in}}\,|\,\psi_H\rangle|^{1/\Delta}}, m\big)$ (where Δ is the desired energy precision), which could be improved given the ability to boost the overlap $|\langle\psi_{\text{in}}\,|\,\psi_H\rangle|$ via an independent GGSP algorithm. Our result partly rules out this possibility by establishing classical hardness of GGSP for a stoquastic Hamiltonian, even in the presence of an arbitrary guiding state. Informally, we establish this classical hardness by proving the following result.

Theorem 3.4 and Proposition 3.3 (Informal). *There exists an m-qubit stoquastic Hamiltonian H for which the GGSP problem can be solved efficiently (in* $\operatorname{poly}(m)$ *time) by a quantum algorithm, whereas no classical algorithm can solve it using at most 2^{m^c} input samples and queries to H, for some constant $c < 1$. This result holds* for any *guiding state provided as input.*

Our second main result is a similar sub-exponential lower bound for classical algorithms solving the ground-state energy estimation problem. We give an informal description of the problem below.

Definition 2.3 (Informal definition of GGSE). *We are given an m-qubit stoquastic Hamiltonian H with a* ground state $|\psi_H\rangle$. *In addition, we are given copies of a* guiding state $|\psi_{\text{in}}\rangle$ *satisfying* $|\langle\psi_{\text{in}}\,|\,\psi_H\rangle| \geq 1/\operatorname{poly}(m)$. *The input is promised to satisfy exactly one of the following:* YES *case:* $\lambda_{\min}(H) \leq a$ *or* NO *case:* $\lambda_{\min}(H) \geq b$ *with* $b - a \geq 1/\operatorname{poly}(m)$. *The goal is to decide between the two cases on every input satisfying the promise. For classical algorithms, the input states are replaced with i.i.d. samples from the distribution* $(\langle x\,|\,\psi_{\text{in}}\rangle^2)_x$.

Existing dequantization results for GGSE (e.g., [GG22; Gal25]) rely on two additional assumptions: query access to the amplitudes of the guiding state and a constant precision $b - a = \Omega(1)$. On the other hand, these classical algorithms apply to arbitrary sparse Hamiltonians. Our result shows that restricting the Hamiltonian to be stoquastic is not enough to eliminate these assumptions. In particular, this suggests that in settings where amplitude access is not naturally available – such as adiabatic quantum computation or certain constructions aimed at oracle separations between quantum and classical proofs – efficient classical algorithms should not be expected for arbitrary stoquastic Hamiltonians. Informally, we prove the following.

Theorem 4.3 and Proposition 4.2 (Informal). *There exists an m-qubit stoquastic Hamiltonian H for which the GGSE problem can be solved efficiently (in* $\operatorname{poly}(m)$ *time) by a quantum algorithm, whereas no classical algorithm can solve it using at most 2^{m^c} input samples and queries to H, for some constant $c < 1$. This result holds for a guiding state $|\psi_{\text{in}}\rangle$ satisfying* $|\langle\psi_{\text{in}}\,|\,\psi_H\rangle| \geq 1 - 1/\operatorname{poly}(m)$.

We stress that, unlike our result for GGSP, the classical lower bound for GGSE holds only for a specific choice of the guiding state. This restriction is necessary because certain guiding states can make the GGSE problem trivial to solve (for instance, by hardcoding the bit corresponding to the answer). We also note that Hastings' result [Has21] can be interpreted as a classical hardness result for GGSE when the guiding state has inverse-polynomial overlap, $1/\operatorname{poly}(m)$, with the ground state. Our construction demonstrates that classical hardness persists even when the overlap is as large as $1 - 1/\operatorname{poly}(m)$.

1.2 Proof overview

1.2.1 Guided Ground-State Preparation

We provide a high-level overview of the ideas behind the proof of Theorem 3.4. Our result holds relative to an oracle representing the adjacency matrix of a graph G constructed in this paper.

Underlying Graph. We construct a graph G (Definition 2.10 and Figure 1) that is obtained by taking a graph E which is a high-degree, high-girth spectral expander and then attaching many copies of self-similar trees $\widehat{T}$ to its nodes (these are similar to the trees studied in [Has21; GHV21]). G is constructed in a way that it is degree-regular everywhere except at the leaves. The expander results in G having a large spectral gap and no short cycles (by definition, the girth is the length of the shortest cycle in the graph). The stoquastic Hamiltonian H acting on m qubits is the negative of the adjacency matrix of G over 2^m vertices, i.e., $H = -A_G$. The only assumption on the guiding state $|\psi_{\text{in}}\rangle$ is that it has at least inverse-polynomial overlap with $|\psi_H\rangle$ since, without this assumption, even a quantum computer may not be able to efficiently solve the GGSP problem. We show that *no such guiding state* can simultaneously enable an efficient classical algorithm.

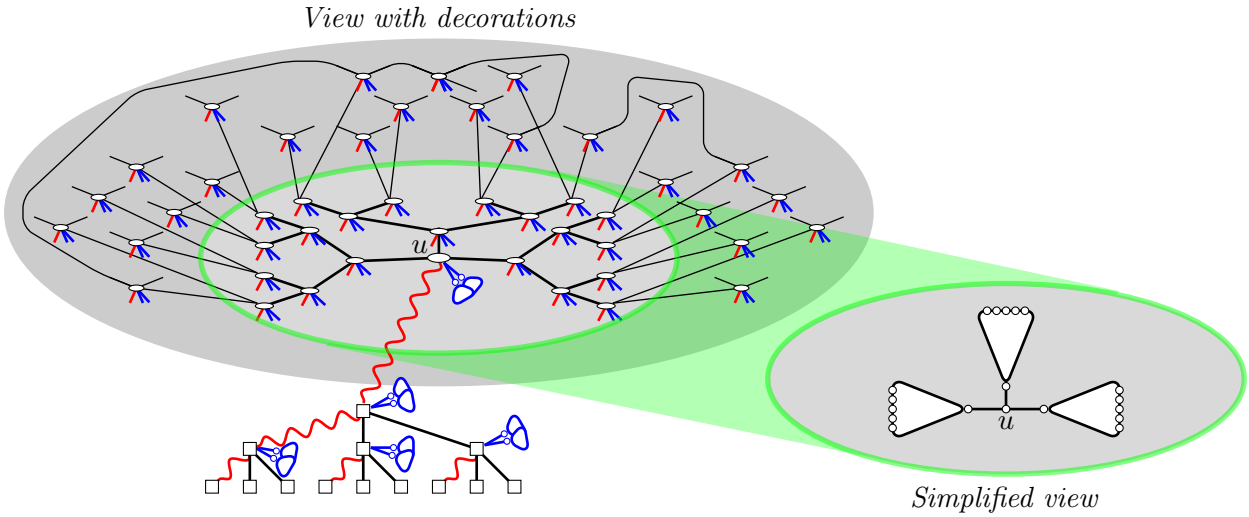


Figure 1: Pictorial representation of the graph G, with the expander E shown inside the gray ellipse (circular vertices) and details of the self-similar trees attached to a vertex u (square vertices). The expander is locally tree-like within a ball of radius less than half the girth around u (shown by the green boundary). The graph G locally resembles a self-similar tree – see Lemma 2.15. The simplified view depicts the local tree structure around u, drawing the incident edges within the expander, but omitting the decorations arising from the blue and red edges.

Hardness of classical exploration. The graph G is designed so that it is difficult for classical algorithms to distinguish locally between the expander part and the tree part (Lemmas 3.12 and 2.15). Indeed, since the expander has high girth and the final graph is degree-regular except at the leaves, it is difficult to distinguish the expander vertices from the tree vertices in a small neighborhood. Moreover, the self-similar trees have a sufficiently confusing structure that makes it difficult to infer one's position in the graph from their leaves (Lemma 3.13 and Section C). For instance, a simple random walk on G is insufficient for sampling from the ground state of H, since its stationary distribution is the degree distribution rather than the ground-state distribution, with the latter being far more concentrated on the vertices close to the expander. While more sophisticated classical algorithms might sample from the correct distribution, our main result is to show that none can do so efficiently.

Localization property. The classical hardness is established by examining a specific *localization* property of the ground state $|\psi_H\rangle$, which must also be approximately satisfied by the guiding state $|\psi_{\text{in}}\rangle$ and the output state $|\psi_{\text{out}}\rangle$. This property states that, upon sampling sub-exponentially many vertices from the ground-state distribution, the distance between their *nearest vertices* on the expander graph E is large with very high probability (Lemma 3.6). This follows essentially from the fact that the marginal ground-state distribution over the expander vertices is uniform (Lemma 2.12). The localization property must also hold for the distributions arising from $|\psi_{\text{in}}\rangle$ and $|\psi_{\text{out}}\rangle$ with constant probability (Proposition 3.10).

Hardness of classical sampling. We demonstrate that any classical algorithm making fewer than 2^{m^c} queries to H for some $c < 1$ (i.e., sub-exponentially many queries) cannot reproduce the above localization property in its output distribution with more than exponentially small probability (Proposition 3.14). By standard arguments, a classical algorithm can be assumed to explore the graph G by growing connected components around the input sampled vertices (Lemma 3.11). However, since the graph is difficult to explore, none of these connected components can reach an expander vertex sufficiently far from its starting point (Lemmas 3.12 and 3.13). It follows that such an algorithm cannot satisfy the localization property with high probability, and therefore cannot solve the GGSP problem on H.

1.2.2 Guided Ground-State Energy Estimation

We provide a quick overview of the main ideas behind the proof of Theorem 4.3, highlighting the new ingredients needed there. The lower bound is proved by constructing two graph families whose negative adjacency matrices define two families of stoquastic Hamiltonians corresponding to the YES and NO instances of GGSE. Thus, solving GGSE on these instances amounts to distinguishing between the two graphs.

Main and envelope graphs. The graph chosen for the YES instance is the same *main graph* G used in the GGSP lower bound. The graph for the NO instance, which we call the *envelope graph* F (Definition 2.10), is designed to be locally similar to G while being spectrally separated from it. Like G, it is obtained from a high-degree, high-girth graph E' by attaching many copies of self-similar trees $\widehat{T}$ to its vertices. Structurally, the main difference between G and F is that the base graph E' of F has smaller degree than the base graph E of G (which gives it a smaller top eigenvalue), while additional self-similar trees are attached to it. In particular, the internal vertices of F have the same degree as those of G. The additional trees are chosen so that the local structure of F mimics that of G up to a certain distance, making the two graphs locally indistinguishable.

Local indistinguishability. The graphs are constructed so that a classical algorithm *starting at an expander vertex* cannot distinguish G from F without exploring sufficiently far. The key novelty establishing this property is an embedding of the local graph structure around any expander vertex of G into the corresponding structure in F (Corollary 4.5). The two local structures agree up to a large distance. Consequently, the adjacency-list responses seen by a local exploration have the same distribution in the two graphs until the exploration reaches a sufficiently far-away vertex. The hardness of such an exploration then follows by the same arguments as in the GGSP lower bound. This holds even when multiple expander vertices are provided as input, as long as they are sufficiently far apart from one another (Proposition 4.6).

Hardness of classical energy estimation. Finally, it remains to show that the graphs are indistinguishable even to algorithms given a guiding state distribution. Unlike in the GGSP lower bound, this distribution cannot be chosen arbitrarily, since it may place too much probability on vertices whose local structures are distinguishable. We show that the uniform superposition over the expander vertices is a valid guiding state, with overlap $1 - O(1/n)$ with the ground state (Lemma 2.12). Under this distribution, the input samples are far-away expander vertices with high probability, which lets us apply the hardness result above. We conclude that the GGSE problem on G and F cannot be solved with probability exceeding $1/2$ by more than a sub-exponentially small amount.

2 Setup and graph construction

2.1 Guided Ground-State Preparation (GGSP) problem

The Guided Ground-State Preparation problem (GGSP) is formally defined below. In the quantum setting, the goal is to prepare the ground state of a gapped Hamiltonian, given access to copies of a guiding state. In the classical setting, quantum states are replaced by classical samples drawn from the corresponding standard basis measurement distribution. We use $\mathrm{poly}(m)$ as shorthand for m^c for some absolute constant c. The constants $3/4$ and $1/5$ in the output condition are chosen arbitrarily.

Definition 2.1. The *Guided Ground-State Preparation* (GGSP) problem is defined as follows. We are given an m-qubit stoquastic Hamiltonian H with bounded norm $\|H\| \leq \mathrm{poly}(m)$, which has a unique *ground state* $|\psi_H\rangle$ and a spectral gap at least $1/\mathrm{poly}(m)$. In addition, we are given the description of t copies of a *guiding state* $|\psi_{\mathrm{in}}\rangle$ that satisfies $|\langle\psi_{\mathrm{in}}\,|\,\psi_H\rangle| \geq 1/\mathrm{poly}(m)$. We should output the description of a *refined state* $|\psi_{\mathrm{out}}\rangle$ such that $|\langle\psi_{\mathrm{out}}\,|\,\psi_H\rangle| \geq 3/4$. The descriptions of the guiding and refined states are given in one of the following two models:

- *(Quantum input-output state)* The algorithm is given the quantum state $|\psi_{\mathrm{in}}\rangle^{\otimes t}$ and must output the quantum state $|\psi_{\mathrm{out}}\rangle$.
- *(Classical input-output sample)* Let p_{in} and p_{out} denote the probability distributions on $\{0,1\}^m$ obtained by measuring $|\psi_{\mathrm{in}}\rangle$ and $|\psi_{\mathrm{out}}\rangle$, respectively, in the standard basis. The algorithm is given i.i.d. samples $x_1, \ldots, x_t \sim p_{\mathrm{in}}$, and must output one sample $y \sim p_{\mathrm{out}}$ such that the joint distribution of $(x_1, \ldots, x_t, y)$ is $1/5$-close in total variation distance to the product distribution $p_{\mathrm{in}}^{\otimes t} \times p_{\mathrm{out}}$.

In the classical setting, the conditions on the input-output states could equivalently be defined in terms of the *fidelities* $F(p_{\mathrm{in}}, p_H) = |\langle\psi_{\mathrm{in}}\,|\,\psi_H\rangle|^2 \geq 1/\mathrm{poly}(m)$ and $F(p_{\mathrm{out}}, p_H) = |\langle\psi_{\mathrm{out}}\,|\,\psi_H\rangle|^2 \geq (3/4)^2$, where $p_H(x) = |\langle x\,|\,\psi_H\rangle|^2$ is the ground-state distribution.

The GGSP problem can be solved quantumly by filtering out the ground-state component of the guiding state. Because this state already has non-negligible overlap with the ground state, quantum algorithms can rapidly amplify this overlap, something classical algorithms cannot do efficiently in general. Below, we describe an elementary approach to achieve this goal.

Proposition 2.2 (Quantum algorithm for the GGSP problem)**.** *There exists a quantum algorithm that solves the GGSP problem by outputting the exact ground state* $|\psi_{\text{out}}\rangle = |\psi_H\rangle$ *of* H *with cost* $\text{poly}(m) \cdot T_H$*, where* T_H *is the cost of simulating the unit-time evolution* e^{-iH}*.*

Proof. The algorithm runs quantum phase estimation on H using the guiding state $|\psi_{\text{in}}\rangle$ as input. Because the spectral gap is at least inverse polynomial, it suffices to use inverse-polynomial precision to distinguish the smallest-eigenvalue component $|\psi_H\rangle$ from the others. Measuring the energy register then collapses the state to that component with probability equal to the squared overlap $|\langle\psi_{\text{in}} \,|\, \psi_H\rangle|^2 \geq 1/\,\text{poly}(m)$. The success probability can be amplified arbitrarily close to 1 by repeating the procedure on fresh copies of $|\psi_{\text{in}}\rangle$ until the state collapses to $|\psi_H\rangle$. □

The simulation cost T_H appearing in the quantum algorithm above depends on the properties and input specification of H. It is polynomial in m for several standard classes of Hamiltonians, including local Hamiltonians, sparse Hamiltonians given oracle access to their nonzero entries (as in the present work), and, more generally, Hamiltonians admitting an efficient block-encoding.

On the independence requirement.

A subtle but important requirement of Definition 2.1 – in the classical setting – is that the output sample y must be (almost) independent of the input samples $x_1, \ldots, x_t$. Without this condition, an algorithm could sometimes reuse the same input as output (for example, when the guiding state already coincides with the ground state). This would pose a significant practical limitation. For instance, standard concentration bounds, such as Chernoff or Chebyshev inequalities, rely on (nearly) independent samples to hold. Producing multiple correlated samples would therefore prevent estimation of statistics of the ground-state distribution via such bounds: if we estimate an expectation value $\langle\psi_H|O|\psi_H\rangle$ via the empirical mean $\frac{1}{t}\sum_i O(y_i)$, where $y_1, \ldots, y_t$ are correlated samples, then the estimator's variance may not shrink with t at all. This obstruction is formalized by our subsequent lower bound for ground-state energy estimation (Theorem 4.3), which shows that estimating $\langle\psi_H|H|\psi_H\rangle$ is no easier than outputting independent samples.

