Approximate Low-Weight Check Codes and Circuit Lower Bounds for Noisy Ground States

Chinmay Nirkhe
Electrical Engineering and Computer Sciences, University of California, Berkeley
387 Soda Hall Berkeley, CA 94720, U.S.A.
https://people.eecs.berkeley.edu/~nirkhe/
nirkhe@cs.berkeley.edu

Umesh Vazirani
Electrical Engineering and Computer Sciences, University of California, Berkeley
387 Soda Hall Berkeley, CA 94720, U.S.A.
https://people.eecs.berkeley.edu/~vazirani/
vazirani@cs.berkeley.edu

Henry Yuen
Electrical Engineering and Computer Sciences, University of California, Berkeley
387 Soda Hall Berkeley, CA 94720, U.S.A.
https://www.henryyuen.net/
hyuen@cs.berkeley.edu

Abstract

The No Low-Energy Trivial States (NLTS) conjecture of Freedman and Hastings (Quantum Information and Computation 2014), which asserts the existence of local Hamiltonians whose low-energy states cannot be generated by constant-depth quantum circuits, identifies a fundamental obstacle to resolving the quantum PCP conjecture. Progress towards the NLTS conjecture was made by Eldar and Harrow (Foundations of Computer Science 2017), who proved a closely related theorem called No Low-Error Trivial States (NLETS). In this paper, we give a much simpler proof of the NLETS theorem and use the same technique to establish superpolynomial circuit size lower bounds for noisy ground states of local Hamiltonians (assuming $\mathsf{QCMA} \neq \mathsf{QMA}$), resolving an open question of Eldar and Harrow. We discuss the new light our results cast on the relationship between NLTS and NLETS.

Finally, our techniques imply the existence of *approximate quantum low-weight check (qLWC) codes* with linear rate, linear distance, and constant weight checks. These codes are similar to quantum LDPC codes except (1) each particle may participate in a large number of checks, and (2) errors only need to be corrected up to fidelity $1 - 1/\mathsf{poly}(n)$. This stands in contrast to the best-known stabilizer LDPC codes due to Freedman, Meyer, and Luo which achieve a distance of $O(\sqrt{n \log n})$.

The principal technique used in our results is to leverage the Feynman-Kitaev clock construction to approximately embed a subspace of states defined by a circuit as the ground space of a local Hamiltonian.

2012 ACM Subject Classification Theory of computation → Quantum complexity theory

Keywords and phrases quantum pcps, local hamiltonians, error-correcting codes

Digital Object Identifier 10.4230/LIPIcs.ICALP.2018.91

Related Version A full version of the paper is available at https://arxiv.org/abs/1802.07419.

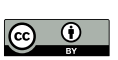

45th International Colloquium on Automata, Languages, and Programming (ICALP 2018).
Editors: Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella;
Article No. 91; pp. 91:1–91:11

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Funding This work was supported by ARO Grant W911NF-12-1-0541 and NSF Grant CCF-1410022.

Acknowledgements We thank Dorit Aharanov, Adam Bouland, Elizabeth Crosson, Bill Fefferman, Lior Eldar, Zeph Landau, Ashwin Nayak, Nicholas Spooner, and Thomas Vidick for helpful discussions.

1 Introduction

The quantum PCP conjecture [4, 2] is a central open question in quantum complexity theory. To understand the statement, it is helpful to review the dictionary translating between classical constraint satisfaction problems (CSPs) and their quantum analogue, the local Hamiltonian problem. A classical CSP on n variables corresponds to a local Hamiltonian $H = H_1 + \cdots + H_m$ acting on n qubits[1]. A solution to the CSP corresponds to an n qubit quantum state, and the number of violated constraints corresponds to the energy (eigenvalue) of that quantum state. The NP-hardness of SAT corresponds to the QMA-hardness of deciding whether the H has minimum eigenvalue at most a or at least b for given a, b such that $b - a = 1/\mathsf{poly}(n)$. The quantum analogue of the PCP theorem, called the qPCP conjecture, asserts that the problem remains QMA-hard even when $b - a \geq cm = c\|H\|$.

Just as the classical PCP theorem connects coding theory to constraint satisfaction problems, it is natural to expect any resolution of the quantum PCP conjecture to rely on — and to reveal — deep connections between the theory of quantum error-correcting codes and ground states (i.e. states of minimum energy) of local Hamiltonians. Examples of quantum error-correcting codes realized as the ground spaces of local Hamiltonians already play a central role in our understanding of the physical phenomenon known as topological order [30]. Moreover, it has been suggested that the qPCP conjecture is closely related to one of the biggest open problems in quantum coding theory: whether quantum low density parity check (qLDPC) codes with linear rate and linear distance exist [23, 13, 33].

