

Fermionic Hamiltonians without trivial low-energy states

Yaroslav Herasymenko^{*,1,2,3}, Anurag Anshu⁴, Barbara M. Terhal^{1,2}, and Jonas Helsen³

¹QuTech, TU Delft, P.O. Box 5046, 2600 GA Delft, The Netherlands

²Delft Institute of Applied Mathematics, TU Delft, 2628 CD Delft, The Netherlands

³QuSoft and CWI, Science Park 123, 1098 XG Amsterdam, The Netherlands

⁴School of Engineering and Applied Sciences, Harvard University, 150 Western Ave., Allston, MA 02134, US

July 27, 2023

Abstract

We construct local fermionic Hamiltonians with no low-energy trivial states (NLTS), providing a fermionic counterpart to the NLTS theorem. Distinctly from the qubit case, we define trivial states via finite-depth *fermionic* quantum circuits. We furthermore allow free access to Gaussian fermionic operations, provided they involve at most $O(n)$ ancillary fermions. The desired fermionic Hamiltonian can be constructed using any qubit Hamiltonian which itself has the NLTS property via well-spread distributions over bitstrings, such as the construction in [ABN22]. We define a fermionic analogue of the class quantum PCP and discuss its relation with the qubit version.

1 Introduction

One of the main problems in quantum complexity is establishing the hardness of quantum computational problems such as the estimation of the ground state energy of a many-body Hamiltonian. If high accuracy is required, –scaling inversely polynomial with the problem size n –, this problem is known to be QMA-complete [KSV02]. It has been an open question whether the estimation of the ground state energy *density* of a Hamiltonian with constant error, would also be QMA-hard [AAV13]. Such a result would mirror the classical PCP theorem, which states that there are classical SAT problems for which it is NP-complete to decide whether all clauses can be satisfied or a constant fraction of the clauses can never be satisfied.

Motivated by this quantum PCP problem, it was conjectured in [FH13] that there exist local sparse qubit Hamiltonians such that any state with sufficiently small energy density cannot be prepared by a constant-depth circuit, hence suggesting that low-energy states have some nontrivial complexity. This “no low-energy trivial state” (NLTS) conjecture was recently proven in [ABN22]. They demonstrate that a family of Hamiltonians based on recently established good quantum LDPC codes [LZ22] has the desired property. While the low energy states of this family of Hamiltonians are non-trivial, their ground state can be efficiently prepared by a Clifford circuit. This means that these Hamiltonians fall short of being genuinely ‘complex’. Hence it is natural to extend the definition of “non-trivial” beyond constant-depth circuits to climb up the complexity ladder, see [GG22, WFC23, CCNN23, AGK23]. In particular, [CCNN23] have shown that an

*Corresponding author: yaroslav@cw.nl

adaptation of the construction in [ABN22] can be used to construct Hamiltonians without either stabilizer low-energy states or constant-depth low-energy states.

Here we similarly extend the construction in [ABN22], but now in the direction of fermionic Hamiltonians. Fermionic Hamiltonians describe the majority of quantum systems of practical interest, namely interacting electrons in materials and chemical compounds. These systems are among the prime targets for quantum simulation using both classical and quantum computers. Fermionic quantum complexity is interesting when it differs from qubit complexity and when qubit arguments are not directly applicable. In addition, while most attention in the quantum complexity literature is devoted to qubit systems, fermions have come more into focus in recent years, see e.g. [BG17, BGKT19, ODMZ22, HO22, HSHT22].

In the present manuscript, we answer the question posed in [HO22] in the affirmative: do there exist fermionic Hamiltonians without low-energy trivial states? Our results directly build on [ABN22]. We give natural notions of triviality through the class of fermionic Gaussian states, which are classically simulatable. Besides our main result, the fermionic NLTS Theorem (Theorem 7), we analyze the definition of the quantum PCP class from a fermionic perspective, presenting a fermionic version (Section 5). The paper is otherwise organized as follows. In Section 2 we provide some useful definitions, in particular we give two different natural notions of trivial fermionic circuits. In Section 3 we introduce the key tools from [ABN22] and in Section 4 we prove our main Theorem. We end the paper with a Discussion (Section 7) on sparse fermion-to-qubit mappings and what lies beyond.