A similar consideration is standard in the classical Markov chain Monte Carlo (MCMC) literature. For instance, when estimating partition functions via simulated annealing, Markov chains are initialized from a warm-start distribution (playing the role of the guiding state) and run long enough to produce fresh, independent samples from the stationary distribution (see, e.g., [JSV04, Section 6] and [ŠVV09, Section 7]).

One may wonder why we do not impose a similar independence requirement in the quantum setting. Due to the no-cloning theorem, it is generally impossible to produce the ground state without entirely consuming a copy of the guiding state. Nevertheless, this obstruction does not preclude certain types of quantum statistical estimators that behave differently from classical ones. For example, quantum speedups for MCMC methods can rely on non-destructive quantum primitives [HW20; CH23], which make it possible to use a single copy of the ground state $|\psi_H\rangle$ restored at the end of the computation, together with multiple applications of the reflection $\mathbb{I} - 2|\psi_H\rangle\langle\psi_H|$.

2.2 Guided Ground-State Energy Estimation (GGSE) problem

The Guided Ground-State Energy-Estimation problem (GGSE) is formally defined below.

Definition 2.3 (Guided Ground-State Energy Estimation)**.** The *Guided Ground-State Energy Estimation* (GGSE) problem is defined as follows. We are given oracle access to an m-qubit stoquastic Hamiltonian H satisfying $\|H\| \leq \text{poly}(m)$, with a ground state $|\psi_H\rangle$. In addition, we are given the description of t copies of a *guiding state* $|\psi_{\text{in}}\rangle$ that satisfies $|\langle\psi_{\text{in}} \,|\, \psi_H\rangle| \geq 1/\,\text{poly}(m)$ (in the same input model as in Definition 2.1). The lowest eigenvalue $\lambda_{\min}(H)$ of the Hamiltonian

is promised to satisfy exactly one of the following:

$$\text{YES: } \lambda_{\min}(H) \leq a \qquad \text{NO: } \lambda_{\min}(H) \geq b,$$

where $b - a \geq 1/\operatorname{poly}(m)$. The goal is to output 1 in the YES case and 0 in the NO case on every input satisfying the promise.

For the sake of completeness, we provide a quick proof of the fact that there exists an efficient quantum algorithm for this problem by applying the quantum phase estimation algorithm.

Proposition 2.4 (Quantum algorithm for GGSE)**.** *There exists a quantum algorithm that solves the GGSE problem with success probability at least* $2/3$ *and at cost* $\operatorname{poly}(m) \cdot T_H$*, where* T_H *is the cost of simulating the unit-time evolution* e^{-iH}*.*

Proof. The algorithm applies quantum phase estimation to H with input $|\psi_{\text{in}}\rangle$ and precision below $b - a$. Repeating the procedure with polynomially many input copies and taking the minimum measured energy yields the correct ground-state energy with constant probability. We note that, unlike in GGSP, the algorithm does not require H to have a large spectral gap. □

2.3 Graph construction

We describe the graph $G_{n,E}$ that underlies both of our lower bounds. The graph combines a high-degree, high-girth spectral expander with a particular family of self-similar trees. We use a free integer parameter n throughout the construction, which will later be chosen as $n = \operatorname{poly}(m)$ when defining the corresponding Hamiltonians.

We first state the properties required of the spectral expanders. A key difficulty is that the vertex degree must grow (logarithmically) with the number of vertices, which rules out constructions with constant degree.

Proposition 2.5 (High-degree, high-girth expanders E_n)**.** *There exists a family of regular graphs* $(E_n)_n$ *such that, for even* n *large enough,* E_n *has* $N_E \in [2^{100n^{5/2}\log^2 n}, 2^{n^3 \log^2 n}]$ *vertices, degree* $d_n = n$*, spectral gap at least* $n/2$ *and girth at least* $g_n = 100n^{5/2}\log n$*.*

Proof. We provide two constructions of such graphs in Section A, one based on the number-theoretic Ramanujan graphs from [LPS88], and the other on the short cycle removal algorithm from [Par21]. The first construction offers a more explicit description of the expander, but only works for even degrees. The second construction provides an algorithm that runs in time polynomial in the number of vertices. □

We expect that the graph constructed in [GHV21] can be derandomized by relying on similar expanders as above, together with high-girth bipartite regular graphs (such as those introduced in [Dah14]) in place of the random bipartite matchings used in that work.

If the graphs do not need to be expanders, as is the case in Section 4 for GGSE, other constructions of high-girth graphs become available. We describe such a sequence of graphs that is used later.

Proposition 2.6 (High-degree, high-girth graphs E'_n)**.** *There exists a family of regular graphs* $(E'_n)_n$ *such that, for* n *large enough,* E'_n *has* $N_E \in [2^{100n^{5/2}\log^2 n}, 2^{n^3 \log^2 n}]$ *vertices, degree* $n - \sqrt{n}$ *and girth at least* $g_n = 100n^{5/2}\log n$*.*

Proof. An early result of Erdős and Sachs [ES63; EJ08] establishes the existence of d-regular graphs on N vertices with girth at least g, provided that $d \geq 3$, $g \geq 4$, $N \geq 4(d-1)^{g-1}$, and N is even. A random d-regular graph has nonzero probability, albeit very small, of having girth at least g whenever $(d-1)^{2g-1} = o(N)$ [MWW04, Corollary 1]. A combinatorial construction of high-girth Cayley graphs is also given in [Dah14]. □

Next, we describe a family of trees with a fractal-like structure, which we call *self-similar trees*. The construction is based on the notion of *graph decoration*: given two graphs G and G', one attaches a distinct copy of G' to each vertex of the base graph G by adding a new edge connecting that vertex to a distinguished vertex of the corresponding copy of G'. Such a process has been studied in the context of stoquastic Hamiltonian constructions in [Has21; GHV21], and our construction of the trees bears strong similarities to these works. One difference is that we do not decorate the leaves of the base graph. This has the advantage of making the final graph slightly smaller, and ensures that it is degree-regular everywhere except at its root and its leaves.

Definition 2.7 (Self-similar trees $\widehat{T}_{n,k}$)**.** For $n \geq 1$ and $1 \leq k \leq \sqrt{n}+1$, consider the degree sequence $d_{n,k} = 2n - k\sqrt{n}$ and the depth sequence $\ell_{n,k} = k \cdot 10n^{3/2} \log n$. Define the following trees,

- The *perfect tree* $C_{n,k}$ is the tree in which every internal vertex has exactly $d_{n,k} - 1$ children and all leaves are at depth $\ell_{n,k}$. The 0*-level decorated tree* $T_{n,k,0}$ is defined to be $C_{n,k}$.
- The r*-level decorated tree* $T_{n,k,r}$ is defined iteratively, for $r = 1, \ldots, k-1$, by adding $d_{n,k-r} - d_{n,k-r+1}$ edges to each internal node of $T_{n,k,r-1}$ and attaching the other endpoint of each edge to the root of a new copy of $C_{n,k-r}$.
- The *full-level decorated tree* $\widehat{T}_{n,k}$ is defined to be $T_{n,k,k-1}$.

The r*-level leaves* is the set of all new leaves added in the construction of $T_{n,k,r}$ by decorating $T_{n,k,r-1}$.

The above definition is equivalent to the following construction, which will prove convenient for some of the later analysis.

Fact 2.8 (Bottom-up construction of $\widehat{T}_{n,k}$)**.** *The full-level decorated tree $\widehat{T}_{n,k}$ can equivalently be obtained by, for each $r = 1, \ldots, k-1$, adding $d_{n,r} - d_{n,r+1} = \sqrt{n}$ edges to each internal node of $C_{n,k}$ and attaching the other endpoint of each edge to the root of a new copy of $\widehat{T}_{n,r}$. The base case is $\widehat{T}_{n,1} = C_{n,1}$.*

These trees have the following properties.

Fact 2.9. *All vertices in $T_{n,k,r}$ (except the root and the leaves) have degree $d_{n,k-r}$. The number of vertices in $\widehat{T}_{n,k}$ is at most $2^{12n^3 \log^2 n}$.*

We now describe the graph $G_{n,E}$, which is obtained by decorating the expander defined above with self-similar trees. We also add dummy isolated vertices for technical reasons, following prior work on classical graph exploration (e.g., [CCD+03]). These isolated vertices are used to force classical algorithms to explore a connected component, since otherwise a queried vertex would be isolated with very high probability.

Definition 2.10 (Decorated graph $G_{n,E}$)**.** Given an integer $n \geq 1$, define the graph constructed as follows:

Fix a pair $(K, E) \in \{(\sqrt{n}, E_n), (\sqrt{n}+1, E'_n)\}$ (see Propositions 2.5 and 2.6).

For $k = 1, \ldots, K-1$, add $d_{n,k} - d_{n,k+1} = \sqrt{n}$ edges to each node of E and attach the other endpoint of each edge to the root of a new copy of $\widehat{T}_{n,k}$ (see Definition 2.7).

Pad the construction with $2^{15n^3 \log^2 n} - N_E \cdot (1 + \sqrt{n} \cdot \sum_{k=1}^{K-1} |\widehat{T}_{n,k}|)$ isolated vertices.

We let $G_{n,E}$ denote the graph constructed above. We call $G_n = G_{n,E_n}$ and $F_n = G_{n,E'_n}$ the *main graph* and the *envelope graph*, respectively.

We represent the vertices of the graph $G_{n,E}$ by the sets $V = V_E \cup V_T \cup V_I$, where V_E are the vertices coming from the subgraph E (called the *expander vertices*), V_T are the vertices coming from the trees $\widehat{T}_{n,k}$ (called the *tree vertices*), and V_I are the other *isolated vertices*. When there is risk of ambiguity between different graphs, we write $V(G_{n,E}) = V$ (and similarly for the other vertex sets).

For each non-isolated vertex $v \in V_E \cup V_T$, we let $\mathrm{ex}(v) \in V_E$ denote the closest expander vertex to v in the graph (this is v itself when $v \in V_E$, and the expander vertex to which the tree containing v is attached when $v \in V_T$).

The adjacency matrix of $G_{n,E}$ is denoted by $A_{n,E}$, and its top eigenvector by $|\psi_{n,E}\rangle$.

Fact 2.11. *All vertices in $G_{n,E}$ (except the leaves and the isolated vertices) have degree $d_{n,1}$. The graph $G_{n,E}$ has the same girth as E. The number of vertices is $|V| = 2^{15n^3 \log^2 n}$.*

There are two crucial properties of $G_{n,E}$ that will play a central role in both of our lower-bound arguments. First, $G_{n,E}$ inherits the spectral properties of the subgraph E, in the sense that it has a large spectral gap if E does, and that its top eigenvector has a large overlap with the top eigenvector of E. We use the notation ψ (rather than $|\psi\rangle$) to denote a state that is not necessarily normalized.

Lemma 2.12 (Spectral properties of the adjacency matrix $A_{n,E}$)**.** *The difference between the two largest eigenvalues of A_{n,E_n} is at least $d_n/2 - 4\sqrt{2n}$.*

The top eigenvector $\psi_{n,E}$ of $A_{n,E}$, where $E \in \{E_n, E'_n\}$, can be written as $\psi_{n,E} = \psi_E + \sum_{v \in V_E} \psi_{v,T}$, where ψ_E is a uniform (unnormalized) superposition over the expander vertices $\{v\}_{v \in V_E}$, and each $\psi_{v,T}$ is supported on the vertices of the trees attached to v. Moreover, the states $\{\psi_{v,T}\}_{v \in V_E}$ are identical up to a permutation of the computational basis.

The base component satisfies $\|\psi_E\|^2 \geq (1 - O(1/n))\|\psi_{n,E}\|^2$.

Proof. The proof is deferred to Section B. It relies on generic spectral properties of decorated graphs, similar to those used in [Has21; GHV21]. □

The second property is that, locally, $G_{n,E}$ resembles the tree $\widehat{T}_{n,K}$ along any edge within the expander graph (for the remainder of this section, K denotes the index paired with E in Definition 2.10). This will make it very difficult for classical algorithms to distinguish between these two graphs, although their spectral properties are radically different. We illustrate a simplified version of that in Figures 1 and 2 and in the next lemma. A stronger version is given in Figure 3 and Lemma 3.12, where we consider multiple connected components rooted at arbitrary vertices. The property is based on the following distance measure dist_E defined on the graph $G_{n,E}$.

Definition 2.13 (Distance measure)**.** Given two vertices $u, v \in V$, the distance $\mathrm{dist}_E(u, v)$ is defined as the number of *expander* edges along a shortest path connecting u and v in $G_{n,E}$ (with the convention $\mathrm{dist}_E(u, v) = +\infty$ if either u or v is an isolated vertex). Given a subset $U \subseteq V$ of vertices, we define $\mathrm{dist}_E(U, v) = \min_{u \in U} \mathrm{dist}_E(u, v)$.

Note that dist_E is not a metric, since the distance between two distinct vertices $u \neq v$ can be zero (if they belong to trees attached to the same expander vertex). Nevertheless, it satisfies the triangle inequality. We use dist_E to study locality in the graph $G_{n,E}$ as follows.

In most of our arguments we work with a *truncated neighborhood* of a set of expander vertices $\mathcal{U}$. This is the subgraph of $G_{n,E}$ induced by all vertices (expander or tree) within distance r of $\mathcal{U}$, together with only the expander vertices at distance exactly $r + 1$. We define it formally here.

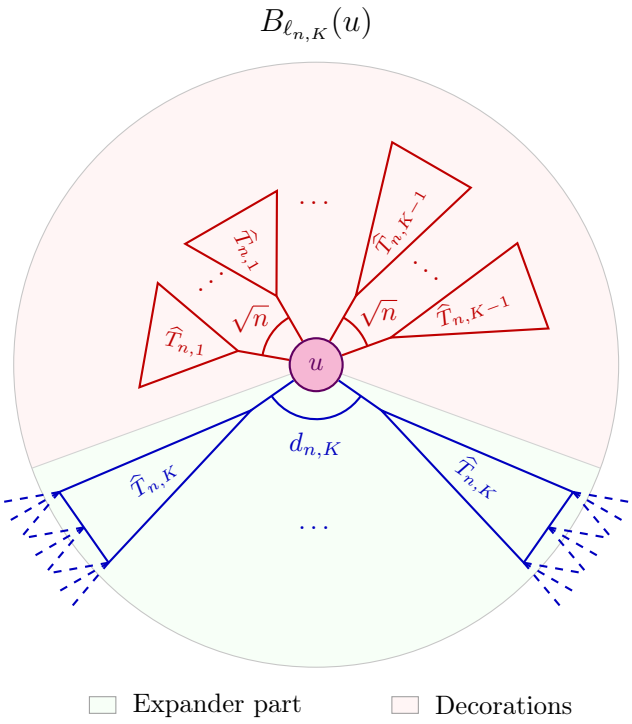


Figure 2: The truncated neighborhood $B_{\ell_{n,K}}(u)$ of an expander vertex $u \in V_E$.

Definition 2.14 (Truncated neighborhood)**.** Let $G \in \{G_n, F_n\}$ be the main or the envelope graph. For a set of expander vertices $\mathcal{U} \subseteq V_E(G)$ and an integer $r \geq 0$, let $B_{G,r}(\mathcal{U})$ be the subgraph of G induced by

$$\{v \in V_E(G) \cup V_T(G) : \mathrm{dist}_E(\mathcal{U}, v) \leq r\} \cup \{v \in V_E(G) : \mathrm{dist}_E(\mathcal{U}, v) = r + 1\}.$$

We write $B_{G,r}(v)$ for $B_{G,r}(\{v\})$ and drop G when it is clear from context.

Lemma 2.15 (Local structure of $G_{n,E}$)**.** *Consider the graph $G_{n,E}$, where K is the index paired with E in Definition 2.10, and let $u \in V_E$ be an expander vertex. Then the truncated neighborhood $B_{\ell_{n,K}}(u)$ (Definition 2.14) is isomorphic to the following graph (Figure 2): a central vertex $\widetilde{u}$ joined by an edge to the root of each of $\sqrt{n}$ copies of $\widehat{T}_{n,k}$ for every $k = 1, \ldots, K-1$, and of $d_{n,K}$ copies of $\widehat{T}_{n,K}$.*

Moreover, the isomorphism can be chosen so that u maps to $\widetilde{u}$ and the expander vertices at distance exactly $\ell_{n,K} + 1$ from u map to the 0-level leaves of the copies of $\widehat{T}_{n,K}$.

Proof. The subgraph $B_{\ell_{n,K}}(u)$ is necessarily a tree, since the subgraph E has no cycles of length less than g_n by the girth property, and $g_n > 2(\ell_{n,K} + 1) + 1$. Hence, the expander vertices in $B_{\ell_{n,K}}(u)$ form a full tree of depth $\ell_{n,K} + 1$ and degree $d_{n,K}$, to which are attached the decorations specified in Definition 2.10. Observe that, for any expander neighbor v of u, the subtree rooted at v matches exactly the definition of the self-similar tree $\widehat{T}_{n,K}$ from Definition 2.7 and Fact 2.8. □

2.4 Hamiltonian construction

We describe the stoquastic Hamiltonian encoding used throughout the paper. Given a graph G on 2^m vertices, we associate to it the negative of its adjacency matrix, with the vertices encoded

via a permutation. In our classical lower bounds, this permutation will be chosen uniformly at random. The random labeling prevents the algorithm from exploiting any canonical description of the graph and is used crucially in our classical lower bounds. The resulting Hamiltonian is specified to the algorithm via an adjacency-list oracle.