The difficulty of the qPCP conjecture motivated Freedman and Hastings to formulate a simpler goal called the *No Low-Energy Trivial States (NLTS) Conjecture* [22]. One way to put one's finger on the additional difficulty of qPCP (beyond the "standard" difficulty of proving a classical PCP theorem) is that solutions of QMA-hard problems are expected to have high description complexity. For example, if NP $\neq$ QMA, then ground states of local Hamiltonians do not have classically checkable polynomial-size descriptions. The NLTS conjecture isolates this aspect of high description complexity by asserting the existence of a family of local Hamiltonians $\{H^{(n)}\}_{n=1}^{\infty}$ where $H^{(n)}$ acts on n particles, such that low-energy states (of energy less than $c\|H\|$) cannot be generated by quantum circuits of constant depth. A much stronger version of the NLTS conjecture is a necessary consequence of the qPCP conjecture: assuming QCMA $\neq$ QMA,[2] low-energy states cannot be described even by polynomial-size quantum circuits. However, one of the advantages of the NLTS conjecture is that it does not involve complexity classes such as QMA, but rather focuses on the entanglement complexity that is intrinsic to low-energy states of local Hamiltonians.

[1] For normalization, we assume that the terms of a local Hamiltonian have spectral norm at most 1.

[2] For precise definitions of the complexity classes QCMA and QMA, we refer the reader to [20, 29]. Roughly speaking, QMA is the class of problems for which the solution is a quantum state that can be efficiently checked by a quantum computer. QCMA is the class of problems where the solution is a *classical* string that can be efficiently checked by a quantum computer.

Like the qPCP conjecture, the NLTS conjecture remains unresolved. In [19], Eldar and Harrow proposed a variant of the NLTS called *No Low-Error Trivial States (NLETS)*, which is itself a necessary consequence[3] of NLTS. The difference was that rather than considering low-energy states of H, they considered a notion of "local corruption error", what they call *ϵ-error states*: these are states that differ from the ground state in at most ϵn qubits. More precisely, σ is ϵ-error for a local Hamiltonian H if there exists a ground state ρ of H and a set S of at most ϵn qudits such that $\mathrm{Tr}_S(\rho) = \mathrm{Tr}_S(\sigma)$. Under this definition they were able to establish a family of Hamiltonians for which any ϵ-error state requires circuit depth of $\Omega(\log n)$. This was welcomed as very encouraging progress towards establishing NLTS, since NLETS could be regarded as a close proxy for NLTS, with a technical change in definition of distance under which to examine the robustness of the ground space.

In this paper, we start by giving a simple argument for the $\Omega(\log n)$ circuit depth lower bound of Eldar and Harrow; our lower bound holds even under a more general error model, which allows any probabilistic mixture of ϵ-error states (we call these states *noisy ground states*). Moreover, we can use the same techniques to answer their open question of whether one can obtain circuit size lower bounds on low-error states that go beyond logarithmic depth: specifically, we show that there exists a family of local Hamiltonians whose noisy ground states require superpolynomial-size circuits, assuming $\mathsf{QCMA} \neq \mathsf{QMA}$.

One way to view these results is that they provide further progress towards the NLTS conjecture and beyond. However, it is instructive to take a step back to consider more closely the basic difference between NLETS and NLTS. This lies in the different notion of approximation: in NLETS, approximation corresponds to local corruptions in ϵn sites, where n is the total number of particles, whereas in NLTS approximation corresponds to energy at most $\epsilon \|H\|$ (intuitively, at most ϵ fraction of the terms of the Hamiltonian are violated). An alternative perspective on our results is that they suggest these two notions of approximation are quite different. This view is reinforced by the fact that our $\Omega(\log n)$-circuit depth lower bounds on noisy ground states holds for a family of 1D Hamiltonians, whereas we know that NLTS and qPCP Hamiltonians cannot live on any constant-dimensional lattice [2]. This suggests that in the context of the qPCP and NLTS conjectures, the correct notion of distance is given by the energy or number of violated terms of the Hamiltonian.

On the other hand, the local corruption distance as defined by Eldar and Harrow for their NLETS result is the natural one that arises in quantum error correction: the distance of a code is *defined* by the maximum number of qubits of a codeword that can be erased while maintaining recoverability. We give a construction of a family of codes (inspired by the construction used in our noisy ground state lower bound) that we call *quantum low weight check (qLWC) codes*. The family of codes we consider are approximate error-correcting codes in the sense of [16, 11]. They are closely related to qLDPC codes, with the difference that they are not stabilizer codes and therefore the low weight checks are not Pauli operators. Specifically, we give a family of approximate qLWCs with linear distance and linear rate. Constructing qLDPC codes with similar parameters is a central open question in coding theory, with the best-known stabilizer LDPC codes due to Freedman, Meyer, and Luo which achieve a distance of $O(\sqrt{n \log n})$ [23].