2 Preliminaries and Definitions

A system of n fermionic modes can be described by $2n$ Majorana fermion operators c_i ($i \in [2n]$) which obey $c_i^2 = \mathbb{I}$, $\forall i \neq j$ $c_i c_j = -c_j c_i$, $c_i = c_i^\dagger$, $\text{Tr}(c_i) = 0$. Any fermionic state on this system is given by a Hermitian positive semidefinite $\rho \geq 0$ with $\text{Tr}(\rho) = 1$ which can be expressed as an *even* polynomial in the Majorana fermion operators.

We define Hermitian operators C_K as ordered products of the operators c_k :

$$C_K = i^{|K|(|K|-1)/2} c_{k_1} c_{k_2} \dots c_{k_{|K|}}, \quad (1)$$

with $K = (k_1 < k_2 < \dots < k_{|K|})$. The class of Gaussian fermionic states is a subclass of fermionic states which are efficiently describable. Gaussian circuits preserve this class of states, and hence are efficiently-classically simulatable. For more background on Gaussian states and Gaussian circuits, –also called fermionic linear optics–, we refer the reader to e.g. [BK02, TD02, Bra05, HO22, HSHT22].

Consider preparing a fermionic state ρ on n modes using m ancillary modes (with Majorana operators c_k , $k \in [2(n+m)] \setminus [2n]$). Starting from a *pure* initial Gaussian state σ_G , one may use a fermionic circuit U which is generally decomposable into gates which use quartic or quadratic interactions between $O(1)$ fermionic operators [BK02], i.e.

$$\rho = \text{Tr}_{[2(n+m)] \setminus [2n]} \left(U \sigma_G U^\dagger \right). \quad (2)$$

We now introduce two slightly different notions of nontrivial fermionic circuit ‘depth’. Our definition will make use of the notion of Gaussian circuits, namely circuits consisting of unitaries of the form $\exp(-\omega c_{k_1} c_{k_2})$. We note that any Gaussian unitary on n fermionic modes can be written as a $\Theta(n)$ -deep circuit [JSK⁺18] of such unitaries. The definition is as follows:

Definition 1. We say that a fermionic circuit U has depth T , if it is given by

$$U = W_T W_{T-1} \dots W_1, \quad (3)$$

with each W_t of the form

$$W_t = \prod_{K: |K| \in \{2,4\}} \exp\left(i\omega_K^{(t)} C_K\right) \quad (4)$$

where the product runs over non-overlapping sets $K \subset [2n]$. Furthermore we say that U is a depth T circuit with free access to Gaussian operations, if it is given by

$$U = G_T W_T G_{T-1} \dots W_2 G_1 W_1, \quad (5)$$

where G_t are arbitrary Gaussian circuits.

Note we take a fixed choice of operators c_s in Eq. (4). Allowing the use of an arbitrary basis of fermionic operators $\{c_k\}$ in each layer of W_t would have allowed more expressive unitaries U . However, we don't consider this model separately, since it is subsumed by free access to Gaussian operations in Eq. (5).

Here is another useful (standard) definition:

Definition 2 (Local sparse n -qubit (resp. n -fermion) Hamiltonian). $H = \sum_{i=1}^m H_i$ is a local sparse n -qubit (resp. n -fermion) Hamiltonian when the maximum number of Pauli operators (resp. Majorana operators) in each term H_i is $O(1)$ (local) and the maximum number of terms involving any qubit (resp. Majorana operator) is $O(1)$. The smallest eigenvalue of H is denoted as $\lambda(H)$.

Throughout this manuscript, we will be considering Hamiltonians which are sparse, local, positive-semidefinite and have terms bounded by one in operator norm unless explicitly said otherwise.