Definition 2.16 (Hamiltonian associated with a labelled graph)**.** Let G be a graph on 2^m vertices with adjacency matrix A_G, and let $\pi : V(G) \to [2^m]$ be a bijection. Define the m-qubit Hamiltonian

$$H_{G,\pi} = -P_\pi A_G P_\pi^{-1},$$

where P_π is the permutation matrix corresponding to π (i.e., $(P_\pi)_{x,u} = 1$ if $\pi(u) = x$, and $(P_\pi)_{x,u} = 0$ otherwise).

The Hamiltonian $H_{G,\pi}$ can be queried via the following adjacency-list oracles:

- *(Classical oracle)* Given $x \in [2^m]$, the oracle $\mathcal{O}_{G,\pi}$ returns the list $\mathcal{O}_{G,\pi}(x) = \Gamma_x \subset [2^m]$ of labels corresponding to the neighbors of $\pi^{-1}(x)$ in G (equivalently, the non-zero entries in the x-th row of $H_{G,\pi}$).
- *(Quantum oracle)* The oracle implements the unitary operation $\mathcal{O}_{G,\pi}|x\rangle|0\rangle \mapsto |x\rangle|\Gamma_x\rangle$.

When clear from the context, we may drop G from the subscript.

3 Hardness of Ground-State Preparation Problem

For the remainder of this section, we specialize the Hamiltonian construction of Definition 2.16 to the *main graph* $G_n = G_{n,E_n}$ derived from Definition 2.10 with $K = \sqrt{n}$, and we let $A_n = A_{n,E_n}$ denote its adjacency matrix. For m sufficiently large, let $n = \mathrm{poly}(m)$ be such that G_n has 2^m vertices. Let $\pi : V(G_n) \to [2^m]$ be a bijection. We write $H_\pi := H_{G_n,\pi} = -P_\pi A_n P_\pi^{-1}$, and denote its ground state by $|\psi_\pi\rangle$. The spectral properties of H_π follow directly from those of the graph G_n (Lemma 2.12). In particular, its ground state $|\psi_\pi\rangle$ coincides, up to the labeling π, with the top eigenvector $|\psi_n\rangle$ of A_n.

Fact 3.1 (Spectral properties of H_π)**.** *The Hamiltonian H_π is stoquastic, it has a non-degenerate ground space with spectral gap at least $n/2 - 4\sqrt{2n}$, its ground state is $|\psi_\pi\rangle = P_\pi|\psi_n\rangle$.*

Finally, given a guiding state $|\psi_{\mathrm{in}}\rangle$ for the unlabelled Hamiltonian $-A_n$, we associate the corresponding family of guiding states $\{|\psi_{\mathrm{in},\pi}\rangle\}_\pi$ for the Hamiltonians H_π.

Definition 3.2 (Guiding state $|\psi_{\mathrm{in},\pi}\rangle$)**.** Given a guiding state $|\psi_{\mathrm{in}}\rangle$ for the Hamiltonian $-A_n$, we let $|\psi_{\mathrm{in},\pi}\rangle = P_\pi|\psi_{\mathrm{in}}\rangle$ denote the corresponding guiding state for the Hamiltonian H_π.

We now prove our main result (Theorem 3.4): any classical algorithm that makes only a sub-exponential number of queries to the adjacency-list oracle of the Hamiltonian H_π, even when provided with sub-exponentially many copies of a guiding state $|\psi_{\mathrm{in},\pi}\rangle$, cannot sample from the ground-state distribution with sufficiently high precision. Before presenting the classical lower bound, we first show that quantum algorithms can solve the problem efficiently.

Proposition 3.3 (Quantum algorithm)**.** *There exists a quantum algorithm that solves GGSP on every Hamiltonian H_π (Definition 2.16), using* $\mathrm{poly}(n)$ *copies of a guiding state $|\psi_{\mathrm{in},\pi}\rangle$, and* $\mathrm{poly}(n)$ *quantum queries to the adjacency-list oracle of H_π.*

Proof. The graph G_n (and hence the Hamiltonian H_π) is $d_{n,1}$-sparse, where $d_{n,1} = O(n)$ (Fact 2.11). The algorithm follows immediately from Proposition 2.2, by noting that simulating a $O(n)$-sparse $\mathrm{poly}(n)$-qubit Hamiltonian H requires $T_H = \mathrm{poly}(n)$ quantum gates and queries to its adjacency-list oracle (see, for instance, [AT07, Lemma 1.4]). □

The main theorem we prove is that classical algorithms using only sub-exponentially many queries and input samples cannot succeed in solving the GGSP problem on H_π with high probability.

Theorem 3.4 (Classical hardness of GGSP)**.** *Let $a \geq 1$ be a constant and $(|\psi_{\mathrm{in},\pi}\rangle)_\pi$ a family of guiding states for $\{H_\pi\}_\pi$ (Definition 3.2) with $|\langle\psi_{\mathrm{in},\pi}\,|\,\psi_\pi\rangle| \geq 1/n^a$ for all π. Suppose that a classical algorithm is making q queries to the adjacency-list oracle of H_π, for a random labeling π, and is given t independent samples $x_1, \ldots, x_t \sim p_{\mathrm{in},\pi}$ from the distribution $p_{\mathrm{in},\pi}(x) = \langle x\,|\,\psi_{\mathrm{in},\pi}\rangle^2$.*

If $q + t < 2^{\sqrt{n}}$ then the algorithm cannot (with non-negligible probability over the choice of π) output a sample y from a distribution $p_{\mathrm{out},\pi}$ that is the measurement distribution of a refined state $|\psi_{\mathrm{out},\pi}\rangle$ satisfying $|\langle\psi_{\mathrm{out},\pi}\,|\,\psi_\pi\rangle| \geq 3/4$, and such that the joint distribution J_π of $(x_1, \ldots, x_t, y)$ is within total variation distance $1/5$ of $p_{\mathrm{in},\pi}^{\otimes t} \times p_{\mathrm{out},\pi}$.

The theorem is proved by analyzing a certain localization property of the ground state $|\psi_\pi\rangle$ of H_π, which must also be satisfied (at least to some extent) by the output state $|\psi_{\mathrm{out},\pi}\rangle$. It concerns the probability that the output vertex lies at distance at least $\ell_{n,K}$ from the *closure* of the input vertices. We measure distance using dist_E (Definition 2.13). Informally, the closure enlarges the input set until every short path between its vertices lies entirely within the set. Such closure operations also arise, for instance, in LP integrality gaps and sum-of-squares lower bounds [CMM09; BCK15].

Definition 3.5 (Closure of a set of vertices)**.** Let $G = (V, E)$ be a finite graph, let $r \geq 2$ be an integer, and let $U \subseteq V$. Starting from U, repeatedly apply the following rule for as long as it is applicable:

> If there is a path of length $\ell \in \{2, \ldots, r\}$ whose two endpoints are distinct vertices of the current set U and whose internal vertices all lie outside U, add its internal vertices to U.

Since each step strictly enlarges U and V is finite, the process terminates. The resulting set does not depend on the order in which the steps are performed. We call it the *closure* of U with parameter r and denote it by $\mathrm{clos}_{G,r}(U)$. We drop the subscript G when the graph is clear from context.

First, we show in Section 3.1 that any input-output distribution that correctly solves the GGSP problem satisfies the localization property with probability at least $1/10$.

Proposition 3.10 (Input-Output state localization)**.** *Fix any labeling π and suppose that a classical algorithm solves the GGSP problem on H_π under the conditions of Theorem 3.4. Let J_π denote the joint distribution of the input-output samples $(x_1, x_2, \ldots, x_t, y)$ of the algorithm under the vertex labeling π. Then, if $t < 2^{\sqrt{n}}$, the probability that the output vertex y lies at distance at least $\ell_{n,K}$ from $\overline{U} = \mathrm{clos}_{2\ell_{n,K}+3}(\mathrm{ex}(\pi^{-1}(\{x_1, x_2, \ldots, x_t\})))$ satisfies,*

$$\Pr_{(x_1,\ldots,x_t,y)\sim J_\pi}\left[\mathrm{dist}_E\left(\overline{U}, \pi^{-1}(y)\right) > \ell_{n,K}\right] \geq 1/10. \tag{1}$$

Conversely, we prove in Section 3.2 that for any classical algorithm using at most $2^{\sqrt{n}}$ queries and input samples, the probability of the above event is exponentially small. We stress that this result holds for a random π, in contrast to the proposition above which is valid for all π.

Proposition 3.14 (Hardness of exploring the graph G_n)**.** *Suppose that a classical algorithm is given access to H_π for a random labeling π, as well as input samples $x_1, \ldots, x_t$ from a distribution as specified in Theorem 3.4 with $t < 2^{\sqrt{n}}$. Suppose that it outputs an additional sample y, and let J_π denote the joint distribution of the input-output samples $(x_1, x_2, \ldots, x_t, y)$. If the algorithm makes fewer than $2^{\sqrt{n}}$ queries, then the probability that the output vertex y lies at distance at least $\ell_{n,K}$ from $\overline{U} = \mathrm{clos}_{2\ell_{n,K}+3}(\mathrm{ex}(\pi^{-1}(\{x_1, x_2, \ldots, x_t\})))$ satisfies,*

$$\Pr_{\pi,(x_1,\ldots,x_t,y)\sim J_\pi}\left[\mathrm{dist}_E\left(\overline{U}, \pi^{-1}(y)\right) > \ell_{n,K}\right] \leq 2^{-n\log n}. \tag{2}$$

We can now conclude the proof of the main theorem.

Proof of Theorem 3.4. Fix $|\psi_{\text{in}}\rangle$ to be any guiding state for the Hamiltonian $-A_n$. Consider a classical algorithm that solves the GGSP problem on every Hamiltonian H_π, using the guiding state $|\psi_{\text{in},\pi}\rangle = P_\pi|\psi_{\text{in}}\rangle$ as specified in Definition 3.2 and Theorem 3.4. Suppose, for the sake of contradiction, that it makes fewer than $2^{\sqrt{n}}$ queries. Then, by Proposition 3.14, for an overwhelming fraction of labelings π, its input-output distribution J_π violates Equation (1). It implies by Proposition 3.10 that the algorithm cannot be correct on H_π with non-negligible probability over the choice of π. □

3.1 Proof of ground-state localization properties

In this section, we lower bound the probability that a successful algorithm solving the GGSP problem samples a vertex at a large distance from its input. First, we prove the statement under the exact ground-state distribution.

Lemma 3.6 (Ground-state localization)**.** *Let $U \subseteq V$ be any subset of vertices of G_n. For any integer ℓ, the probability that a vertex sampled from the ground-state distribution $p_n(v) = \langle v \,|\, \psi_n\rangle^2$ of $-A_n$ lies at distance at least ℓ from U is at least,*

$$\Pr_{v\sim p_n}[\text{dist}_E(U,v) \geq \ell] \geq 1 - \frac{|U|\cdot n^\ell}{N_E}$$

where N_E is the size of the expander E_n in Proposition 2.5.

Proof. First, we prove the lemma when $U = \{u\}$ is of size 1. Let p_E be the conditional distribution of p_n on the set of expander vertices V_E. By Lemma 2.12, when v is sampled according to p_n, the corresponding expander vertex $\text{ex}(v)$ follows the distribution p_E. Hence,

$$\Pr_{v\sim p_n}[\text{dist}_E(u,v) < \ell] = \Pr_{v\sim p_E}[\text{dist}_E(u,v) < \ell]. \tag{3}$$

By Lemma 2.12, p_E is the uniform distribution over V_E. Moreover, by construction, the expander graph over V_E has degree $d_n = n$. Hence, the number of expander vertices v that are at distance less than ℓ from $\text{ex}(u)$ is at most $1 + n + n^2 + \cdots + n^{\ell-1} \leq n^\ell$. Consequently, the probability of sampling v at distance $< \ell$ from $\text{ex}(u)$, under the distribution p_E, can be upper bounded by the ratio between the number of expander vertices at distance less than ℓ from $\text{ex}(u)$ and the total number N_E of expander vertices:

$$\Pr_{v\sim p_E}[\text{dist}_E(u,v) < \ell] \leq \frac{n^\ell}{N_E}. \tag{4}$$

Equations (3) and (4) imply that $\Pr_{v\sim p_n}[\text{dist}_E(u,v) \geq \ell] \geq 1 - \frac{n^\ell}{N_E}$.

When U is of arbitrary size, the lemma follows by applying the above result together with a union bound over the $|U|$ vertices $u \in U$. □

The localization property we use is stated for the closure of a set of vertices (Definition 3.5). As a first step, we show that when the girth is large enough, the closure of a set U with $|U| = t$ and parameter r has size linear in t and r. A similar bound was shown in [BCK15, Lemma 4.4] under different conditions.

Lemma 3.7 (Size of closure)**.** *Let $U \subseteq V$ be a set of vertices such that $|U| = t \geq 1$ and $r \geq 2$ be an integer. Suppose G has girth at least $g \geq 4r\log(2rt)+4$. Then the procedure of Definition 3.5 makes fewer than $2t$ additions and the final number of vertices is $|\text{clos}_r(U)| < 2rt$.*

Proof. Let U_i be the set after i steps, and let $G[U_i]$ denote the subgraph of G induced by U_i. We write $v_i = |U_i|$ and e_i for the number of vertices and edges of $G[U_i]$, respectively. A step adds a path of length $\ell \le r$, hence $\ell - 1 \le r - 1$ new vertices and at least ℓ new edges (each path edge touches a new vertex). So $v_i \le t + i(r-1)$ and $e_i - v_i \ge i - t$.

Suppose the process makes at least $2t$ steps. Then $v_{2t} < 2rt$ and $e_{2t} - v_{2t} \ge t$, so $G[U_{2t}]$ has average degree $d \ge 2 + 2t/v_{2t} > 2 + 1/r$. Being a subgraph of G, it also has girth at least g. Hence the average-degree Moore bound [AHL02] and $1 + x \ge \exp\left(\frac{x}{1+x}\right)$ with $x = 1/r$ give

$$v_{2t} \ge (d-1)^{\lfloor (g-1)/2 \rfloor} \ge \left(1 + \tfrac{1}{r}\right)^{2r\log(2rt)} \ge \exp\!\left(\tfrac{2r\log(2rt)}{r+1}\right) \ge 2rt,$$

where we used $\lfloor (g-1)/2 \rfloor \ge 2r\log(2rt)$ and $r + 1 \le 2r$. This contradicts $v_{2t} < 2rt$.

Hence fewer than $2t$ steps are made, and $|\mathrm{clos}_r(U)| \le t + (2t-1)(r-1) < 2rt$. □

Combining Lemma 3.6 with the size bound on the closure gives the version of the localization bound that we use below: a sampled vertex is far from the closure of U, not just from U itself. This requires the girth of E_n to be large enough for the size bound to apply, which holds for the parameters we use by Proposition 2.5.

Lemma 3.8. *Let $U \subseteq V$ be a set of vertices of G_n of size $t < 2^{\sqrt{n}}$. Then the probability that a vertex sampled from the ground-state distribution $p_n(v) = \langle v \,|\, \psi_n \rangle^2$ of $-A_n$ lies at distance at least $\ell_{n,K}$ from $\overline{U} = \mathrm{clos}_{2\ell_{n,K}+3}(\mathrm{ex}(U))$ is at least*

$$\Pr_{v \sim p_n}\left[\mathrm{dist}_E(\overline{U}, v) > \ell_{n,K}\right] \ge 1 - 2^{-n^2}.$$

Proof. Set $r = 2\ell_{n,K} + 3$ and $t = |U|$. By Proposition 2.5, the girth g_n of E_n satisfies $g_n > 4r\log(2rt) + 4$, so Lemma 3.7 applies and gives $|\overline{U}| < 2rt = 2(2\ell_{n,K}+3)t$. Applying Lemma 3.6 to the set $\overline{U}$, we obtain

$$\Pr_{v \sim p_n}\left[\mathrm{dist}_E(\overline{U}, v) \ge \ell_{n,K} + 1\right] \ge 1 - \frac{|\overline{U}|\, n^{\ell_{n,K}+1}}{N_E} \ge 1 - \frac{2(2\ell_{n,K}+3)t\, n^{\ell_{n,K}+1}}{N_E}.$$

Recall that $\ell_{n,K} = 10n^2 \log n$ and $N_E \ge 2^{100n^{5/2}\log^2 n}$. Since $t < 2^{\sqrt{n}}$, for all sufficiently large n,

$$\frac{2(2\ell_{n,K}+3)\, t\, n^{\ell_{n,K}+1}}{N_E} \le 2^{-n^2}.$$ □

The previous lemma analyzes the ideal case in which the output vertex is sampled according to the *ideal* ground-state distribution p_n. However, the GGSP problem allows outputting vertices from a perturbed distribution p_{out}, which may adversarially affect the occurrence of the above event. The next result shows that the event still occurs in this setting, at least with constant probability, when $|U| = t$ is of size at most $2^{\sqrt{n}}$ and ℓ is sufficiently smaller than the size of the expander.