What is common to the above results is the technique. We start with the observation that the complicated part of the Eldar and Harrow proof is constructing a local Hamiltonian whose ground states share some of the properties of the cat state $|\circledast_n\rangle = (|0\rangle^{\otimes n} + |1\rangle^{\otimes n})/\sqrt{2}$.

[3] The local Hamiltonian family must be of bounded-degree, meaning no particle participates in more than a constant number of Hamiltonian terms.

To do so, they constructed a local Hamiltonian corresponding to a quantum error-correcting code (based on the Tillich-Zemor hypergraph product construction) and showed that its ground states have non-expansion properties similar to those of the cat state [19]. Our starting point is the observation that the Tillich-Zemor construction is unnecessary and that one can make the cat state *approximately* a ground state of a local Hamiltonian in the following sense: we construct a Feynman-Kitaev clock Hamiltonian corresponding to the circuit that generates $|\boxtimes_n\rangle$ from $|0\rangle^{\otimes n}$.[4] The ground state of this Hamiltonian is the history state of this computation, and we directly argue that the circuit depth necessary to generate this history state is at least $\Omega(\log n)$. This same argument even allows us to lower bound the circuit depth of *approximate* noisy ground states (i.e. states that are close in trace distance to a noisy ground state).

The Feynman-Kitaev clock Hamiltonian plays a central role in our construction of qLWCs, with history states playing the role of codewords. The fact that such a construction yields an error-correcting code flies in the face of classical intuition. After all, it is the brittleness of the Cook-Levin tableau [15, 32] (the classical analogue of the history state) that motivates the elaborate classical PCP constructions [7, 6, 18]. The difference is that *time* is in superposition in a quantum history state. We do not yet understand the implications of this observation. For example, is it possible that it might lead to new ways of constructing qLDPC codes with super-efficient decoding procedures? There are precedents for such connections between computational phenomena and codes, most notably with the PCP theorem and the construction of locally testable and locally checkable codes.

Furthermore, while quantum error-correcting codes have typically provided a wealth of examples of interesting local Hamiltonians, our construction of qLWCs also suggest that a fruitful connection exists in the opposite direction: by considering techniques to construct local Hamiltonians (such as the Feynman-Kitaev clock construction), we can construct an interesting example of a quantum error-correcting code. We note that this reverse connection is starting to take hold in other areas of quantum information theory and physics: see [12, 28].

2 Summary of Results

Before we present our results, we motivate our definition of *noisy ground states*.

2.1 Noisy ground states

The NLETS Theorem and NLTS conjecture describe different ways in which the ground space entanglement is robust. The ground states of NLETS Hamiltonians are robust against local corruptions in ϵn sites, where n is the total number of particles. NLTS Hamiltonians are robust against low-energy excitations in the sense that all states with energy at most $\epsilon\|H\|$ retain nontrivial circuit complexity.

In this paper, we study another way that ground space entanglement can be robust. We focus on the concept of *noisy ground state*, which is a generalization of low-error states: an ϵ-noisy ground state σ of a local Hamiltonian H is a probabilistic mixture of ϵ-error states $\{\sigma_i\}$.

This notion of noisy ground state is naturally motivated by the following situation: consider a ground state ρ of H. On each particle independently apply the following process

[4] A similar construction of a clock Hamiltonian was also considered by Crosson and Bowen in the context of idealized adiabatic algorithms [17]. The construction is inspired by techniques of [14, 10].

$\mathcal{M}$: with probability ϵ, apply a noisy channel $\mathcal{N}$, and with probability $1-\epsilon$ apply the identity channel $\mathcal{I}$. The resulting state is

$$\begin{aligned}\mathcal{M}(\rho) &= \left((1-\epsilon)\mathcal{I} + \epsilon\mathcal{N}\right)^{\otimes n}(\rho) \\ &= \sum_{S\subseteq[n]} (1-\epsilon)^{n-|S|}\epsilon^{|S|}\mathcal{N}^S(\rho) \\ &\approx \sum_{S:|S|\leq 2\epsilon n} (1-\epsilon)^{n-|S|}\epsilon^{|S|}\mathcal{N}^S(\rho)\end{aligned} \tag{1}$$

where $\mathcal{N}^S$ denotes the tensor product of the map $\mathcal{N}$ acting on the particles indexed by S. The last approximate equality follows from the fact that with overwhelmingly large probability, $\mathcal{N}^S$ acts on at most $2\epsilon n$ particles. Notice that the expression on the right hand side is (up to normalization) a 2ϵ-noisy ground state, because when $|S| \leq 2\epsilon n$, the state $\mathcal{N}^S(\rho)$ is a 2ϵ-error state.