3 NLTS Hamiltonians and well-spread distributions

It was proven in [ABN22] that there exist qubit Hamiltonians with the following NLTS property

Definition 3. A local sparse n -qubit Hamiltonian $H \geq 0$ has the NLTS property with parameter ϵ , if any family of n -qubit states ρ_n with energy $\text{Tr}(\rho_n H) < \epsilon n$ requires a quantum circuit, which uses an arbitrary number of ancilla qubits, of depth at least $T = \Omega(\log n)$.

A way to obtain NLTS Hamiltonians passes through the notion of well-spread quantum states, which are states that, when measured in a ‘‘trivial’’ basis, have certain statistical properties. We now define well-spreadness both for qubits and fermions; the fermionic definition will play a key role in our result.

Definition 4. Consider an n -qubit POVM $\pi_R(s)$, defined as

$$\pi_R(s) = R^\dagger \prod_{j=1}^n \frac{1 + (-1)^{s_j} Z_j}{2} R, \quad (6)$$

where s is the outcome bitstring $s = (s_1, \dots, s_n) \in \{0, 1\}^n$ and R is a tensor product of single qubit unitaries. Given an n -qubit state ρ , define a probability distribution over $s \in \{0, 1\}^n$,

$$p_{R,\rho}(s) = \text{Tr}(\pi_R(s)\rho). \quad (7)$$

Definition 5. Consider a POVM with $2n$ Majorana fermion operators

$$\pi_G(s) = G^\dagger \prod_{j=1}^n \frac{\mathbb{I} + (-1)^{s_j} c_{2j-1} c_{2j}}{2} G, \quad (8)$$

where G is a Gaussian unitary and s is the outcome bitstring. A state ρ on n fermionic modes yields a probability distribution over $s \in \{0, 1\}^n$

$$p_{G,\rho}(s) = \text{Tr}(\pi_G(s)\rho). \quad (9)$$

Definition 6. A qubit (resp. fermionic) state ρ is said to be (μ, L) -spread if exists a unitary R (respectively, G) as defined above, and two sets of bitstrings S_1 and $S_2 \subset \{0, 1\}^n$ such that $p(s) \equiv p_{R,\rho}(s)$ (respectively, $p_{G,\rho}(s)$) obeys

$$\sum_{s \in S_1} p(s) \geq \mu, \quad \sum_{s \in S_2} p(s) \geq \mu, \quad \text{and} \quad \min_{s_1, s_2 \in S_{1,2}} |s_1 - s_2| = L. \quad (10)$$

with Hamming distance $|\cdot|$. An n -qubit (n -fermion) Hamiltonian H is referred to as having well-spread low-energy states, if there exists a constant ϵ such that any ρ with energy $\text{tr}(\rho H) < \epsilon n$ is (μ, L) -spread for $\mu = \Omega(1)$ and $L = \Omega(n)$.

Ref. [ABN22] gives a quantum LDPC (CSS) code construction of an n -qubit Hamiltonian H with well-spread low-energy states. This property implies that H is NLTS (Fact 4 in [ABN22]).

4 Fermionic NLTS Hamiltonians

We show that qubit Hamiltonians with well-spread low energy states (by Definition 6) can be used to construct fermionic local Hamiltonians without trivial low-energy states. Terminology related to the fermionic circuits was introduced in Definition 1.

Theorem 7. For even n , consider an n -qubit $H^{(q)}$ with well-spread states below energy ϵn . Construct a $3n/2$ -fermion Hamiltonian H with Majorana operators $\{c_{x,j}, c_{y,j}, c_{z,j}\}_{j \in [n]}$, by replacing the Pauli operators in $H^{(q)}$ using

$$X_j \mapsto ic_{y,j}c_{z,j}, \quad Y_j \mapsto ic_{x,j}c_{z,j}, \quad Z_j \mapsto ic_{x,j}c_{y,j}. \quad (11)$$

The resulting Hamiltonian H has two fermionic NLTS properties. In particular, for an arbitrary $3n/2$ -fermion state ρ such that $\text{Tr}(\rho H) < \epsilon n$, the following holds

1. Using arbitrary Gaussian initialization and an arbitrary number of ancillas, any fermionic circuit that prepares ρ has depth $T = \Omega(\log(n))$
2. Using a fermionic circuit with free access to Gaussian operations and at most $m = O(n)$ ancillary fermionic modes, any circuit that prepares ρ has depth $T = \Omega(\log n)$.