Corollary 3.9. *Consider a joint distribution J over $t+1$ vertices $(U = \{u_1, u_2, \ldots, u_t\}, v)$ of G_n. Suppose that J is within total variation distance $1/5$ of a product distribution $p_{\mathrm{in}}^{\otimes t} \times p_{\mathrm{out}}$, where $F(p_{\mathrm{out}}, p_n) \ge (3/4)^2$. Then, for n large enough and $t < 2^{\sqrt{n}}$, the probability that the vertex v lies at distance at least $\ell_{n,K}$ from $\overline{U} = \mathrm{clos}_{2\ell_{n,K}+3}(\mathrm{ex}(U))$ satisfies,*

$$\Pr_{(u_1,\ldots,u_t,v)\sim J}\left[\mathrm{dist}_E(\overline{U}, v) > \ell_{n,K}\right] \ge 1/10.$$

Proof. The result is shown by upper-bounding the total variation distance between the distributions $p_{\mathrm{in}}^{\otimes t} \times p_n$ and J, and then applying the result from Lemma 3.8.

The total variation distance is bounded as follows,

$$\begin{aligned}
\mathrm{TV}(p_{\mathrm{in}}^{\otimes t} \times p_n, J) &\leq \mathrm{TV}(p_{\mathrm{in}}^{\otimes t} \times p_n, p_{\mathrm{in}}^{\otimes t} \times p_{\mathrm{out}}) + \mathrm{TV}(p_{\mathrm{in}}^{\otimes t} \times p_{\mathrm{out}}, J) \quad \text{(by the triangle inequality)} \\
&\leq \sqrt{1 - F(p_{\mathrm{out}}, p_n)} + \mathrm{TV}(p_{\mathrm{in}}^{\otimes t} \times p_{\mathrm{out}}, J) \\
&\qquad \text{(by relating the TV distance and the fidelity)} \\
&\leq \sqrt{1 - \frac{9}{16}} + \frac{1}{5} \leq 0.87. \qquad \text{(by the assumptions)}
\end{aligned}$$

Hence, with $r = 2\ell_{n,K} + 3$, the probability under J of sampling a vertex at distance at least $\ell_{n,K}$ from $\mathrm{clos}_r(\mathrm{ex}(U))$ is at least

$$\begin{aligned}
&\Pr_{(u_1,\ldots,u_t,v)\sim J}\left[\mathrm{dist}_E(\overline{U}, v) > \ell_{n,K}\right] \\
&\quad \geq \Pr_{(u_1,\ldots,u_t,v)\sim p_{\mathrm{in}}^{\otimes t}\times p_n}\left[\mathrm{dist}_E(\overline{U}, v) > \ell_{n,K}\right] - \mathrm{TV}(p_{\mathrm{in}}^{\otimes t} \times p_n, J) \\
&\qquad \text{(by definition of the TV distance)} \\
&\quad \geq \min_{W\subseteq V:|W|\leq t} \Pr_{v\sim p_n}\left[\mathrm{dist}_E(\mathrm{clos}_r(\mathrm{ex}(W)), v) > \ell_{n,K}\right] - \mathrm{TV}(p_{\mathrm{in}}^{\otimes t} \times p_n, J) \\
&\qquad \text{(by the independence of the distributions)} \\
&\quad \geq 1 - 2^{-n^2} - 0.87 \qquad \text{(by Lemma 3.8)} \\
&\quad \geq 1/10. \qquad \text{(for } n \text{ large enough)}
\end{aligned}$$

□

As a direct corollary, the same result holds for the input-output samples of any algorithm solving the GGSP problem, since the distributions J and J_π are the same object up to a relabeling of the vertices.

Proposition 3.10 (Input-Output state localization). *Fix any labeling π and suppose that a classical algorithm solves the GGSP problem on H_π under the conditions of Theorem 3.4. Let J_π denote the joint distribution of the input-output samples $(x_1, x_2, \ldots, x_t, y)$ of the algorithm under the vertex labeling π. Then, if $t < 2^{\sqrt{n}}$, the probability that the output vertex y lies at distance at least $\ell_{n,K}$ from $\overline{U} = \mathrm{clos}_{2\ell_{n,K}+3}(\mathrm{ex}(\pi^{-1}(\{x_1, x_2, \ldots, x_t\})))$, satisfies,*

$$\Pr_{(x_1,\ldots,x_t,y)\sim J_\pi}\left[\mathrm{dist}_E\left(\overline{U}, \pi^{-1}(y)\right) > \ell_{n,K}\right] \geq 1/10.$$

3.2 Proof of hardness of GGSP

This section proves a converse result to the previous one. Namely, it demonstrates that *no* classical algorithm can output a vertex far from the input vertices unless it makes at least sub-exponentially many queries to H_π. For simplicity, we consider throughout that the output is always a non-isolated vertex since the ground state is not supported on isolated vertices.

First, we show that any classical algorithm operating on the graph G_n is forced to explore the graph by querying connected components around its input vertices. This is similar to the argument in [CCD+03].

Lemma 3.11 (Local exploration). *Let $U \subseteq V$ be any subset of vertices of G_n with $|U| \leq |V|/2$. Suppose that a classical algorithm is given access to the adjacency-list oracle $\mathcal{O}_\pi$, for a random π, together with $\pi(U)$, and that it makes q queries. Then, except with probability at most $q \cdot 2^{-2n^3 \log^2 n}$, the set of all vertices queried by the algorithm at any point during its execution induces a collection of connected components rooted at the vertices in U, together with some isolated vertices in V_I.*

Proof. This is because the set of isolated vertices V_I in Definition 2.10 constitutes an overwhelming majority of the vertices of the graph G_n. Hence, if an algorithm attempts to query a vertex that is not connected to any of its previously queried vertices, then, due to the randomness of the labeling π, it has overwhelming probability of querying an isolated vertex.

In more detail, by Definition 2.10, the ratio between the number of non-isolated vertices and the total number of (non-input) vertices in G_n is at most $(|V_T| + |V_E|)/(|V| - |U|) \leq 2(|V_T| + |V_E|)/|V| = 2N_E(1 + \sqrt{n}\sum_{k=1}^{K-1} |\widehat{T}_{n,k}|)/2^{15n^3 \log^2 n} \leq 2^{-2n^3 \log^2 n}$. Hence, by a union bound over the q queries made by an algorithm, the probability that it queries a non-isolated vertex that is not connected to any of its previously explored vertices is at most $q \cdot 2^{-2n^3 \log^2 n}$. □

We now analyze how local graph exploration around an input set of vertices U in G_n behaves. The case $|U| = 1$ was addressed in Figures 1 and 2 and Lemma 2.15, where we showed that, within truncated neighborhood $B_{\ell_{n,K}}(u)$ of u, the expander edges outgoing from u form multiple copies of the tree $\widehat{T}_{n,K}$. The next lemma treats the case $|U| > 1$. The analysis is more involved, as different connected components may intersect and cycles may appear in the exploration graph. We show that such events must be confined to the closure of $\mathrm{ex}(U)$, a small superset $\overline{U} \supseteq \mathrm{ex}(U)$ of the input vertices, as defined in Definition 3.5. The remainder of the local exploration departing from these vertices can then be described by the same trees as before. This is illustrated in Figure 3.

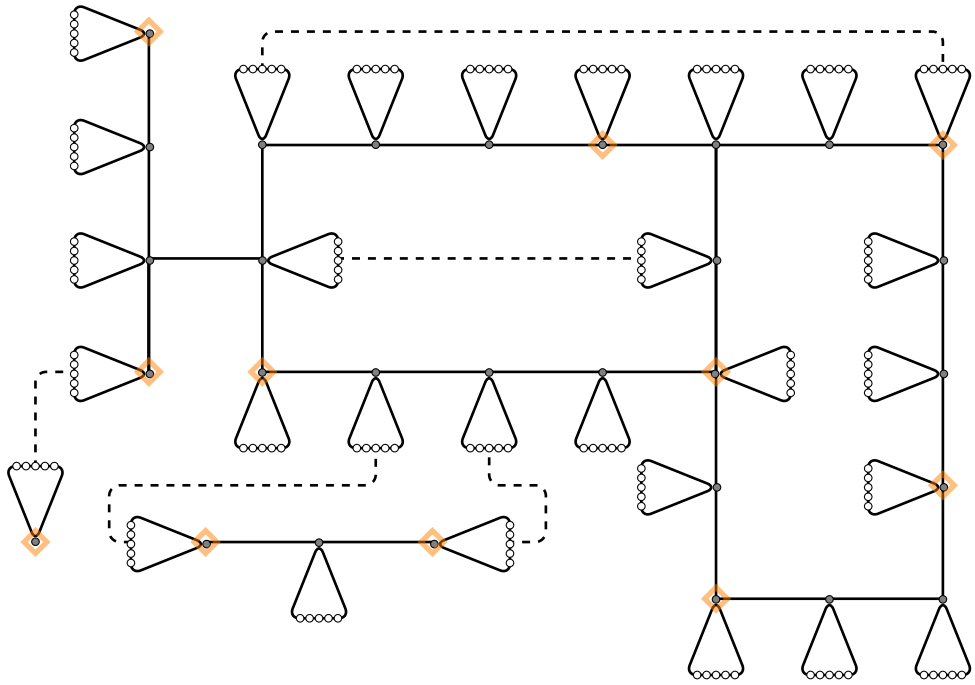


Figure 3: Pictorial representation of the local structure of the graph G_n (only expander vertices are shown). The orange diamonds represent the vertices in the set $\mathrm{ex}(U)$. The set $\overline{U} = \mathrm{clos}_{2\ell_{n,K}+3}(\mathrm{ex}(U))$, shown by the greyed vertices, is the fixed point of the following process: as long as two vertices of the current set are joined by a path of length at most $2\ell_{n,K} + 3$ whose internal vertices lie outside the set, those internal vertices are added to it. The solid edges are such connecting paths, including paths between vertices added in earlier rounds. The triangles depict the local tree structure of the truncated neighborhood $B_{\ell_{n,K}}(\overline{U})$ of $\overline{U}$. The white circles represent the vertices identified with the 0-level leaves of $\widehat{T}_{n,K}$. The dashed lines indicate other paths in the expander graph that may create cycles but lie too far from $\overline{U}$ to be detected by a local exploration.

Lemma 3.12 (Local structure of G_n around multiple roots)**.** *Let $U_0 \subseteq V$ be any subset of* expander *vertices of G_n. Let $\overline{U} = \mathrm{clos}_{2\ell_{n,K}+3}(U_0)$ denote the closure of U_0 with parameter $2\ell_{n,K}+3$. Consider the truncated neighborhood $B_{\ell_{n,K}}(\overline{U}) \subseteq G_n$ of $\overline{U}$ as defined in Definition 2.14. Then $B_{\ell_{n,K}}(\overline{U})$ is isomorphic to the graph obtained by taking the graph induced by $\overline{U}$ in G_n and then attaching to each $u \in \overline{U}$ as many copies of $\widehat{T}_{n,k}$ as in G_n (for $k = 1, \ldots, K-1$), and as many copies of $\widehat{T}_{n,K}$ as the number of expander neighbors of u outside $\overline{U}$.*

Moreover, the isomorphism can be chosen so that the expander vertices at distance exactly $\ell_{n,K}+1$ from $\overline{U}$ in G_n map to the 0-level leaves of the attached copies of $\widehat{T}_{n,K}$.

Proof. Let $R := \ell_{n,K}+1$. Fix $u \in \overline{U}$ and an expander neighbor $v \in V_E \setminus \overline{U}$ of u, and let $T(u,v)$ be the connected component of v in $B_{\ell_{n,K}}(\overline{U})$ *after deleting* the edge (u,v). We show that the subgraphs $\{T(u,v)\}_{u,v}$ are disjoint trees, each isomorphic to $\widehat{T}_{n,K}$. Take a breadth-first forest of the expander vertices of $B_{\ell_{n,K}}(\overline{U})$ rooted at $\overline{U}$; by Definition 2.14, its trees have depth at most R. Any expander edge of $B_{\ell_{n,K}}(\overline{U})$ not in this forest and not having both endpoints in $\overline{U}$ either joins two vertices of the same tree, closing a cycle of length at most $2R+1$, or joins two different trees, giving a path of length between 2 and $2R+1 = 2\ell_{n,K}+3$ between distinct vertices of $\overline{U}$ with all internal vertices outside $\overline{U}$. The first case contradicts the girth bound $g_n > 4\ell_{n,K}+8 > 2R+1$, and the second contradicts $\overline{U}$ being closed. Hence every such edge lies in the forest, and since decoration edges only attach trees, no cycles arise and the $T(u,v)$ are disjoint trees.

As in Lemma 2.15, the expander vertices of $T(u,v)$ form a full tree of depth $\ell_{n,K}$ and degree $d_{n,K}$, to which the decorations of Definition 2.10 are attached. This matches Definition 2.7 and Fact 2.8, so $T(u,v)$ is isomorphic to $\widehat{T}_{n,K}$. □

The next result studies the hardness of exploring the tree $\widehat{T}_{n,K}$, starting from its root and attempting to reach a 0-level leaf.

Lemma 3.13 (Hardness of exploring self-similar trees)**.** *Let $K = \sqrt{n}$, as in Definition 2.10. Suppose a classical algorithm explores the tree $\widehat{T}_{n,K}$, starting at its root and querying only vertices adjacent to those it has already visited. If the algorithm makes fewer than $2^{\sqrt{n}}$ queries, the probability that it ever reaches a 0-level leaf of $\widehat{T}_{n,K}$ is at most $2^{-2n\log n}$.*

Proof. The proof essentially follows [GHV21]. We provide a detailed version in Section C, adapted to our choice of parameters and to our definition of self-similar trees. The lemma is reformulated there as Proposition C.6. □

By combining the three lemmas above, we show that exploring the main graph G_n to reach expander vertices far from a given set U of input vertices is hard.

Proposition 3.14 (Hardness of exploring the graph G_n)**.** *Suppose that a classical algorithm is given access to H_π for a random labeling π, as well as input samples $x_1, \ldots, x_t$ from a distribution as specified in Theorem 3.4 with $t < 2^{\sqrt{n}}$. Suppose that it outputs an additional sample y, and let J_π denote the joint distribution of the input-output samples $(x_1, x_2, \ldots, x_t, y)$. If the algorithm makes fewer than $2^{\sqrt{n}}$ queries, then the probability that the output vertex y lies at distance at least $\ell_{n,K}$ from $\overline{U} = \mathrm{clos}_{2\ell_{n,K}+3}(\mathrm{ex}(\pi^{-1}(\{x_1, x_2, \ldots, x_t\})))$ satisfies,*

$$\Pr_{\pi,(x_1,\ldots,x_t,y)\sim J_\pi}\left[\mathrm{dist}_E\left(\overline{U}, \pi^{-1}(y)\right) > \ell_{n,K}\right] \leq 2^{-n\log n}.$$

Proof. Let $U = \pi^{-1}(\{x_1, \ldots, x_t\})$ denote the (random) set of input vertices sampled under J_π. By Lemma 3.11, except with probability $2^{\sqrt{n}} \cdot 2^{-2n^3\log^2 n} < 2^{-2n\log n}$, we may assume that the consecutive queries of the algorithm form a set of connected components in G_n rooted at the vertices in U.

Let $v = \pi^{-1}(y)$ be the vertex in G_n output by the algorithm. By Lemma 3.12, if the algorithm were to explore a vertex at distance at least $\ell_{n,K}+1$ from $\overline{U}$, then there must exist an exploration

tree attached to one of the vertices in $\overline{U}$ that can be embedded into $\widehat{T}_{n,K}$ and reaches one of its 0-level leaves. Fix one such copy T. Although the algorithm's queries to T may be interleaved adaptively with its queries elsewhere, we may simulate every query of the algorithm outside T using the rest of the graph together with the given input labels, which expose no internal vertex of T, and only charge to T the at most $2^{\sqrt{n}}$ queries that $\mathcal{A}$ makes inside it. Applying Lemma 3.13 to this restricted exploration bounds the probability that the algorithm reaches a 0-level leaf of T by $2^{-2n\log n}$, independently of what it learned from the other copies. By Lemma 3.7, the number of possible roots to consider for the embedding is at most $|\overline{U}| < 2n(2\ell_{n,K}+3)t$. A union bound then implies that the overall success probability is at most $2n(2\ell_{n,K}+3)t \cdot 2^{-2n\log n} \le 2^{-\frac{3}{2}n\log n}$, for n large enough. □

4 Hardness of Ground-State Energy Estimation Problem

In this section, we prove the classical hardness of GGSE by studying two families of stoquastic Hamiltonians obtained from two closely related graphs in Definition 2.10. We fix $K = \sqrt{n}$ in this section. The first family corresponds to the main graph $G_n = G_{n,E_n}$ (studied already in Section 3), which gives the YES instances. The second family corresponds to the *envelope graph* $F_n = G_{n,E'_n}$, and gives the NO instances. Solving GGSE on these instances amounts to distinguishing the two graph families by the next lemma.