This justifies the name "noisy ground state", as the operation $\mathcal{M}$ is is a reasonable model of noise that occurs in physical processes (and is frequently considered in work on quantum fault-tolerance). Furthermore, we believe that our model arises naturally in the context of noisy adiabatic quantum computation.

As mentioned before, noisy ground states are a generalization of low-error states but are a special case of low-energy states: since low-error states are themselves low-energy states, a convex combination of them is also low-energy.

We prove several results about the robustness of entanglement in noisy ground states.

2.2 Logarithmic circuit depth lower bound

First, we generalize Eldar and Harrow's logarithmic circuit depth lower bound [19] to encompass noisy ground states. Furthermore, we present a family of Hamiltonians that is *one-dimensional*; in other words, the particles of the Hamiltonian are arranged on a line and the Hamiltonian terms act on neighboring particles.

We call this the *Logarithmic Noisy Ground States (LNGS) Theorem*[5].

▶ **Theorem 1** (Logarithmic lower bound). *There exists a family of 3-local Hamiltonians $\{H^{(n)}\}$ on a line, acting on particles of dimension 3, such that for all $n \in \mathbb{N}$, for all $0 \leq \epsilon < 1/48$, $0 \leq \delta < \frac{1}{8} - 6\epsilon$, the δ-approximate circuit depth of any ϵ-noisy ground state σ for $H^{(n)}$ is at least $\frac{1}{2}\log(n/2)$.*

Here, the δ-approximate circuit depth of ρ means the circuit depth needed to produce a state that is δ-close to ρ in trace distance.

Our proof of Theorem 1 is simple and self-contained. As a consequence of our simpler local Hamiltonian construction, we obtain improved parameters over those in [19]. Furthermore, as we will discuss below in Section 2.4, the fact that our LNGS Hamiltonian is one-dimensional gives a strong separation between NLETS/LNGS and NLTS Hamiltonians.

2.2.1 Superpolynomial circuit size lower bound

A question that was left open by [19] is whether one can obtain circuit lower bounds on low-error states that are better than logarithmic – say polynomial or even exponential. We show that there exists a family of local Hamiltonians whose noisy ground states require

[5] We pronounce this "Longs."

superpolynomial[6] size circuits, assuming QCMA $\neq$ QMA. Since low-error states are noisy ground states, this provides an answer to Eldar and Harrow's open question.

We call this the *Superpolynomial Noisy Ground States (SNGS) Theorem*[7].

▶ **Theorem 2** (Superpolynomial Noisy Ground States (SNGS)). *If* QCMA $\neq$ QMA*, then there exists* $q, \epsilon > 0$ *and a family of* 7*-local Hamiltonians* $\{H^{(n)}\}$ *acting on dimension-*q *qudits such that for all* $0 \leq \delta < 1/5$*, the* δ*-approximate circuit complexity of any family* $\{\sigma_n\}$ *of* ϵ*-noisy ground states for* $\{H^{(n)}\}$ *grows faster than any polynomial in* n*.*

We call such a family $\{H^{(n)}\}$ *SNGS Hamiltonians*. The following is a proof sketch. Let $L = (L_{yes}, L_{no})$ be the QMA-complete language consisting of descriptions of polynomial-size verifier circuits acting on a witness state and ancilla qubits. We convert each circuit $C \in L$, into a circuit C' where C' applies in order: (a) a unitary V to encode the state in an error-correcting code[8], (b) a collection of identity gates, (c) the unitary $V^\dagger$ to decode the state, and (d) the gate circuit C. The construction maintains that the circuits C' and C are equivalent. We then generate the Feynman-Kitaev clock Hamiltonian for C'. Let H_C be this Hamiltonian. The family of SNGS Hamiltonians is precisely $\{H_C : C \in L_{yes}\}$.

In order to prove that all noisy ground states of this Hamiltonian must have superpolynomial circuit size, we show that if there was a noisy ground state with a polynomial-size generating circuit, then the description of the generating circuit would suffice as a *classical* witness for the original QMA-complete problem. In the *yes* case, the construction of C' from C enforces that tracing out the time register of the noisy ground state will yield a state close to a convex combination of $\{\mathrm{Enc}(|\xi_i, 0\rangle)\}$ where $\mathrm{Enc}(\cdot)$ is the encoding function for the error-correcting code and $\{|\xi_i\rangle\}$, a collection of accepting witness. Therefore, given the description of the generating circuit for the noisy ground state, we can generate the noisy ground state and decode the original witness state. It suffices then to check the witness by running the original circuit C. The *no* case follows easily from the definition of L_{no}. This proves that $L \in$ QCMA, proving QCMA = QMA, contradicting the original assumption.