Before proceeding to the proof, we would like to briefly discuss the above result. The restriction on the access to ancillary modes given access to Gaussian operations, which we see in Theorem 7.2, may be of fundamental origin. On the one hand, ancillary non-Gaussian states in combination with Gaussian operations allow the injection of non-Gaussian gates [Bra06]. On the other hand, this procedure also requires a capacity for adaptive measurements, which is not included in our scenario. We leave as an open question whether the particular upper bound $m = O(n)$ is optimal, or if it can be improved using other proof techniques.

Our constructed Hamiltonian H has the same spectrum as $H^{(q)}$, with an additional $2^{n/2}$ -fold degeneracy. As will be clear from the proof, the degeneracy does not affect the well-spreadness property. The key point is that the mapping we use preserves the locality of H (unlike, say, a direct mapping of n qubits onto n fermions via the inverse Jordan-Wigner transformation). We believe that a similar result can be achieved with the local embedding of n qubits into $4n$ Majorana operators, as introduced by [Kit06], or a mapping used in [LCV07].

Proof. The Hamiltonian H constructed using Eqs. 11 can be understood as an image of $H^{(q)} \otimes \mathbb{I}^{(f)}$ under the unitary ‘‘qubit assimilation’’ mapping (Lemma 15 in Appendix A). Here $\mathbb{I}^{(f)}$ acts on auxiliary $n/2$ fermionic modes.

We now show that ρ is $(\Omega(1), \Omega(n))$ -spread. Then we use this fact to lower-bound the circuit complexity of ρ .

For the state ρ , construct $\rho^{(\text{qf})}$ on n qubits and $n/2$ fermions by applying inverse qubit assimilation to ρ (replacements listed in Eqs. 35). Next, trace out the $n/2$ -fermion sector from $\rho^{(\text{qf})}$ and thus obtain an n -qubit state $\rho^{(\text{q})}$. Observe

$$\text{Tr}(\rho^{(\text{q})} H^{(\text{q})}) = \text{Tr}(\rho^{(\text{qf})} H^{(\text{q})} \otimes \mathbb{I}^{(\text{f})}) = \text{Tr}(\rho H) < \epsilon n. \quad (12)$$

By definition of $H^{(\text{q})}$, this implies that $\rho^{(\text{q})}$ is $(\Omega(1), \Omega(n))$ -spread. By Definition 6, there exists a product of single-qubit rotations R , a POVM $\{\pi_R(s)\}_{s \in \{0,1\}^n}$, and two sets of bitstrings $S_{1,2} \subset \{0,1\}^n$ such that $p_R(s) = \text{Tr}(\rho^{(\text{q})} \pi_R(s))$ obeys

$$\sum_{s \in S_{1,2}} p_R(s) = \Omega(1), \quad \min_{s_1, s_2 \in S_{1,2}} |s_1 - s_2| = \Omega(n). \quad (13)$$

Considering a general R rotating Z_j to $X_j \sin \theta_j \cos \phi_j + \sin \theta_j \sin \phi_j Y_j + \cos \theta_j Z_j$ for each qubit, we write it as

$$R = \prod_{j=1}^n \exp\left(\frac{i\theta_j}{2} (\sin \phi_j X_j - \cos \phi_j Y_j)\right). \quad (14)$$

By qubit assimilation (Lemma 15), X_j , Y_j , and Z_j are mapped respectively onto $ic_{j,z}c_{j,y}$, $ic_{j,x}c_{j,z}$, and $ic_{j,x}c_{j,y}$ for $j \in [n]$. Therefore R is mapped onto

$$G = \prod_j \exp\left(\frac{\theta_j}{2} (-c_{j,z}c_{j,y} \sin \phi_j + c_{j,x}c_{j,z} \cos \phi_j)\right). \quad (15)$$