Lemma 4.1 (GGSE INSTANCES)**.** *There exist thresholds $a < b$ with $b - a = \Omega(\sqrt{n})$ such that, for every labeling π, $H_{G_n,\pi}$ satisfies the* YES *promise of GGSE (i.e., $\lambda_{\min}(H_{G_n,\pi}) \le a$) and $H_{F_n,\pi}$ satisfies the* NO *promise (i.e., $\lambda_{\min}(H_{F_n,\pi}) \ge b$).*

Proof. Let $\lambda_{\max}(G_n)$ and $\lambda_{\max}(F_n)$ denote the top eigenvalues of G_n and F_n respectively. Then, by Section B, there exists an absolute constant $C > 0$ such that $\lambda_{\max}(G_n) \ge n$ and $\lambda_{\max}(F_n) \le n - \sqrt{n} + C$. Set $a := -n$, $b := -n + \sqrt{n} - C$. Then $\lambda_{\min}(H_{G_n,\pi}) \le a$, $\lambda_{\min}(H_{F_n,\pi}) \ge b$, and $b - a = \sqrt{n} - C = \Omega(\sqrt{n})$. □

As in the GGSP setting, the corresponding quantum problem is easy.

Proposition 4.2 (QUANTUM ALGORITHM)**.** *There exists a quantum algorithm that solves GGSE on every Hamiltonian $H_{G,\pi}$, for $G \in \{G_n, F_n\}$, using* $\mathrm{poly}(n)$ *copies of a guiding state $|\psi_{\mathrm{in},\pi}\rangle$, and* $\mathrm{poly}(n)$ *quantum queries to the adjacency-list oracle of $H_{G,\pi}$.*

Proof. Both G_n and F_n, and hence their associated Hamiltonians, are $O(n)$-sparse. The result follows from Proposition 2.4, since a $O(n)$-sparse $\mathrm{poly}(n)$-qubit Hamiltonian can be simulated using $\mathrm{poly}(n)$ quantum gates and queries to its adjacency-list oracle. □

The main and envelope graphs differ in the degree of their base expanders and in the decorations attached to each expander vertex. These changes compensate for one another, so every internal vertex has degree $d_{n,1} = 2n - \sqrt{n}$ in both graphs. The constructions also have the same total number of vertices and the same girth lower bound. Thus these elementary graph parameters do not distinguish the two families. The main result of this section is to show that, in fact, no classical algorithm using only sub-exponentially many queries and input samples can distinguish between the two graphs (and hence solve GGSE).

Theorem 4.3 (CLASSICAL HARDNESS OF GGSE)**.** *Suppose that a classical algorithm is making q queries to the adjacency-list oracle of $H_{G,\pi}$, for a random $G \in \{G_n, F_n\}$ and a random labeling π, and is given t independent samples $x_1, \ldots, x_t \sim p_{E,\pi}$ from the distribution $p_{E,\pi}(x) = \langle x \,|\, \psi_{E,\pi}\rangle^2$, where $|\psi_{E,\pi}\rangle$ is the uniform superposition over the expander vertices of G. If $q + t < 2^{\sqrt{n}}$ then the algorithm cannot solve the GGSE problem on G with success probability greater than $1/2 + 2^{-n}$. The guiding state has overlap $|\langle \psi_{E,\pi} \,|\, \psi_{G,\pi}\rangle| = 1 - O(1/n)$ with the ground state $|\psi_{G,\pi}\rangle$ of $H_{G,\pi}$.*

4.1 Local graph structures

This section studies the local structures around expander vertices in G_n and F_n. Although these structures differ, we show that the local structure in G_n embeds into that in F_n up to a natural frontier defined by a subset of the 0-level leaves of $\widehat{T}_{n,K}$.

By Lemma 2.15 and Figure 2, the truncated neighborhood of an expander vertex u differs between the two graphs in that the expander-generated branches are connected to copies of $\widehat{T}_{n,K}$ in G_n, but to copies of $\widehat{T}_{n,K+1}$ in F_n. We relate the two local structures by first showing that $\widehat{T}_{n,K}$ can itself be embedded into a copy of $\widehat{T}_{n,K+1}$. The fact that $\widehat{T}_{n,K}$ is a subgraph of $\widehat{T}_{n,K+1}$ is straightforward, since the latter graph contains multiple copies of the former. The much subtler property, which we demonstrate next, is that such an embedding is possible even when the roots of the two graphs are matched together. This is illustrated in Figure 4.

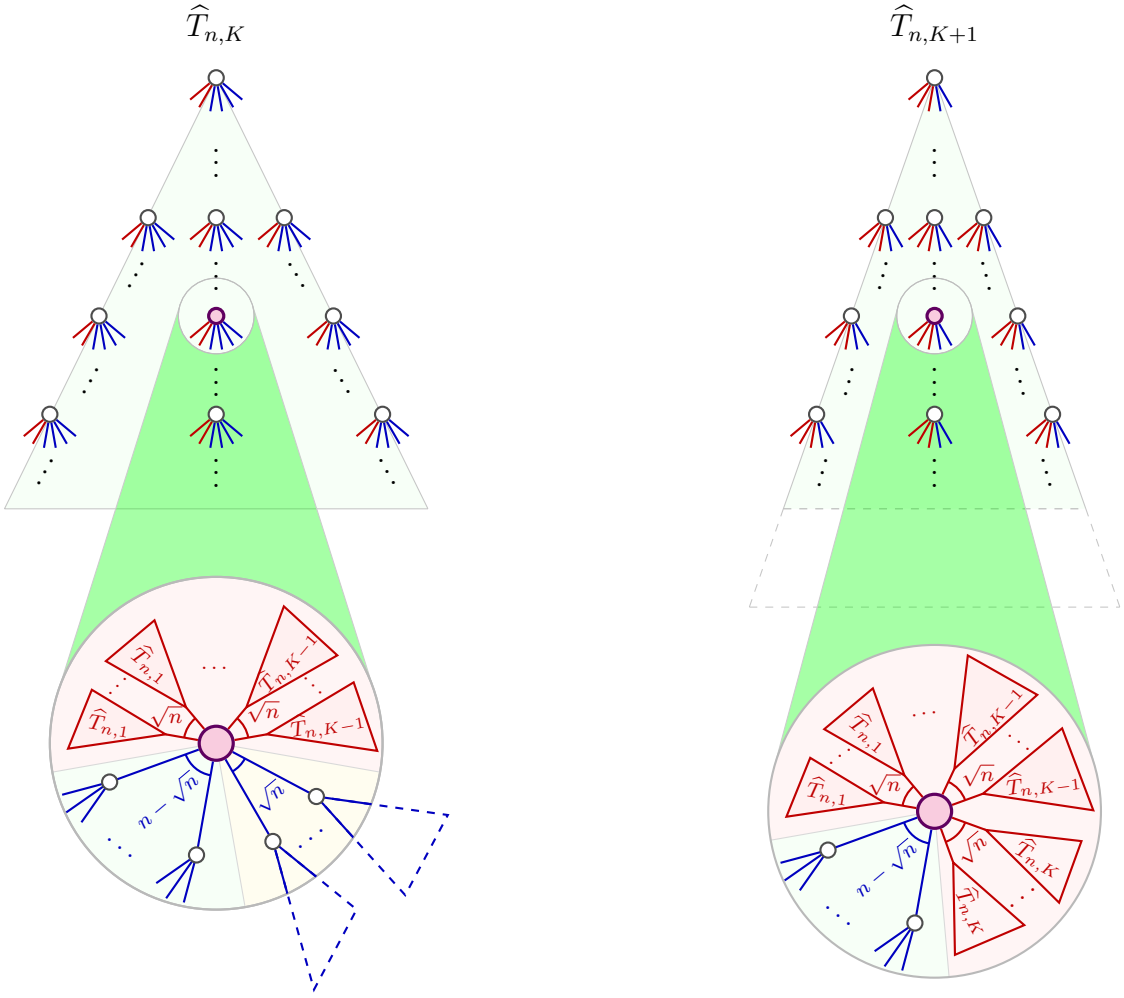


Figure 4: Schematic illustration of the rooted embedding of $\widehat{T}_{n,K}$ into the truncated copy of $\widehat{T}_{n,K+1}$. The zoomed views show one recursive step at corresponding 0-level internal vertices. The attached copies of $\widehat{T}_{n,k}$, for $k = 1, \ldots, K-1$, are mapped by canonical rooted isomorphisms (red regions). Of the remaining branches, $n - \sqrt{n} - 1$ continue recursively through 0-level children (green regions). For each of the other $\sqrt{n}$ children, shown in the yellow region on the left, the entire subtree rooted at that child is mapped onto one of the copies of $\widehat{T}_{n,K}$ attached as a decoration in the red region on the right.

Lemma 4.4 (Embedding of self-similar trees)**.** *For each $1 \le k \le \sqrt{n}$, there exists an*

embedding of $\widehat{T}_{n,k}$ into $\widehat{T}_{n,k+1}$ that maps the roots to each other and preserves the degrees of the nodes, except at some of the 0-level leaves of $\widehat{T}_{n,k}$.

Proof. By Fact 2.8, $\widehat{T}_{n,k}$ and $\widehat{T}_{n,k+1}$ are obtained by decorating the perfect trees $C_{n,k}$ and $C_{n,k+1}$ respectively (blue edges in Figure 4). The tree $C_{n,k+1}$ is deeper but has smaller degree $d_{n,k+1} = d_{n,k} - \sqrt{n}$. On the other hand, each node carries $\sqrt{n}$ additional decorations in $\widehat{T}_{n,k+1}$, each a copy of $\widehat{T}_{n,k}$.

We construct the embedding in two steps. First, consider a perfect subtree of $C_{n,k}$ of degree $d_{n,k+1}$ and depth $\ell_{n,k}$, together with all decorations attached to it. The perfect subtree embeds into the first $\ell_{n,k}$ levels of $C_{n,k+1}$ (blue edges in green region in Figure 4), and its decorations embed into those of $C_{n,k+1}$ since they form a subset of the latter (red region in Figure 4).

Next, it remains to embed, for each node of $C_{n,k}$, the $d_{n,k} - d_{n,k+1} = \sqrt{n}$ subtrees rooted at 0-level children not yet accounted for (yellow region in Figure 4). Consider one such child, at some depth $1 \le \ell \le \ell_{n,k}$. The subtree rooted at it is isomorphic to $C_{n,k}$ truncated at depth $\ell_{n,k} - \ell$ and decorated as in $\widehat{T}_{n,k}$. Each such subtree therefore embeds into one of the $\sqrt{n}$ additional copies of $\widehat{T}_{n,k}$ attached at the corresponding node of $C_{n,k+1}$.

Because of the depth truncated at level ℓ, the 0-level leaves of the subtree are mapped to internal nodes of the copy of $\widehat{T}_{n,k}$. These are the only vertices whose degree is not preserved by the embedding. □

The above result extends readily to an embedding of the local structure of G_n into that of F_n, as we demonstrate next and illustrate in Figure 5.

Corollary 4.5 (Embedding of local structures)**.** *Let $u \in V_E(G_n)$ and $v \in V_E(F_n)$ be expander vertices. There exists an embedding of truncated neighborhoods, $B_{G_n,\ell_{n,K}}(u)$ into $B_{F_n,\ell_{n,K}}(v)$ that maps u to v and preserves the degrees of the nodes, except at some of the expander vertices at distance $\ell_{n,K}+1$ from u.*

Proof. By Lemma 2.15 and Figure 2, the two graphs differ in the $n - \sqrt{n}$ copies of $\widehat{T}_{n,K+1}$ attached in $B_{F_n,\ell_{n,K}}(v)$ in place of $\widehat{T}_{n,K}$ in $B_{G_n,\ell_{n,K}}(u)$. By Lemma 4.4, there exists a rooted embedding of the latter into the former, under which the degree of every vertex is preserved except at some 0-level leaves of the copies of $\widehat{T}_{n,K}$. By Lemma 2.15, these leaves correspond in $B_{G_n,\ell_{n,K}}(u)$ to expander vertices at distance $\ell_{n,K}+1$ from u. □

4.2 Proof of hardness of GGSE

We use the embedding established in the previous section to show the hardness of distinguishing G_n from F_n, even when given multiple expander vertices that are far apart.

Proposition 4.6. *Fix an integer t. For each $G \in \{G_n, F_n\}$, let $U_G \subset V_E(G)$ be a set of expander vertices of size $|U_G| = t$ such that $\mathrm{dist}_E(u,v) > 2(\ell_{n,K}+1)$ for all $u, v \in U_G$, where $K = \sqrt{n}$.*

Suppose that a classical algorithm is given access to the adjacency-list oracle $\mathcal{O}_{G,\pi}$, for a random $G \in \{G_n, F_n\}$ and a random π, together with $\pi(U_G)$, and that it makes q queries. If $q + t < 2^{\sqrt{n}}$ then the algorithm cannot correctly identify G with probability greater than $1/2 + 2^{-3n/2}$.

Proof. By Lemma 3.11, except with probability $q \cdot 2^{-n^3 \log^2 n}$, we may assume that the vertices queried by the algorithm induce a collection of connected components rooted at the vertices in U_G, together with some isolated vertices.

Conditioned on this event, since $\mathrm{dist}_E(u,v) > 2(\ell_{n,K}+1)$ for all $u, v \in U_G$, the local exploration structures around the vertices in U_G remain disjoint at least until the algorithm queries an expander vertex at distance $\ell_{n,K}+1$ from some $u \in U_G$. Moreover, by Corollary 4.5, the local structure around u is indistinguishable between $G = G_n$ and $G = F_n$ at least until

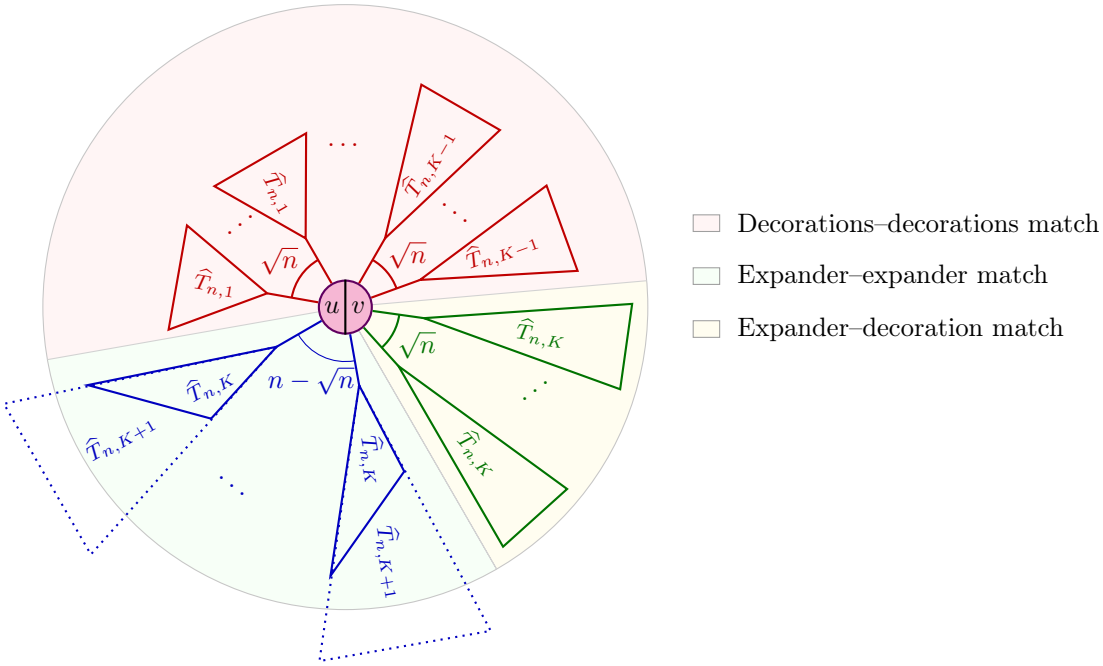


Figure 5: Schematic of the rooted embedding of $B_{G_n,\ell_{n,K}}(u)$ into $B_{F_n,\ell_{n,K}}(v)$. The decorations $\widehat{T}_{n,k}$, $k = 1, \ldots, K-1$, attached to u are mapped canonically to the corresponding decorations at v (red region). Of the n expander-generated copies of $\widehat{T}_{n,K}$ at u, $n - \sqrt{n}$ are mapped into the truncated copies of $\widehat{T}_{n,K+1}$ at v using Lemma 4.4 (green region); the solid blue inner triangles indicate the embedded copies of $\widehat{T}_{n,K}$ inside these larger dotted trees. The remaining $\sqrt{n}$ expander-generated copies are mapped canonically onto the $\sqrt{n}$ copies of $\widehat{T}_{n,K}$ attached to v as decorations (yellow region). Thus, red, green, and yellow shading indicate decoration–decoration, expander–expander, and expander–decoration matches, respectively.

such a vertex is queried. By Lemma 2.15, the probability of this event is upper-bounded by the probability of finding a 0-level leaf in the tree $\widehat{T}_{n,K}$, which is at most $2^{-2n\log n}$ by Lemma 3.13. By a union bound over all $u \in U_G$, the probability of querying any vertex at distance $\ell_{n,K}+1$ from U_G is therefore at most $t \cdot 2^{-2n\log n}$.