2.2.2 Semi-explicit SNGS Hamiltonians via oracle separations

It is an open question in quantum complexity theory of whether QCMA is equal to QMA. Aaronson and Kuperberg gave the first complexity-theoretic evidence that they are different by constructing a *quantum* oracle $\mathcal{O}$ such that $\mathsf{QCMA}^{\mathcal{O}} \subsetneq \mathsf{QMA}^{\mathcal{O}}$ [1]. Fefferman and Kimmel later showed that one can obtain the same oracle separation with *in-place* oracles $\mathcal{O}$, which are permutation matrices in the standard basis [20]. The separations of [1, 20] hold as long as the locality of the oracles $\mathcal{O}$ is $\omega(\log n)$ (i.e. superlogarithmic in the problem size).

We show that any oracle separation between QCMA and QMA can be leveraged to obtain a semi-explicit family of SNGS Hamiltonians:

▶ **Theorem 3.** *There exists* $q, \epsilon > 0$*, a function* $k(n) = O(\log^{1+\alpha} n)$ *for arbitrarily small* $\alpha > 0$ *and a family of* k*-local Hamiltonians* $\{H^{(n)}\}$ *acting on dimension-*q *qudits such that the following holds: The circuit complexity of any family* $\{\sigma_n\}$ *of* ϵ*-noisy ground states for* $\{H^{(n)}\}$ *grows faster than any polynomial in* n*. Furthermore, there is exactly one term in* $H^{(n)}$ *that is* $k(n)$*-local; all other terms are* 7*-local.*

[6] Here, "superpolynomial" refers to functions $f(n)$ that grow faster than any polynomial in n.
[7] We pronounce this "Songs".
[8] Such asymptotically good codes are known to exist (e.g, [8, 24]).

Unlike Theorem 2, the superpolynomial lower bound on the circuit complexity of noisy ground states does not require any complexity-theoretic assumption! The caveat is that this family is only known to exist via a counting argument; there is exactly *one* term of the Hamiltonian that has $\omega(\log n)$-locality and does not have an explicit description. However, however, all of other the terms of the local Hamiltonians are 7-local and have explicit descriptions.

The essential idea is to apply the proof of Theorem 2 to the $\mathsf{QMA}^{\mathcal{O}}$ verifier that decides a language L which is not in $\mathsf{QCMA}^{\mathcal{O}}$. In both [1, 20], this verifier only makes a single call to the oracle $\mathcal{O}$. Thus there is one term in the Feynman-Kitaev clock Hamiltonian corresponding to the propagation of that oracle call. Since we do not have an explicit description of a separating oracle $\mathcal{O}$, this Hamiltonian term is non-explicit.

2.3 Asymptotically good approximate low-weight check codes

The techniques from the previous sections also give rise to what we call *approximate quantum low-weight check (qLWC) codes*. These are closely related to quantum low-density parity check (qLDPC) codes, which are stabilizer codes where each parity check acts on a bounded number of particles, and each particle participates in a bounded number of parity checks. It is a long-standing open question of whether asymptotically good qLDPC codes exist (i.e. constant locality, constant rate, and constant relative distance). The qLDPC conjecture posits that such codes exist.

We show that if one relaxes the conditions of (a) each particle participating in a small number of constraints, and (b) that we can *exactly* recover from errors, we can obtain locally defined quantum error-correcting codes with such good parameters. First, we define our notion of approximate qLWC codes:

▶ **Definition 4** (Approximate qLWC code)**.** A local Hamiltonian $H = H_1 + \cdots + H_m$ acting on n dimension-q qudits is a $[[n, k, d]]_q$ *approximate quantum LWC code* with error δ and locality w iff each of the terms H_i act on at most w qudits and there exists encoding and decoding maps Enc, Dec such that

1. $\langle\Psi| H |\Psi\rangle = 0$ if and only if $|\Psi\rangle\langle\Psi| = \mathrm{Enc}(|\xi\rangle\langle\xi|)$ for some $|\xi\rangle \in (\mathbb{C}^q)^{\otimes k}$.
2. For all $|\phi\rangle \in (\mathbb{C}^q)^{\otimes k} \otimes \mathcal{R}$ where $\mathcal{R}$ is some purifying register, for all completely positive trace preserving maps $\mathcal{E}$ acting on at most $(d-1)/2$ qudits,

$$\|\mathrm{Dec} \circ \mathcal{E} \circ \mathrm{Enc}(|\phi\rangle\langle\phi|) - |\phi\rangle\langle\phi|\|_1 \leq \delta. \tag{2}$$

Here, the maps Enc, $\mathcal{E}$, and Dec do not act on register $\mathcal{R}$.