Relabeling $c_{j,y} \rightarrow c_{2j-1}$, $c_{j,x} \rightarrow c_{2j}$ and $c_{j,z} \rightarrow c_{2n+j}$, define a POVM

$$\pi_G(\tilde{s}) = G^\dagger \left(\prod_{j=1}^{3n/2} \frac{\mathbb{I} + (-1)^{\tilde{s}_j} ic_{2j-1}c_{2j}}{2} \right) G, \quad (16)$$

where $\tilde{s} \in \{0,1\}^{3n/2}$. We adopted the notation $\pi_G(\tilde{s})$ from Definition 5, since the unitary G is Gaussian. We see that qubit assimilation maps $\pi_R(s)$ onto:

$$\begin{aligned} G^\dagger \left(\prod_{j=1}^n \frac{\mathbb{I} + (-1)^{s_j} ic_{2j-1}c_{2j}}{2} \right) G &= \sum_{\substack{\tilde{s}_j = \pm 1, \\ j \in (n+1, \dots, 3n/2)}} G^\dagger \left(\prod_{j=1}^{3n/2} \frac{\mathbb{I} + (-1)^{\tilde{s}_j} ic_{2j-1}c_{2j}}{2} \right) G \Big|_{\tilde{s}_j = s_j, j \in [n]} \\ &= \sum_{\substack{\tilde{s}_j = \pm 1, \\ j \in (n+1, \dots, 3n/2)}} \pi_G(\tilde{s}) \Big|_{\tilde{s}_j = s_j, j \in [n]}. \end{aligned} \quad (17)$$

Note that in Eq. (17) bits s_j for $j \in (1, \dots, n)$ came from the pre-image $\pi_R(s)$, while s_j for $j \in (n+1, \dots, 3n/2)$ are the newly introduced dummy variables. Applying the constructed POVM $\pi_G(s)$, we find that the $3n/2$ -fermion state ρ is $(\Omega(1), \Omega(n))$ by Definition 6. Indeed, using $S_{1,2}$ from Eq. (13) we can directly construct two sets of bitstrings $\tilde{S}_{1,2} \equiv \{(s_1, \dots, s_n, s'_1, \dots, s'_{n/2}) \mid s \in S_{1,2}, s' \in \{0,1\}^{n/2}\} \subset \{0,1\}^{3n/2}$, such that $p_G(s) = \text{Tr}(\rho \pi_G(s))$ obeys

$$\sum_{\tilde{s} \in \tilde{S}_{1,2}} p_G(\tilde{s}) = \sum_{s \in S_{1,2}} p_R(s) = \Omega(1), \quad \min_{\tilde{s}_1, \tilde{s}_2 \in \tilde{S}_{1,2}} |\tilde{s}_1 - \tilde{s}_2| = \min_{s_1, s_2 \in S_{1,2}} |s_1 - s_2| = \Omega(n). \quad (18)$$

Consider constructing ρ with a depth- T fermion circuit with free access to Gaussian operations, using $m = O(n)$ ancillary modes. To prove that $T = \Omega(\log n)$ depth is required, we employ $(\Omega(1), \Omega(n))$ -spreadness of ρ and Lemma 8 stated below. This directly yields $T = \Omega(\log n)$.