By standard arguments relating the probability of distinguishing two graphs to the probability that their query responses first differ (see, e.g., [Has21, Lemma 5]), we conclude that the probability of correctly identifying G is at most $q \cdot 2^{-n^3 \log^2 n} + (1/2 + t \cdot 2^{-2n\log n}/2) < 1/2 + 2^{-3n/2}$ since $q + t < 2^{\sqrt{n}}$. □

We are now ready to prove that the two families of GGSE instances G_n and F_n cannot be distinguished efficiently by a classical algorithm.

Proof of Theorem 4.3. We first show that the input samples satisfy the spread out condition needed to apply Proposition 4.6 with high probability. Let $u_i = \pi^{-1}(x_i)$ for all i be the vertices sampled under the guiding state $|\psi_{E,\pi}\rangle$ (i.e., the uniform superposition over the expander vertices). By the same arguments as in Lemma 3.6, the probability that two vertices are close to each other is, $\Pr[\mathrm{dist}_E(u_i, u_j) \le 2(\ell_{n,K}+1)$ for some $i \neq j] \le \binom{t}{2} \cdot \frac{n^{2(\ell_{n,K}+1)+1}}{N_E} \le 2^{-n^2/2}$ since $t < 2^{\sqrt{n}}$, $\ell_{n,K} = 10n^2 \log n$ and $N_E \ge 2^{100n^2\log^2 n}$. Thus we can assume that with probability at least $1 - 2^{-n^2/2}$, the set $U = \{u_1, \ldots, u_t\}$ of input samples satisfies the spread out condition needed to apply Proposition 4.6. Conditioned on this event, the classical algorithm cannot distinguish the two graphs with probability greater than $1/2 + 2^{-3n/2}$. Thus, the overall distinguishing

probability is at most $1/2 + 2^{-3n/2} + 2^{-n^2/2} < 1/2 + 2^{-n}$.

The proof that the guiding state has overlap $|\langle \psi_{E,\pi} \,|\, \psi_{G,\pi} \rangle| = 1 - O(1/n)$ follows directly from Lemma 2.12. □

Acknowledgements

We acknowledge the use of AI-based tools for language editing and proofreading, and for auditing the proofs after the first version of this paper was posted on arXiv. This audit helped us identify and correct minor errors.

This work was supported by the Maison du Quantique de Nouvelle-Aquitaine "HybQuant", as part of the HQI initiative and France 2030, under the French National Research Agency (ANR) grant ANR-22-PNCQ-0002. It was also supported by the PEPR integrated project EPiQ under ANR grant ANR-22-PETQ-0007.

References

[AF02] D. Aldous and J. A. Fill. *Reversible Markov Chains and Random Walks on Graphs*. Unfinished monograph, recompiled 2014, available at http://www.stat.berkeley.edu/~aldous/RWG/book.html. 2002 (cit. on p. 1).

[AHL02] N. Alon, S. Hoory, and N. Linial. "The Moore Bound for Irregular Graphs". In: *Graphs and Combinatorics* 18.1 (2002), pp. 53–57 (cit. on p. 16).

[AK07] S. Aaronson and G. Kuperberg. "Quantum Versus Classical Proofs and Advice". In: *Theory of Computing* 3.7 (2007), pp. 129–157 (cit. on p. 3).

[AL18] T. Albash and D. A. Lidar. "Adiabatic Quantum Computation". In: *Reviews of Modern Physics* 90 (2018), p. 015002 (cit. on p. 1).

[AT07] D. Aharonov and A. Ta-Shma. "Adiabatic Quantum State Generation". In: *SIAM Journal on Computing* 37.1 (2007), pp. 47–82 (cit. on pp. 3, 13).

[BBT06] S. Bravyi, A. J. Bessen, and B. M. Terhal. *Merlin-Arthur Games and Stoquastic Complexity*. arXiv:quant-ph/0611021. 2006 (cit. on p. 2).

[BCG+24] S. Ben-David, A. M. Childs, A. Gilyén, W. Kretschmer, S. Podder, and D. Wang. "Symmetries, Graph Properties, and Quantum Speedups". In: *SIAM Journal on Computing* 53.6 (2024), FOCS20-368–FOCS20-415 (cit. on p. 3).

[BCGL23] S. Bravyi, G. Carleo, D. Gosset, and Y. Liu. "A Rapidly Mixing Markov Chain from Any Gapped Quantum Many-Body System". In: *Quantum* 7 (2023), p. 1173 (cit. on p. 2).

[BCK15] B. Barak, S. O. Chan, and P. K. Kothari. "Sum of Squares Lower Bounds from Pairwise Independence". In: *Proceedings of the 47th Symposium on Theory of Computing (STOC)*. 2015, pp. 97–106 (cit. on pp. 14, 15).

[BDOT08] S. Bravyi, D. P. DiVincenzo, R. Oliveira, and B. M. Terhal. "The Complexity of Stoquastic Local Hamiltonian Problems". In: *Quantum Information & Computation* 8.5 (2008), pp. 361–385 (cit. on pp. 2, 3).

[BFM23] R. Bassirian, B. Fefferman, and K. Marwaha. "On the Power of Nonstandard Quantum Oracles". In: *Proceedings of the 18th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC)*. 2023, 11:1–11:25 (cit. on p. 3).

[BHNZ26] J. Bostanci, J. Haferkamp, C. Nirkhe, and M. Zhandry. "Separating QMA from QCMA with a Classical Oracle". In: *Proceedings of the 58th Symposium on Theory of Computing (STOC)*. 2026, pp. 595–606 (cit. on p. 3).

[BHV26] J. Bostanci, A. Huang, and V. Vaikuntanathan. *Separating Quantum and Classical Advice with Good Codes*. arXiv:2602.09385 [quant-ph]. 2026 (cit. on p. 3).

[BLH25] S. Balasubramanian, T. Li, and A. W. Harrow. "Exponential Speedups for Quantum Walks in Random Hierarchical Graphs". In: *Communications in Mathematical Physics* 406.9 (2025), p. 209 (cit. on p. 3).

[Bra15] S. Bravyi. "Monte Carlo Simulation of Stoquastic Hamiltonians". In: *Quantum Information & Computation* 15.13-14 (2015), pp. 1122–1140 (cit. on p. 2).

[BT10] S. Bravyi and B. Terhal. "Complexity of Stoquastic Frustration-Free Hamiltonians". In: *SIAM Journal on Computing* 39.4 (2010), pp. 1462–1485 (cit. on pp. 2, 3).

[CCD+03] A. M. Childs, R. Cleve, E. Deotto, E. Farhi, S. Gutmann, and D. A. Spielman. "Exponential Algorithmic Speedup by a Quantum Walk". In: *Proceedings of the 35th Symposium on Theory of Computing (STOC)*. 2003, pp. 59–68 (cit. on pp. 3, 10, 17).

[CFG+23] C. Cade, M. Folkertsma, S. Gharibian, R. Hayakawa, F. Le Gall, T. Morimae, and J. Weggemans. "Improved Hardness Results for the Guided Local Hamiltonian Problem". In: *Proceedings of the 50th International Colloquium on Automata, Languages, and Programming (ICALP)*. 2023, 32:1–32:19 (cit. on p. 2).

[CH12] J. Cullinan and F. Hajir. "Primes of Prescribed Congruence Class in Short Intervals". In: *Integers* 12 (2012), A56 (cit. on p. 27).

[CH23] A. Cornelissen and Y. Hamoudi. "A Sublinear-Time Quantum Algorithm for Approximating Partition Functions". In: *Proceedings of the 34th Symposium on Discrete Algorithms (SODA)*. 2023, pp. 1245–1264 (cit. on p. 8).

[CMM09] M. Charikar, K. Makarychev, and Y. Makarychev. "Integrality Gaps for Sherali-Adams Relaxations". In: *Proceedings of the 41st Symposium on Theory of Computing (STOC)*. 2009, pp. 283–292 (cit. on p. 14).

[Dah14] X. Dahan. "Regular Graphs of Large Girth and Arbitrary Degree". In: *Combinatorica* 34.4 (2014), pp. 407–426 (cit. on p. 9).

[EJ08] G. Exoo and R. Jajcay. "Dynamic Cage Survey". In: *The Electronic Journal of Combinatorics* Dynamic Survey 16 (2008) (cit. on p. 9).

[ES63] P. Erdős and H. Sachs. "Reguläre Graphen Gegebener Taillenweite mit Minimaler Knotenzahl". In: *Wiss. Z. Martin-Luther-Univ. Halle-Wittenberg Math.-Natur. Reihe* 12 (1963), pp. 251–257 (cit. on p. 9).

[Gal25] F. Le Gall. "Classical Algorithms for Constant Approximation of the Ground State Energy of Local Hamiltonians". In: *Proceedings of the 33rd European Symposium on Algorithms (ESA 2025)*. 2025, 73:1–73:19 (cit. on pp. 2, 4).

[GG22] S. Gharibian and F. Le Gall. "Dequantizing the Quantum Singular Value Transformation: Hardness and Applications to Quantum Chemistry and the Quantum PCP Conjecture". In: *Proceedings of the 54th Symposium on Theory of Computing (STOC)*. 2022, pp. 19–32 (cit. on pp. 2, 4).

[GHLS14] S. Gharibian, Y. Huang, Z. Landau, and S. W. Shin. "Quantum Hamiltonian Complexity". In: *Foundations and Trends in Theoretical Computer Science* 10.3 (2014), pp. 159–282 (cit. on p. 2).

[GHV21] A. Gilyén, M. B. Hastings, and U. Vazirani. "(Sub)Exponential Advantage of Adiabatic Quantum Computation with No Sign Problem". In: *Proceedings of the 53rd Symposium on Theory of Computing (STOC)*. 2021, pp. 1357–1369 (cit. on pp. 1, 3–5, 9–11, 19, 31, 34, 35).

[GLMR25] F. Le Gall, M. Luce, J. Marchand, and M. Roget. *Exponential Quantum Advantage for Message Complexity in Distributed Algorithms*. arXiv:2510.01657 [quant-ph]. 2025 (cit. on p. 3).

[GR26] A. B. Grilo and M. Rozos. *The Complexity of Stoquastic Sparse Hamiltonians*. arXiv:2605.02845 [cs.CC]. 2026 (cit. on p. 2).

[Has21] M. B. Hastings. "The Power of Adiabatic Quantum Computation with No Sign Problem". In: *Quantum* 5 (2021), p. 597 (cit. on pp. 1, 3–5, 10, 11, 23, 31).

[HW20] A. W. Harrow and A. Y. Wei. "Adaptive Quantum Simulated Annealing for Bayesian Inference and Estimating Partition Functions". In: *Proceedings of the 31st Symposium on Discrete Algorithms (SODA)*. 2020, pp. 193–212 (cit. on p. 8).

[JGL10] S. P. Jordan, D. Gosset, and P. J. Love. "Quantum-Merlin-Arthur–Complete Problems for Stoquastic Hamiltonians and Markov Matrices". In: *Physical Review A* 81 (2010), p. 032331 (cit. on p. 3).

[Jia25] J. Jiang. "Local Hamiltonian Problem with Succinct Ground State is MA-Complete". In: *PRX Quantum* 6 (2025), p. 020312 (cit. on p. 2).

[JSV04] M. Jerrum, A. Sinclair, and E. Vigoda. "A Polynomial-Time Approximation Algorithm for the Permanent of a Matrix with Nonnegative Entries". In: *Journal of the ACM* 51.4 (2004), pp. 671–697 (cit. on p. 8).

[JZ23] S. Jeffery and S. Zur. "Multidimensional Quantum Walks". In: *Proceedings of the 55th Symposium on Theory of Computing (STOC)*. 2023, pp. 1125–1130 (cit. on p. 3).

[KSV02] A. Y. Kitaev, A. H. Shen, and M. N. Vyalyi. *Classical and Quantum Computation*. Vol. 47. American Mathematical Society, 2002 (cit. on pp. 1, 2).

[Liu21] Y. Liu. "StoqMA Meets Distribution Testing". In: *Proceedings of the 16th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC)*. 2021, 4:1–4:22 (cit. on p. 2).

[LLL24] G. Li, L. Li, and J. Luo. "Recovering the Original Simplicity: Succinct and Exact Quantum Algorithm for the Welded Tree Problem". In: *Algorithmica* 86.12 (2024), pp. 3719–3758 (cit. on p. 3).

[LMY25] J. Liu, S. Mutreja, and H. Yuen. "QMA vs QCMA and Pseudorandomness". In: *Proceedings of the 57th Symposium on Theory of Computing (STOC)*. 2025, pp. 1520–1531 (cit. on p. 3).

[LPS88] A. Lubotzky, R. Phillips, and P. Sarnak. "Ramanujan Graphs". In: *Combinatorica* 8.3 (1988), pp. 261–277 (cit. on pp. 9, 27, 28).

[Lut11] A. Lutomirski. *Component Mixers and a Hardness Result for Counterfeiting Quantum Money*. arXiv:1107.0321 [quant-ph]. 2011 (cit. on p. 3).

[LWWZ25] J. Leng, K. Wu, X. Wu, and Y. Zheng. *(Sub)Exponential Quantum Speedup for Optimization*. arXiv:2504.14841 [quant-ph]. 2025 (cit. on p. 3).

[LZ25] J. Li and S. Zur. "Multidimensional Electrical Networks and their Application to Exponential Speedups for Graph Problems". In: *Quantum* 9 (2025), p. 1733 (cit. on p. 3).

[MWW04] B. D. McKay, N. C. Wormald, and B. Wysocka. "Short Cycles in Random Regular Graphs". In: *The Electronic Journal of Combinatorics* 11.1 (2004) (cit. on p. 9).

[NN24] A. Natarajan and C. Nirkhe. "A Distribution Testing Oracle Separation between QMA and QCMA". In: *Quantum* 8 (2024), p. 1377 (cit. on p. 3).

[Par21] P. Paredes. "Spectrum Preserving Short Cycle Removal on Regular Graphs". In: *Proceedings of the 38th Symposium on Theoretical Aspects of Computer Science (STACS)*. 2021, 55:1–55:19 (cit. on pp. 9, 28, 29).

[SNK12] R. D. Somma, D. Nagaj, and M. Kieferová. "Quantum Speedup by Quantum Annealing". In: *Physical Review Letters* 109 (2012), p. 050501 (cit. on p. 3).

[ŠVV09] D. Štefankovič, S. S. Vempala, and E. Vigoda. "Adaptive Simulated Annealing: A Near-Optimal Connection between Sampling and Counting". In: *Journal of the ACM* 56.3 (2009), pp. 1–36 (cit. on p. 8).

[TCC+22] J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth, and J. Tennyson. "The Variational Quantum Eigensolver: A Review of Methods and Best Practices". In: *Physics Reports* 986 (2022), pp. 1–128 (cit. on p. 1).

[Wai25] G. Waite. *The Guided Local Hamiltonian Problem for Stoquastic Hamiltonians*. arXiv:2509.25829 [quant-ph]. 2025 (cit. on p. 2).

[WFC24] J. Weggemans, M. Folkertsma, and C. Cade. "Guidable Local Hamiltonian Problems with Implications to Heuristic Ansatz State Preparation and the Quantum PCP Conjecture". In: *Proceedings of the 19th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC)*. 2024, 10:1–10:24 (cit. on p. 2).

A Construction of high-degree, high-girth expanders

A.1 Construction via LPS Ramanujan graphs

Lubotzky, Phillips and Sarnak [LPS88] constructed a family of high-degree, high-girth Ramanujan Cayley graphs $(X^{p,q})_{p,q}$, parametrized by pairs of prime numbers p and q congruent to 1 modulo 4. We show that these prime numbers can be chosen so that deleting suitable pairs of generators yields the graphs required by Proposition 2.5. First, we recall the result from [LPS88].

Theorem A.1 ([LPS88])**.** *Let $p < q$ be two prime numbers congruent to 1 modulo 4. Then, the Ramanujan graph $X^{p,q}$ constructed in [LPS88] has the following properties:*

- *The number of vertices is either $q(q^2-1)/2$ or $q(q^2-1)$,*
- *The graph is regular of degree $p+1$,*
- *The spectral gap between its two largest adjacency eigenvalues is at least $(p+1) - 2\sqrt{p}$,*
- *The girth is at least $2\log_p q$.*

We now prove that Proposition 2.5 holds for every sufficiently large even n.

Proposition A.2. *For every sufficiently large even integer n, there exists a graph E_n on $N_E \in [2^{100n^{5/2}\log^2 n}, 2^{n^3\log^2 n}]$ vertices with degree $d_n = n$, spectral gap at least $n/2$ and girth at least $g_n = 100n^{5/2}\log n$.*

Proof. The proof makes use of the following variant of Bertrand's postulate for primes of prescribed congruence classes.