The first condition of the above definition enforces that the ground space of the Hamiltonian H of an approximate qLWC code is a q^k-dimensional codespace; it is the exactly the image of the encoding map Enc. The second condition corresponds to the approximate error-correcting condition, where we only require that the decoded state is *close* to the original state (i.e., we no longer insist that $\mathrm{Dec} \circ \mathcal{E} \circ \mathrm{Enc}$ is exactly the identity channel $\mathcal{I}$). Although there are few results on approximate quantum error-correcting codes, we do know that relaxing the exact decoding condition yields codes with properties that cannot be achieved using exact codes [31, 16, 11].

Our proof of Theorem 2 yields a construction of an approximate quantum LWC code with distance $\Omega(n)$.We believe this may be of independent interest.

▶ **Theorem 5** (Good approximate qLWC codes exist)**.** *For all error functions $\delta(n)$ there exist a family of $[[n,k,d]]_q$ approximate quantum LWC codes with the following parameters:*

$$\begin{aligned} \textit{Qudit dimension} \quad & q = O(1), \\ \textit{Error} \quad & \delta = \delta(n), \\ \textit{Locality} \quad & w = 3 + 2r, \\ \textit{Blocklength} \quad & n = O(rk), \\ \textit{Distance} \quad & d = \Omega(n/r) \end{aligned}$$

where

$$r = O\left(\frac{\log\left(1 + 4/\delta^2\right)}{\log n}\right) + 2. \tag{3}$$

Furthermore, the encoding and decoding maps for these codes are explicit and efficiently computable.

Observe that when $\delta(n) = 1/\mathsf{poly}(n)$, the parameter $r = O(1)$.

By comparison, the best-known qLDPC codes (of the stabilizer variety) with constant locality have distance bounded by $O(\sqrt{n \log n})$ [23]. Hastings constructs a qLDPC stabilizer code with constant locality that has distance $n^{1-\epsilon}$ for any $\epsilon > 0$, assuming a conjecture in high-dimensional geometry [26, 25]. Bacon, et al. were able to construct sparse subsystem codes (a generalization of stabilizer codes) with constant locality and distance $n^{1-o(1)}$ [9]. We note that, interestingly, the codes of [9] are constructed from fault-tolerant quantum circuits that implement a stabilizer code — this is similar to the way we construct our approximate qLWC codes!

2.4 Implications for NLTS, quantum PCP and quantum LDPC

Our investigation into noisy ground states and approximate low-weight check codes is motivated by a number of important open questions in quantum information theory: NLTS, quantum PCP, and quantum LDPC. We believe that our results help clarify the status of these open problems, and the relationships between them.

A separation between LNGS/SNGS and NLTS Hamiltonians.

First, our logarithmic circuit-depth lower bound for noisy ground states (Theorem 1) gives a strong separation between the notions of entanglement robustness in NLETS and NLTS: we showed that a one-dimensional local Hamiltonian is NLETS. However, it is easy to see that one-dimensional Hamiltonians (or any Hamiltonian on a constant-dimensional lattice) cannot be NLTS. To see this, consider taking a n-particle ground state $|\Psi\rangle$ of a 1D Hamiltonian H; divide up the n particles into contiguous chunks of length L. Let $\sigma = \rho_1 \otimes \rho_2 \otimes \cdots \otimes \rho_{n/L}$ where ρ_i is the reduced density matrix of $|\Psi\rangle$ on the i'th chunk. This state σ violates $O(n/L)$ terms of the Hamiltonian (since H is one-dimensional). Therefore, it is a ϵ-energy state of H for $L = \Theta(1/\epsilon)$. On the other hand, σ is a tensor product state that can be generated by $2^{O(1/\epsilon)}$-depth circuits, which is constant for constant ϵ. This indicates that the form of entanglement robustness as expressed in NLETS and in our LNGS/SNGS Hamiltonian constructions is much weaker than the entanglement robustness required by the NLTS conjecture and quantum PCP, where one has to look for Hamiltonians on high-dimensional geometries.

Quantum LDPC codes and the Quantum PCP conjecture.

Resolving the qPCP conjecture would likely involve a transformation from H to H' that (at the very least) has the property that exact ground states of H (or closeby states in trace distance) can be recovered from *low-energy* states of H'. It has been suggested that such a transformation would involve some kind of qLDPC code [22, 3, 26, 19]. In fact, it is believed that a special kind of qLDPC code, called a *quantum locally testable code (qLTC)*, is necessary [3]. However, the existence of qLTCs (or even qLDPC codes) with constant relative distance is a major open problem.

We believe our results on approximate quantum LWC codes present two take-home messages for the qPCP and qLDPC conjectures. First, it is important that a qPCP (or a qLTC) Hamiltonian be local, but it is *not* necessary that the Hamiltonian be bounded degree (meaning that each particle only participates in a few terms). The bounded degree condition is useful in the original context for qLDPCs, where an important motivation is to find fast decoding algorithms. In the context of qPCP/qLTC, however, decoding efficiency is not an immediate concern; thus resolving the qPCP conjecture *need not* resolve the qLDPC conjecture.