Finally, consider producing ρ in a depth T fermionic circuit U (see Eqs. (2, 3)), starting from an arbitrary Gaussian initialization σ_G . Since the output state ρ is obtained as a marginal of $U\sigma_G U^\dagger$ on Majorana operators $[3n]$, we are not concerned with the action of U on the entire system of Majorana operators $[3n + 2m]$. Without changing ρ , U can be replaced with U' : its backward light cone stemming from $[3n]$. This backward light cone will be supported on at most $3n \cdot 4^T$ Majoranas, because each gate in W_t (Eq. (4)) involves at most three Majoranas in addition to each one from the light cone at layer $t + 1$. Naively, an arbitrary number $2m$ of ancillary Majoranas can still be non-trivially involved in the initialization state σ_G . However, one can replace this state with its marginal ρ_G on $3n \cdot 4^T$ Majoranas supporting U' , without changing ρ . In turn, being a mixed Gaussian state on $3n \cdot 4^T$ Majoranas, ρ_G can be purified back onto $O(n \cdot 4^T)$ fermions. Therefore, any $3n/2$ -fermionic state prepared using a depth- T fermionic circuit, arbitrary Gaussian initialization, and arbitrary m ancillary fermionic modes, can be prepared in the same setting with $3n/2 + m$ reduced to $O(n \cdot 4^T)$. Employing a next Lemma 8 and $(\Omega(1), \Omega(n))$ -spreadness of ρ once again, we obtain $T = \Omega(\log n)$, finalizing the proof of the Theorem. \square

Lemma 8. *Consider a (μ, L) -spread l -fermion state ρ . Any fermionic circuit using m ancillary fermionic modes and free access to Gaussian operations that produces ρ must be of depth at least*

$$T = \frac{1}{2} \log_3 \left(\frac{L^2}{1600(l+m) \log(1/\mu)} \right). \quad (19)$$

Proof. The following parallels the proof of Fact 4 from [ABN22]. The key additional observation is that Gaussian circuits, however deep, do not change the locality of fermionic operators. Consider preparing ρ as

$$\rho = \text{Tr}_{[2(m+l)] \setminus [2l]} \left(U \sigma_G U^\dagger \right), \quad (20)$$

with a fermionic circuit U and the pure Gaussian starting state

$$\sigma_G = \frac{1}{2^{l+m}} \prod_{k=1}^{l+m} (\mathbb{I} + i c_{2j-1} c_{2j}). \quad (21)$$

Consider the operator $Q = \frac{1}{(l+m)} \sum_{j=1}^{l+m} \frac{1}{2} (\mathbb{I} - i c_{2j-1} c_{2j})$. The spectrum of Q is $(0, 1/(l+m), 2/(l+m), \dots, 1)$, with the non-degenerate ground state ρ_G . The Hamiltonian UQU^\dagger has the same spectrum as Q and has $U\sigma_G U^\dagger$ as the ground state; we now prove that this Hamiltonian is also 23^T -local, where T is the depth of U . Indeed, (i) any Gaussian circuit G_t in U (cf. Eq. (5)) does not change the locality of a fermionic operator $C_{K'}$ in UQU^\dagger , while (ii) every layer W_t of a non-Gaussian circuit can increase the weight of $C_{K'}$ at most by a factor 3. To show point (i), consider the two options for a transformation of any Majorana monomial $C_{K'}$ with an elementary Gaussian operation:

$$e^{-i w_{k_1, k_2} c_{k_1} c_{k_2}} C_{K'} e^{i w_{k_1, k_2} c_{k_1} c_{k_2}} = \begin{cases} C_{K'} (\cos(2w_{k_1, k_2}) + i \sin(2w_{k_1, k_2}) c_{k_1} c_{k_2}) & \text{if } \{C_{K'}, c_{k_1} c_{k_2}\} = 0, \\ C_{K'} & \text{if } [C_{K'}, c_{k_1} c_{k_2}] = 0. \end{cases} \quad (22)$$

In the first option, k_1 or k_2 must belong to K' and therefore $C_{K'} c_{k_1} c_{k_2}$ has the same weight as $C_{K'}$. The second option trivially conserves locality. To show point (ii), consider the transformation similar to Eq. (22), using an elementary gate $\exp(i w_K C_K)$ with $|K| = 4$; this transformation can only add weight 2 to $C_{K'}$,

and only if $\{C_{K'}, C_K\} = 0$. Since W_t only contains non-overlapping generators C_K , there will be at most $|K'|$ generators such that $\{C_{K'}, C_K\} = 0$. Therefore, a transformation of $C_{K'}$ with W_t will produce operators with a weight of at most $|K'| + |K'| \cdot 2 = 3|K'|$.