Lemma A.3 ([CH12, Theorem 1])**.** *For any $x \geq 887$, the interval $(x, 1.048x]$ contains a prime congruent to 1 mod 4.*

Let n be a sufficiently large even integer. By Lemma A.3, choose a prime $p \equiv 1 \pmod 4$ in $(n, 1.048n]$. Set $B = 60n^{5/2}\log^2 n$ and choose a prime $q \equiv 1 \pmod 4$ in $(2^B, 1.048 \cdot 2^B] \subset (2^B, 2^{B+1}]$ using the same lemma.

The generating set of the Cayley graph $X^{p,q}$ splits into $(p+1)/2$ pairs $\{s, s^{-1}\}$, each consisting of a generator and its inverse [LPS88, Section 3]. Removing such a pair preserves the undirectedness of the graph and decreases its degree by two. Retain any $n/2$ of these pairs and let E_n be the corresponding Cayley graph on the same vertex set. The deleted edges form a d'-regular graph Y, where $d' = p + 1 - n \le 0.048n + 1$, so $A_{E_n} = A_{X^{p,q}} - A_Y$ and $\|A_Y\| \le d'$. Applying the variational principle on the subspace orthogonal to the uniform vector, its second-largest adjacency eigenvalue satisfies $\lambda_2(E_n) \le 2\sqrt{p} + d' \le 2\sqrt{1.048n} + 0.048n + 1 \le n/2$ for sufficiently large n. Thus the spectral gap of E_n is at least $n/2$.

By Theorem A.1, the number of vertices satisfies $N_E \in \{q(q^2-1)/2, q(q^2-1)\}$, and hence $\log N_E = 3B + O(1) = 180n^{5/2}\log^2 n + O(1)$. Finally, deleting edges cannot decrease the girth, so it is at least $2\log_p q \ge 120n^{5/2}\log^2(n)/\log(1.048n) \ge 100n^{5/2}\log n$. □

A.2 Construction via short cycle removal

We use the following results of Paredes [Par21]. We first recall the configuration model and the relevant definitions. Throughout, loops contribute two to the degree of a vertex.

Definition A.4 (Configuration model, [Par21, Definition 14])**.** Let $N > d > 0$ be integers with Nd even. The d-regular N-vertex configuration model is obtained by taking the set of half-edges $[N] \times [d]$ and choosing a uniformly random perfect matching of these Nd half-edges. Each matched pair is then joined to form an edge between the corresponding vertices. The resulting object is a d-regular undirected multigraph, where loops and parallel edges are allowed.

For a d-regular multigraph G, let $d = \lambda_1(G) \ge \lambda_2(G) \ge \cdots \ge \lambda_N(G)$ denote the eigenvalues of its adjacency matrix, and write $\lambda(G) := \max\{\lambda_2(G), |\lambda_N(G)|\}$.

Definition A.5 ([Par21, Definitions 6 and 7])**.** A multigraph G is *bicycle-free at radius* r if the radius-r neighborhood of every vertex contains at most one cycle.

For positive integers r and τ, we call a (multi)graph G an (r, τ)-(multi)graph if it is bicycle-free at radius at least r and contains at most τ cycles of length at most r. We denote the collection of such cycles by $\mathrm{Cyc}_r(G)$.

The following result gives bicycle-freeness and spectral expansion simultaneously, while allowing the degree to grow with the number of vertices.

Theorem A.6 ([Par21, Theorem 32])**.** *There exists a sufficiently large universal constant α such that the following holds. Let $N > 0$, and suppose that $3 \le d \le \alpha^{-1}\sqrt{\log N}$, $c < 1/4$, and*

$$\alpha^3 \left(\frac{\log\log N}{\log_{d-1} N}\right)^2 \le \epsilon \le 1.$$

If G is sampled from the d-regular N-vertex uniform configuration model, then, except with probability at most $N^{-0.99}$, G is bicycle-free at radius $c\log_{d-1} N$ and

$$\lambda(G) \le 2\sqrt{d-1}(1+\epsilon).$$

We also use the following bound on the number of short cycles. We use the multigraph versions of the following results, as justified by the configuration-model extension in [Par21, Remark 29].

Proposition A.7 ([Par21, Proposition 34])**.** *Let G be a d-regular N-vertex multigraph that is bicycle-free at radius $2r$. Then*

$$|\mathrm{Cyc}_r(G)| \le \frac{N}{(d-1)^r}.$$

Finally, we use the short-cycle removal construction of [Par21, Theorem 8]. Since in our application d grows with n, we state the bounds obtained by keeping explicit the dependence on d in its construction and proof.

Theorem A.8 (Short-cycle removal, adapted from [Par21, Theorem 8 and its proof])**.** *Let G be a d-regular N-vertex (r,τ)-multigraph satisfying $r \le (2/3)\log_{d-1}(N/\tau) - 5$. Set $h := \lceil \log_{d-1}\tau \rceil + \lceil r/2 \rceil + 1$ and*

$$T(d,h) := 1 + d\sum_{j=0}^{h-1}(d-1)^j = 1 + \frac{d((d-1)^h - 1)}{d-2}.$$

Then the deterministic algorithm fix *of [Par21] runs in time* $\mathrm{poly}(N)$ *and outputs a d-regular graph* $\mathrm{fix}(G)$ *of girth at least r such that*

$$|V(\mathrm{fix}(G))| \le N + 2T(d,h)$$

and

$$\lambda(\mathrm{fix}(G)) \le \max\{\lambda(G), 2\sqrt{d-1}\} + \frac{2(\sqrt{d-1}+1)}{\lfloor r/2 \rfloor} + \frac{2dT(d,h)}{N}.$$

We can now prove Proposition 2.5, which we restate below.

Proposition A.9. *There exists a family of n-regular graphs $(E_n)_n$ on $N_E = 2^{n^3}(1+o(1))$ vertices such that, for all sufficiently large n, E_n has spectral gap at least $n/2$ and girth at least $g_n = 100n^{5/2}\log n$.*

Proof. Let $N := 2^{n^3}$, $d := n$, and $L := \log_{d-1} N$. Let α be the universal constant from Theorem A.6. We apply that theorem with $c = 1/5$ and $\epsilon = 1$.

We first verify its assumptions. Since $N = 2^{n^3}$, $\sqrt{\log N} = \Theta(n^{3/2})$, and therefore

$$\frac{d}{\sqrt{\log N}} = O(n^{-1/2}) \longrightarrow 0.$$

Since α is a fixed constant, it follows that $d \le \alpha^{-1}\sqrt{\log N}$ for all sufficiently large n. Moreover,

$$\log_{d-1} N = \Theta\left(\frac{n^3}{\log n}\right), \qquad \log\log N = \Theta(\log n),$$

and hence

$$\alpha^3\left(\frac{\log\log N}{\log_{d-1} N}\right)^2 = O\left(\alpha^3\frac{\log^4 n}{n^6}\right) = o(1).$$

Thus the latter quantity is at most $\epsilon = 1$ for all sufficiently large n. The remaining condition $c < 1/4$ is immediate.

Theorem A.6 therefore implies that, with positive probability, a realization G of the d-regular N-vertex uniform configuration model is bicycle-free at radius $L/5$ and satisfies

$$\lambda(G) \le 4\sqrt{d-1}.$$

In particular, there exists such a realization G.

Set $s := \lfloor L/10 \rfloor$. Since $2s \le L/5$, Proposition A.7 gives

$$|\mathrm{Cyc}_s(G)| \le \frac{N}{(d-1)^s}.$$

Define $\tau := \lceil N/(d-1)^s \rceil$ and $r := \lfloor L/20 \rfloor$. Since $r \le s$, the graph G is an (r,τ)-multigraph.

We next verify the hypothesis of Theorem A.8. By definition, we have $\tau \le 2N/(d-1)^s$ and therefore $\log_{d-1}(N/\tau) \ge s - \log_{d-1} 2$. As $s = (L/10) + O(1)$, $(2/3)\log_{d-1}(N/\tau) - 5 \ge (L/15) - O(1)$. On the other hand, $r \le L/20$. Since L grows with N, for all sufficiently large n we have

$$r \le \frac{2}{3}\log_{d-1}(N/\tau) - 5.$$

Thus Theorem A.8 applies. Set $E_n := \mathrm{fix}(G)$. By Theorem A.8, E_n is d-regular and $\mathrm{girth}(E_n) \ge r$. Moreover, $r = \Theta\left(n^3/\log n\right) > 100n^{5/2}\log n$ for all sufficiently large n. Hence $\mathrm{girth}(E_n) \ge g_n$. However it should be noted that this algorithm runs in $\mathrm{poly}(N) = 2^{\mathrm{poly}(n)}$ time and is thus exponential in n, the degree of the graph.

We next bound the number of vertices. Recall from Theorem A.8 that the two d-regular trees used in the construction have height

$$h = \lceil \log_{d-1}\tau \rceil + \lceil r/2 \rceil + 1.$$

Since $\tau \le 2N/(d-1)^s$,

$$\log_{d-1}\tau \le L - s + \log_{d-1} 2.$$

Using $s = L/10 + O(1)$ and $r = L/20 + O(1)$, we obtain

$$h \le L - s + \frac{r}{2} + O(1) = \frac{37}{40}L + O(1).$$

A d-regular tree of height h contains

$$T(d,h) = 1 + \frac{d((d-1)^h - 1)}{d-2} \le 1 + 3(d-1)^h$$

vertices for $d \ge 3$. Consequently,

$$T(d,h) \le (d-1)^{\frac{37}{40}L + O(1)} = N^{37/40 + o(1)},$$

where we used $d = n = N^{o(1)}$. The explicit size bound in Theorem A.8 therefore gives

$$N \le |V(E_n)| \le N + N^{37/40+o(1)} = N(1 + o(1)).$$

Thus $N_E = |V(E_n)| = 2^{n^3}(1 + o(1))$.

Finally, we bound the spectrum. By the explicit form of Theorem A.8,

$$\lambda(E_n) \le \max\{\lambda(G), 2\sqrt{d-1}\} + \frac{2(\sqrt{d-1}+1)}{\lfloor r/2 \rfloor} + \frac{2dT(d,h)}{N}.$$

Since $\lambda(G) \le 4\sqrt{d-1}$, the first term is at most $4\sqrt{d-1}$. Furthermore,

$$\frac{\sqrt{d}}{r} = O\left(\frac{\log n}{n^{5/2}}\right) = o(1),$$

while

$$\frac{dT(d,h)}{N} \le n\, N^{-3/40+o(1)} = N^{-3/40+o(1)} = o(1).$$

It follows that $\lambda(E_n) \le 4\sqrt{d-1} + o(1)$. In particular, since $\lambda_2(E_n) \le \lambda(E_n)$,

$$d - \lambda_2(E_n) \ge d - 4\sqrt{d-1} - o(1) \ge \frac{d}{2}$$

for all sufficiently large n. Since $d = n$, the spectral gap of E_n is at least $n/2$, as claimed. □

The same result holds even when $d = n - \sqrt{n}$ with the same girth by a similar argument. This lets us prove Proposition 2.6.

B Spectral properties of the decorated graph

In this section, we prove the spectral properties of the graph $G_{n,E}$ that are summarized in Lemma 2.12. The study of the top eigenvector of its adjacency matrix is closely related to the analysis in Appendix B of [GHV21]. However, their argument proceeds via an iterative decoration process in which identical complete trees are attached to all vertices at each stage. Our construction differs in that we do not decorate the leaves created at intermediate levels of the iteration as in Definitions 2.7 and 2.10. Consequently, the spectral analysis in [GHV21] does not apply directly to our setting. We therefore present a self-contained proof that analyzes the top-eigenvector of the main graph directly, rather than via an iterative argument.

Below is a lemma proved in [Has21] about the spectral gap of the sum of two Hermitian matrices of the same dimension in terms of the spectral gap of one and the spectral norm of another.

Lemma B.1 (Spectral gap of sum of matrices, (Lemma 4 in [Has21]; Footnote 15 in [GHV21])). *Let A_1 and A_2 be any two Hermitian matrices of the same dimension. If A_1 has a spectral gap of δ and A_2 has spectral norm at most γ, then $A_1 + A_2$ has a spectral gap of at least $\delta - 2\gamma$.*

Applying the above result to the main graph, we obtain the following, which proves the first part of Lemma 2.12:

Lemma B.2 (Spectral gap of G_{n,E_n}). *The difference between the two largest eigenvalues of A_{n,E_n} is at least $d_n/2 - 4\sqrt{2n}$.*

Proof. The adjacency matrix A_{n,E_n} of the graph G_{n,E_n} can be written as $A_{n,E_n} = A_E + A_T$, with A_E denoting the expander part and A_T denoting the tree edges added to it. By the assumptions of Proposition 2.5, the spectral gap of A_E is at least $n/2$. The graph formed by the added tree edges is a forest of maximum degree $2n$, and thus the corresponding adjacency matrix A_T has spectral norm bounded by $2\sqrt{2n}$.

Applying Lemma B.1 with $A_1 = A_E$, $A_2 = A_T$, $\delta = n/2$ and $\gamma = 2\sqrt{2n}$ proves the lemma. □

We next develop a general description of the Perron–Frobenius eigenpair of a graph obtained by attaching the same collection of rooted trees to every vertex of a connected base graph.

Let E be a connected graph whose adjacency matrix is denoted by A_E and $\mathcal{T}$ be a multiset of finite rooted trees where the adjacency matrix of $T \in \mathcal{T}$ is denoted by A_T. Consider the graph G obtained by attaching every $T \in \mathcal{T}$ to $v \in V_E$ through an edge $\{v, t\}$ where t is the root of T. Let A_G be its adjacency matrix.

The results below express the Perron–Frobenius eigenvalue and eigenvector of G in terms of those of E and the attached trees. We will later apply them to $G_{n,E}$.

Lemma B.3 (Top eigenvector under rooted-tree decorations). *Fix a Perron–Frobenius vector ψ_E with arbitrary normalization, and let λ_E denote the corresponding eigenvalue of the adjacency matrix A_E. Let $\mathcal{T}$ be a finite multiset of finite rooted trees. For $T \in \mathcal{T}$, denote its root by t. Denote the Perron–Frobenius eigenvalue of A_G, the adjacency matrix of G, by λ_G.*

Then we have the following:

1. $\lambda_G = \lambda_E + \sum_{T\in\mathcal{T}} \left[(\lambda_G \cdot I - A_T)^{-1}\right]_{t,t}$.
2. *The restriction of any Perron–Frobenius vector of G to E is proportional to ψ_E. Let ψ_G be the Perron–Frobenius vector of G whose restriction to E exactly equals ψ_E. Then,*

$$\psi_G = \psi_E \oplus \bigoplus_{v\in V_E} \psi_E(v) \left(\bigoplus_{T\in\mathcal{T}} (\lambda_G \cdot I - A_T)^{-1} e_t \right).$$

Proof. For each $T \in \mathcal{T}$, we use the same notation for every copy of T attached to a vertex of E. In particular, we denote the root of each copy by t, and its adjacency matrix by A_T.

The adjacency matrix of G has the block form

$$A_G = \begin{pmatrix} A_E & B \\ B^\top & \bigoplus_{v \in V_E} \bigoplus_{T \in \mathcal{T}} A_T \end{pmatrix},$$

where B records the edges joining each vertex $v \in V_E$ to the root t of every copy of T attached to it.

Let $\psi_G = (x, y)$ be a Perron–Frobenius eigenvector of A_G, where $x \in \mathbb{R}^{V_E}$ and $y = \bigoplus_{v \in V_E} \bigoplus_{T \in \mathcal{T}} y_{v,T}$. The eigenvalue equation $A_G \psi_G = \lambda_G \cdot \psi_G$ is equivalent to

$$A_E x + B y = \lambda_G x$$

and

$$B^\top x + \left(\bigoplus_{v,T} A_T \right) y = \lambda_G y.$$

We first solve the second equation. Since the blocks are independent, for every pair (v, T) we obtain $(\lambda_G \cdot I - A_T) y_{v,T} = e_t x_v$. The matrix $\lambda_G \cdot I - A_T$ is invertible because $\lambda_G > \lambda_T$.

To see this, let z be a Perron–Frobenius eigenvector of A_T, and extend it by zero outside T to obtain a vector $\widetilde{z}$ on G. The Rayleigh quotient of $\widetilde{z}$ is λ_T, so $\lambda_T \leq \lambda_G$. If equality held, then $\widetilde{z}$ would be a top eigenvector of A_G. However, if v is the vertex of E attached to the root of this copy of T, then $\widetilde{z}_v = 0$, whereas $(A_G \widetilde{z})_v = \widetilde{z}_t > 0$, contradicting $A_G \widetilde{z} = \lambda_G \widetilde{z}$. Hence $\lambda_T < \lambda_G$.