Second, we believe this gives evidence that considering codes other than stabilizer codes — such as approximate codes or subsystem codes — may be useful in the quest for both qPCP and qLDPC. Most work on qLDPC codes has focused on constructing CSS and stabilizer codes, but it may be fruitful to branch out beyond the CSS/stabilizer setting for the purposes of understanding the possibilities (or limits) of qPCP/qLDPC. For example, our qLWC codes are unconventional in a few ways: they are defined by non-commuting Hamiltonians, they only admit approximate recovery, and each particle participates in a large number of checks.

2.5 Open questions

We list a few open problems.

1. Are there SNGS Hamiltonians or (approximate) qLWC codes that are geometrically local (with respect to, say, the Euclidean metric)? Our construction of a one-dimensional NLGS Hamiltonian uses a simplification of a technique of Aharanov et. al. [5] of converting a quantum circuit into a two-dimensional local Hamiltonian. This technique works because of the specific structure of the circuit generating the $|\boxtimes\rangle$ state. In general, the transformation involves increasing the number of qudits by more than a constant factor. If this factor is $\Theta(n^\alpha)$, then the ground states are resilient to errors of size at most $n^{1-\alpha}$.
2. Is there a family of local Hamiltonians such that any *superposition* (not just convex combination) of low-error states have large circuit complexity? This notion is a generalization of a noisy state; such states have small *quantum Hamming distance* to the ground space. This is an interesting notion in the context of quantum locally testable codes (qLTCs) because low-energy states are equivalent to states with low quantum Hamming distance to the codespace (see [19] for definitions of quantum Hamming distance and qLTCs).
3. Are there applications of our qLWC constructions?
4. There has been a number of recent results about approximate quantum error-correcting codes in a variety of areas including many-body physics [12], the AdS/CFT correspondence [28], and quantum resource theories [27]. Could approximate error-correcting codes play a role in trying to resolve the qPCP and qLDPC conjectures?
5. Eldar and Harrow showed that quantum locally testable codes of the CSS type are NLTS [19]. Can this argument be extended to general qLTCs?

6. Is it possible for qLDPC codes (not necessarily stabilizer or exact error-correcting codes) to be defined as the codespace of a geometrically local Hamiltonian? There are a few no-go results that give limitations on codes living on lattices [13, 21], but they apply to special classes of codes such as stabilizer codes or locally-correctible codes. Our qLWC codes, by contrast, are neither.
7. Could the *combinatorial NLTS conjecture* be easier to prove than the NLTS conjecture? This conjecture posits that there exist a family of local Hamiltonians where states that have non-zero energy penalty on only a small constant fraction of Hamiltonian terms must have non-trivial circuit complexity.

References

1 Scott Aaronson and Greg Kuperberg. Quantum versus classical proofs and advice. In *Computational Complexity, 2007. CCC'07. Twenty-Second Annual IEEE Conference on*, pages 115–128. IEEE, 2007.

2 Dorit Aharonov, Itai Arad, and Thomas Vidick. Guest column: The quantum PCP conjecture. *SIGACT News*, 44(2):47–79, 2013. doi:10.1145/2491533.2491549.

3 Dorit Aharonov and Lior Eldar. Quantum locally testable codes. *SIAM Journal on Computing*, 44(5):1230–1262, 2015. doi:10.1137/140975498.

4 Dorit Aharonov and Tomer Naveh. Quantum NP - a survey, 2002. arXiv:arXiv:quant-ph/0210077.

5 Dorit Aharonov, Wim van Dam, Julia Kempe, Zeph Landau, Seth Lloyd, and Oded Regev. Adiabatic quantum computation is equivalent to standard quantum computation. *SIAM J. Comput.*, 37(1):166–194, 2007. doi:10.1137/S0097539705447323.

6 Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998. doi:10.1145/278298.278306.

7 Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of np. *J. ACM*, 45(1):70–122, 1998. doi:10.1145/273865.273901.

8 Alexei Ashikhmin, Simon Litsyn, and Michael A Tsfasman. Asymptotically good quantum codes. *Physical Review A*, 63(3):032311, 2001.

9 Dave Bacon, Steven T Flammia, Aram W Harrow, and Jonathan Shi. Sparse quantum codes from quantum circuits. *IEEE Transactions on Information Theory*, 63(4):2464–2479, 2017.

10 Johannes Bausch and Elizabeth Crosson. Analysis and limitations of modified circuit-to-hamiltonian constructions, 2016. arXiv:arXiv:1609.08571.

11 Cédric Bény and Ognyan Oreshkov. General conditions for approximate quantum error correction and near-optimal recovery channels. *Physical review letters*, 104(12):120501, 2010.

12 Fernando GSL Brandao, Elizabeth Crosson, M Burak Şahinoğlu, and John Bowen. Quantum error correcting codes in eigenstates of translation-invariant spin chains. *arXiv preprint arXiv:1710.04631*, 2017.

13 Sergey Bravyi, David Poulin, and Barbara Terhal. Tradeoffs for reliable quantum information storage in 2d systems. *Physical review letters*, 104(5):050503, 2010.

14 Libor Caha, Zeph Landau, and Daniel Nagaj. The Feynman-Kitaev computer's clock: bias, gaps, idling and pulse tuning. *arXiv preprint arXiv:1712.07395*, 2017.

15 Stephen A. Cook. The complexity of theorem-proving procedures. In *Proceedings of the Third Annual ACM Symposium on Theory of Computing*, STOC '71, pages 151–158, New York, NY, USA, 1971. ACM. doi:10.1145/800157.805047.

16 Claude Crépeau, Daniel Gottesman, and Adam Smith. Approximate quantum error-correcting codes and secret sharing schemes. In *Proceedings of the 24th Annual International Conference on Theory and Applications of Cryptographic Techniques*, EUROCRYPT'05, pages 285–301, Berlin, Heidelberg, 2005. Springer-Verlag. doi:10.1007/11426639_17.

17 Elizabeth Crosson and John F Bowen. Quantum ground state isoperimetric inequalities for the energy spectrum of local hamiltonians. *arXiv quant-ph*, 2017. arXiv:1703.10133.

18 Irit Dinur. The PCP theorem by gap amplification. *J. ACM*, 54(3), jun 2007. doi:10.1145/1236457.1236459.

19 Lior Eldar and Aram Wettroth Harrow. Local hamiltonians whose ground states are hard to approximate. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 427–438, 2017. doi:10.1109/FOCS.2017.46.

20 Bill Fefferman and Shelby Kimmel. Quantum vs classical proofs and subset verification. *CoRR*, abs/1510.06750, 2015. arXiv:1510.06750.

21 Steven T. Flammia, Jeongwan Haah, Michael J. Kastoryano, and Isaac H. Kim. Limits on the storage of quantum information in a volume of space. *Quantum*, 1:4, apr 2017. doi:10.22331/q-2017-04-25-4.

22 Michael H. Freedman and Matthew B. Hastings. Quantum systems on non-k-hyperfinite complexes: a generalization of classical statistical mechanics on expander graphs. *Quantum Information & Computation*, 14(1-2):144–180, 2014. URL: http://www.rintonpress.com/xxqic14/qic-14-12/0144-0180.pdf.

23 Michael H Freedman, David A Meyer, and Feng Luo. Z2-systolic freedom and quantum codes. *Mathematics of quantum computation, Chapman & Hall/CRC*, pages 287–320, 2002.

24 Daniel Eric. Gottesman. *Stabilizer codes and quantum error correction*. PhD thesis, California Institute of Technology, Pasadena, California, 1997.

25 Matthew B Hastings. Weight reduction for quantum codes. *arXiv preprint arXiv:1611.03790*, 2016.

26 Matthew B. Hastings. Quantum codes from high-dimensional manifolds. In *8th Innovations in Theoretical Computer Science Conference, ITCS 2017, January 9-11, 2017, Berkeley, CA, USA*, pages 25:1–25:26, 2017. doi:10.4230/LIPIcs.ITCS.2017.25.

27 Patrick Hayden and Geoffrey Penington. Approximate quantum error correction revisited: Introducing the alphabit. *arXiv preprint arXiv:1706.09434*, 2017.

28 Isaac H Kim and Michael J Kastoryano. Entanglement renormalization, quantum error correction, and bulk causality. *Journal of High Energy Physics*, 2017(4):40, 2017.

29 A.Y. Kitaev, A. Shen, and M.N. Vyalyi. *Classical and Quantum Computation*. Graduate studies in mathematics. American Mathematical Society, 2002.

30 A.Yu. Kitaev. Fault-tolerant quantum computation by anyons. *Annals of Physics*, 303(1):2–30, 2003. doi:10.1016/S0003-4916(02)00018-0.

31 Debbie W Leung, Michael A Nielsen, Isaac L Chuang, and Yoshihisa Yamamoto. Approximate quantum error correction can lead to better codes. *Physical Review A*, 56(4):2567, 1997.

32 L. A. Levin. Universal sequential search problems. *Problems of Information Transmission*, 9(3):265–266, 1973.

33 Jean-Pierre Tillich and Gilles Zémor. Quantum ldpc codes with positive rate and minimum distance proportional to the square root of the blocklength. *IEEE Transactions on Information Theory*, 60(2):1193–1202, 2014.