Therefore, $y_{v,T} = (\lambda_G \cdot I - A_T)^{-1} e_t \, x_v$. Substituting this into the first block equation gives, for every $v \in V_E$,

$$(A_E x)_v + \sum_{T \in \mathcal{T}} \left[(\lambda_G \cdot I - A_T)^{-1} \right]_{t,t} x_v = \lambda_G x_v.$$

Equivalently,

$$\left(A_E + \sum_{T \in \mathcal{T}} \left[(\lambda_G \cdot I - A_T)^{-1} \right]_{t,t} I \right) x = \lambda_G x.$$

The vector x is positive, so by the Perron–Frobenius theorem it is proportional to ψ_E. Comparing the corresponding eigenvalues gives

$$\lambda_G = \lambda_E + \sum_{T \in \mathcal{T}} \left[(\lambda_G \cdot I - A_T)^{-1} \right]_{t,t},$$

which proves part (1).

Finally, choose the normalization $x = \psi_E$. Then

$$y_{v,T} = \psi_E(v) (\lambda_G \cdot I - A_T)^{-1} e_t.$$

Combining the expressions for x and y gives

$$\psi_G = \psi_E \; \oplus \bigoplus_{v \in V_E} \psi_E(v) \left(\bigoplus_{T \in \mathcal{T}} (\lambda_G \cdot I - A_T)^{-1} e_t \right),$$

which proves part (2). □

Next we show that most of the Perron mass is supported on the base graph whenever the number and spectral norms of the decorations are sufficiently small relative to λ_G.

Corollary B.4 (Overlap with the base graph)**.** *Consider the setting of Lemma B.3. Let λ_G be the Perron eigenvalue of G, let ψ_G be a corresponding Perron–Frobenius eigenvector, and write $\psi_E := \psi_G|_{V_E}$. Define $\eta := \max_{T\in\mathcal{T}} \frac{\|A_T\|}{\lambda_G}$. Then,*

$$\frac{\|\psi_E\|_2^2}{\|\psi_G\|_2^2} \geq \left(1 + \frac{|\mathcal{T}|}{\lambda_G^2(1-\eta)^2}\right)^{-1}.$$

Proof. For every $T \in \mathcal{T}$, the strict inequality $\|A_T\| < \lambda_G$ follows from the proof of Lemma B.3. Hence $\eta < 1$. For $v \in V_E$ and $T \in \mathcal{T}$, let T_v denote the copy of T attached to v with root t, and let ψ_{T_v} denote the restriction of ψ_G to T_v. By Lemma B.3,

$$\psi_{T_v} = \psi_E(v)(\lambda_G \cdot I - A_T)^{-1} e_t.$$

Since $\|A_T\| < \lambda_G$, we have that $(\lambda_G \cdot I - A_T)$ is positive definite. Therefore,

$$\|\psi_{T_v}\|_2 \leq |\psi_E(v)| \left\|(\lambda_G \cdot I - A_T)^{-1}\right\|_2 \leq \frac{|\psi_E(v)|}{(\lambda_G - \|A_T\|)} \leq \frac{|\psi_E(v)|}{\lambda_G(1-\eta)}.$$

Since every decoration vertex belongs to a unique copy attached to a unique base vertex,

$$\|\psi_G\|_2^2 - \|\psi_E\|_2^2 \leq \sum_{v\in V_E} \sum_{T_v\in\mathcal{T}} \frac{\psi_E(v)^2}{\lambda_G^2(1-\eta)^2} \leq \frac{|\mathcal{T}|}{\lambda_G^2(1-\eta)^2}\|\psi_E\|_2^2.$$

Rearranging proves the claim. □

We now specialize the generic estimates to the two graph families used in the construction. We begin with the main graph G_n, for which the base graph is n-regular and the decorations are the trees $\widehat{T}_{n,k}$ with $k \in [\sqrt{n} - 1]$.

Corollary B.5 (Top-eigenvalue and expander overlap for G_n and F_n)**.** *Let $\lambda_{n,E}$ be the Perron eigenvalue of $G_{n,E}$ for $(K, E) \in \{(\sqrt{n}, E_n), (\sqrt{n}+1, E'_n)\}$ (cf. Definition 2.10), and let $\psi_{n,E}$ be the corresponding Perron–Frobenius eigenvector. Write $\psi_E := \psi_{n,E}|_{V(E)}$. Then,*

1. $\lambda_{n,E} = d_{n,K} + O(1)$.
2. $\|\psi_E\|_2^2 \geq (1 - O(1/n)) \cdot \|\psi_{n,E}\|_2^2$.

Proof. At each vertex of E, the multiset of attached trees is $\mathcal{T} = \biguplus_{k=1}^{K-1} \biguplus_{j=1}^{\sqrt{n}} \{\widehat{T}_{n,k}\}$, and hence $|\mathcal{T}| = \sqrt{n} \cdot (K-1)$. By Lemma B.3, $\lambda_{n,E} = d_{n,K} + \sum_{T\in\mathcal{T}} \left[(\lambda_{n,E} \cdot I - A_T)^{-1}\right]_{t,t}$. Every tree in $\mathcal{T}$ has maximum degree at most $2n$, so $\|A_T\| \leq 2\sqrt{2n}$. Since $\lambda_{n,E} \geq d_{n,K}$, the matrix $\lambda_{n,E} \cdot I - A_T$ is positive definite, and

$$0 \leq \left[(\lambda_{n,E} \cdot I - A_T)^{-1}\right]_{t,t} \leq \|(\lambda_{n,E} \cdot I - A_T)^{-1}\| \leq \frac{1}{d_{n,K} - 2\sqrt{2n}}.$$

Here the first inequality follows by writing the diagonal entry as $e_t^\top (\lambda_{n,E} \cdot I - A_T)^{-1} e_t$, while the last follows since the eigenvalues of $(\lambda_{n,E} \cdot I - A_T)^{-1}$ are $(\lambda_{n,E} - \mu)^{-1}$ for μ in the spectrum of A_T. Therefore $0 \leq \lambda_{n,E} - d_{n,K} \leq \sqrt{n} \cdot (K-1)/(d_{n,K} - 2\sqrt{2n}) = O(1)$. When $E \in \{E_n, E'_n\}$, we get that $\lambda_{n,E} = d_{n,K} + O(1)$ and $\max_{T\in\mathcal{T}} \|A_T\|/\lambda_{n,E} \leq 2\sqrt{2n}/d_{n,K} = O(1/\sqrt{n})$. The overlap bound follows from Corollary B.4. □

C Hardness of exploring self-similar trees

This section is devoted to proving that any classical algorithm requires a sub-exponential number of queries to explore a 0-level leaf in the full-level decorated tree $\widehat{T}_{n,k}$ starting from its root (see Definition 2.7 for a description of these objects). The main result is Lemma 3.13 (restated below as Proposition C.6). The proof proceeds by induction on the index k.

We begin by introducing several probabilistic events and random variables that play a central role in the argument. Observe that $\widehat{\textsc{Exit}}_{n,K}$ for $K = \sqrt{n}$ is precisely the event whose probability we aim to bound in order to prove Lemma 3.13.

Definition C.1 (Events and random variables)**.** Given a randomized algorithm that explores the 1-level decorated tree $T_{n,k,1}$ starting at its root, we define:

- $\textsc{Exit}_{n,k}$: the event that the algorithm queries a 0-level leaf,
- $W_{n,k}$: the random variable denoting the number of queried 1-level leaves that belong to *distinct* decorations.

The same objects are denoted $\widehat{\textsc{Exit}}_{n,k}$ and $\widehat{W}_{n,k}$ when the algorithm is instead exploring the full-level decorated tree $\widehat{T}_{n,k}$.

For a randomized algorithm $\mathcal{A}$, let $q_{\mathcal{A}}$ denote the maximum number of queries made by that algorithm (over all possible random executions). Our goal is to identify the largest value of $q_{n,k}$ for which the event $\widehat{\textsc{Exit}}_{n,k}$ is unlikely to occur when the algorithm makes at most $q_{n,k}$ queries. This amounts to upper-bounding the quantity $\max_{\mathcal{A}:q_{\mathcal{A}}<q_{n,k}} \Pr[\widehat{\textsc{Exit}}_{n,k}]$. The inductive proof will relate this to $\max_{\mathcal{A}:q_{\mathcal{A}}<q_{n,k-1}} \Pr[\widehat{\textsc{Exit}}_{n,k-1}]$, where the number of queries is decreased to $q_{n,k-1} < q_{n,k}$ and the input tree is $\widehat{T}_{n,k-1}$.

First, it is easy to see that full-level decorated trees are harder to explore than 1-level decorated trees. In particular, we will make use of the following lemma.

Lemma C.2 (Monotonicity of hardness under decorations)**.** *For any number q of queries and integer w, we have*

$$\max_{\mathcal{A}:q_{\mathcal{A}}<q} \Pr[\widehat{\textsc{Exit}}_{n,k} \text{ and } \widehat{W}_{n,k} < w] \leq \max_{\mathcal{A}:q_{\mathcal{A}}<q} \Pr[\textsc{Exit}_{n,k} \text{ and } W_{n,k} < w].$$

Proof. Any algorithm with query access to $T_{n,k,1}$ can simulate query access to the full-level decorated tree $\widehat{T}_{n,k}$ by simply simulating the subsequent levels of decoration itself. Hence, the success probability cannot be larger on $\widehat{T}_{n,k}$ than on $T_{n,k,1}$. □

The next central lemma upper bounds the probability of the event $\textsc{Exit}_{n,k}$ occurring when only a few 1-level leaves of the 1-level decorated tree $T_{n,k,1}$ are explored. Note that the input graph need not be fully-decorated, nor does the number $q_{\mathcal{A}}$ of queries need to be bounded to apply this lemma.

Lemma C.3 (Hardness of avoiding the decorations)**.** *For any integer w,*

$$\max_{\mathcal{A}:q_{\mathcal{A}}<\infty} \Pr[\textsc{Exit}_{n,k} \text{ and } W_{n,k} < w] \leq \left(\frac{d_{n,k}}{d_{n,k-1}}\right)^{(\ell_{n,k}-\ell_{n,k-1})/w}.$$

Proof. The cases $w = 1$ and $w = 2$ correspond essentially to [GHV21, Lemma 3] and [GHV21, Lemma 4, Case 1]. We do not prove the case $w > 2$ here since it is not needed for our application.

We sketch the proof for $w = 1$. We view an algorithm $\mathcal{A}$ as a local exploration that starts at the root of $T_{n,k,1}$ and may only query edges adjacent to previously queried vertices. The events $\textsc{Exit}_{n,k}$ and $W_{n,k} < 1$ require the first queried leaf in $T_{n,k,1}$ to be a 0-level leaf. Thus,

the algorithm must grow an exploration path of length $\ell_{n,k}$ that reaches a 0-level leaf before any 1-level leaf is queried. A path that grows to length $\ell_{n,k}$ avoids reaching a 1-level leaf if and only if it does not enter any decoration before depth $\ell_{n,k} - \ell_{n,k-1}$, since the decorations have depth $\ell_{n,k-1}$. For each node in the base graph, the proportion of its children that do not belong to a decoration is at most $d_{n,k}/d_{n,k-1}$. Hence, the probability of avoiding any decoration for $\ell_{n,k} - \ell_{n,k-1}$ consecutive steps is at most $(d_{n,k}/d_{n,k-1})^{\ell_{n,k}-\ell_{n,k-1}}$. □

The next lemma exploits the self-similar structure of $\widehat{T}_{n,k}$ to relate the probability of finding multiple 1-level leaves in $\widehat{T}_{n,k}$ to the probability of the event $\widehat{\textsc{Exit}}_{n,k-1}$. Most importantly, the number of queries allowed to explore $\widehat{T}_{n,k-1}$ is reduced by a factor of $1/w$, which will make the recursion effective later on.

Lemma C.4 (Self similarity)**.** *For any integer w,*

$$\max_{\mathcal{A}:q_{\mathcal{A}}<q} \Pr[\widehat{W}_{n,k} \geq w] \leq q \cdot \max_{\mathcal{A}:q_{\mathcal{A}}<q/w} \Pr[\widehat{\textsc{Exit}}_{n,k-1}].$$

Proof. This generalizes [GHV21, Lemma 4, Case 2]. Essentially, finding leaves of w distinct outer trees in $\widehat{T}_{n,k}$ amounts to the event $\widehat{\textsc{Exit}}_{n,k-1}$ occurring w times in w distinct subtrees $\widehat{T}_{n,k-1}$ of $\widehat{T}_{n,k}$. By a pigeonhole argument, at least one of these events must happen after making only q/w queries in the corresponding subtree $\widehat{T}_{n,k-1}$. Since at most q subtrees are explored in total (this is a very crude bound), by a union bound argument, the overall probability must be at most $q \cdot \max_{\mathcal{A}:q_{\mathcal{A}}<q/w} \Pr[\widehat{\textsc{Exit}}_{n,k-1}]$. □

Combining the above three results, we obtain the inductive step of the proof:

Corollary C.5 (Inductive argument)**.** *For any integer w,*

$$\max_{\mathcal{A}:q_{\mathcal{A}}<q_{n,k}} \Pr[\widehat{\textsc{Exit}}_{n,k}] \leq \left(\frac{d_{n,k}}{d_{n,k-1}}\right)^{(\ell_{n,k}-\ell_{n,k-1})/w} + q_{n,k} \cdot \max_{\mathcal{A}:q_{\mathcal{A}}<q_{n,k}/w} \Pr[\widehat{\textsc{Exit}}_{n,k-1}]. \tag{5}$$

Proof. It is simply a matter of applying Lemmas C.2 to C.4, and basic laws of probability,

$$\begin{aligned}
\max_{\mathcal{A}:q_{\mathcal{A}}<q_{n,k}} \Pr[\widehat{\textsc{Exit}}_{n,k}] &\leq \max_{\mathcal{A}:q_{\mathcal{A}}<q_{n,k}} \Pr[\widehat{\textsc{Exit}}_{n,k} \text{ and } \widehat{W}_{n,k} < w] + \max_{\mathcal{A}:q_{\mathcal{A}}<q_{n,k}} \Pr[\widehat{\textsc{Exit}}_{n,k} \text{ and } \widehat{W}_{n,k} \geq w] \\
&\leq \max_{\mathcal{A}:q_{\mathcal{A}}<\infty} \Pr[\textsc{Exit}_{n,k} \text{ and } W_{n,k} < w] + \max_{\mathcal{A}:q_{\mathcal{A}}<q_{n,k}} \Pr[\widehat{W}_{n,k} \geq w] \\
&\leq \left(\frac{d_{n,k}}{d_{n,k-1}}\right)^{(\ell_{n,k}-\ell_{n,k-1})/w} + q_{n,k} \cdot \max_{\mathcal{A}:q_{\mathcal{A}}<q_{n,k}/w} \Pr[\widehat{\textsc{Exit}}_{n,k-1}].
\end{aligned}$$

□

Finally, we solve the above recurrence relation using $w = 2$ and the set of parameters given in Lemma 3.13.

Proposition C.6 (Finding exit is hard - Lemma 3.13 restated)**.** *For $k \in \{1, \ldots, \sqrt{n}\}$, let $d_{n,k} = 2n - k\sqrt{n}$, $\ell_{n,k} = k \cdot 10n^{3/2} \log n$, $q_{n,k} = 2^k$ and $w = 2$. Then,*

$$\max_{\mathcal{A}:q_{\mathcal{A}}<q_{n,k}} \Pr[\widehat{\textsc{Exit}}_{n,k}] \leq 2^{-2n\log n + (\sqrt{n}+1)(k-\sqrt{n})}.$$

In particular, for $K = \sqrt{n}$, the probability of $\widehat{\textsc{Exit}}_{n,K}$ is at most $2^{-2n\log n}$ for any classical algorithm making fewer than $2^{\sqrt{n}}$ queries. Moreover, the tree $\widehat{T}_{n,K}$ contains fewer than $2^{12n^3\log^2 n}$ vertices.

Proof. The base case $k = 1$ is trivial. For the induction, we have,

$$\begin{aligned}\left(\frac{d_{n,k}}{d_{n,k-1}}\right)^{(\ell_{n,k}-\ell_{n,k-1})/2} &= \left(1 - \frac{1}{2\sqrt{n}-k+1}\right)^{(10n^{3/2}\log n)/2} \\ &\le \left(1 - \frac{1}{2\sqrt{n}}\right)^{5n^{3/2}\log n} \\ &\le 2^{-\frac{5}{2}n\log n}.\end{aligned}$$

Hence, by Corollary C.5,

$$\begin{aligned}\max_{\mathcal{A}:q_{\mathcal{A}}<q_{n,k}} \Pr[\widehat{\textsc{Exit}}_{n,k}] &\le 2^{-\frac{5}{2}n\log n} + 2^k \cdot 2^{-2n\log n+(\sqrt{n}+1)(k-1-\sqrt{n})} \\ &\le 2^{-\frac{5}{2}n\log n} + 2^{\sqrt{n}} \cdot 2^{-2n\log n+(\sqrt{n}+1)(k-1-\sqrt{n})} \\ &\le 2^{-2n\log n+(\sqrt{n}+1)(k-\sqrt{n})}.\end{aligned}$$

The number of vertices in $\widehat{T}_{n,k}$ is $d_{n,k}^{\ell_{n,k}} \cdot \prod_{i=k-1}^{1}(d_{n,i} - d_{n,i+1})d_{n,i}^{\ell_{n,i}} \le 2^{12n^3\log^2 n}$. □