We now use a matrix polynomial approximation $P(UQU^\dagger)$ to the ground state of degree f (concretely, the polynomial construction in [KLS96, BCDZ99, AAG22]) such that

$$\| |\psi\rangle\langle\psi| - P(UQU^\dagger) \|_\infty \leq \exp\left(-\frac{f^2}{100(l+m)}\right). \quad (23)$$

As UQU^\dagger is $2 \cdot 3^T$ -local, $P(UQU^\dagger)$ is $2f3^T$ -local. Fixing $f = \frac{L}{4 \cdot 3^T}$, we obtain $P(UQU^\dagger)$ that is $L/2$ -local. The right-hand-side of Eq. (23) then becomes $\exp(-\frac{L^2}{1600(l+m) \cdot 3^{2T}})$.

Next, from ρ being $(\Omega(1), \Omega(n))$ -spread there exists a Gaussian G and $S_1, S_2 \subset \{0, 1\}^l$ with $\min_{s_1, s_2 \in S_{1,2}} |s_1 - s_2| \geq L$ such that

$$\left\| \sum_{s_1, s_2 \in S_{1,2}} \pi_G(s_1) |\psi\rangle\langle\psi| \pi_G(s_2) \right\|_\infty \geq \mu. \quad (24)$$

On the other hand, from the $L/2$ -locality of $P(UQU^\dagger)$ it follows that

$$\left\| \sum_{s_1, s_2 \in S_{1,2}} \pi_G(s_1) P(UQU^\dagger) \pi_G(s_2) \right\|_\infty = 0, \quad (25)$$

since $\pi_G(s_1) C \pi_G(s_2) = 0$ for any $|s_1 - s_2| \leq L$ and $L/2$ -local C . The latter condition is true since the Gaussian transformation G conserves the locality of operator C , and because this condition manifestly holds for the case $G = \mathbb{I}$. Collecting Eqs. (23)-(25), we arrive at:

$$\exp\left(-\frac{L^2}{1600(l+m) \cdot 3^{2T}}\right) \geq \mu, \quad (26)$$

which amounts to the claimed lower bound on T . □

5 Qubit PCP versus Fermionic PCP

The quantum PCP complexity class (QPCP) was defined in e.g. [AALV09]. The quantum PCP conjecture [AAV13] asks whether $\text{QPCP}[q = O(1)]$ is equal to QMA, where $q = O(1)$ is the number of qubits of the proof which are checked. Here we introduce a new class $\text{FermPCP}[q = O(1)]$ which accesses q fermionic modes of a (fermionic) proof and argue that this class could be larger than $\text{QPCP}[q = O(1)]$ due to the limitations of fermion-to-qubit mappings.

In the definitions below we do not make explicit reference to the number of random bits of the verifier: it is sufficient if this is at most $\text{poly}(\log n)$ when the number of qubits (or fermionic modes) that the verifier accesses is $q = O(1)$ or $q = O(\log n)$ which is what we will use in this section. In addition, we allow the verifier to simulate randomness in the quantum circuit by having $(\text{poly}(n))$ ancillary qubits. We start by reproducing the definition of QPCP:

Definition 9 (QPCP). *A promise problem $L = L_{\text{yes}} \cup L_{\text{no}} \in \text{QPCP}[q]$ if there exists a quantum polynomial verifier and a polynomial $p(\cdot)$ with the following properties. The verifier receives as input a classical string x and a $p(x)$ -qubit density matrix ρ . The verifier randomly picks q qubits out of $p(x)$ according to some*

scheme using her random bits, and then runs a polynomial-sized quantum circuit $V(x, \rho)$, which can use $\text{poly}(n)$ ancillary qubits, and the circuit V accesses only the chosen q qubits as inputs. The accept/reject output of V is obtained by a measurement in the Z -basis of one of the ancilla qubits. Then $L \in \text{QPCP}[q]$ when,

- If $x \in L$, $\exists \rho$ such that $\text{Prob}(V(x, \rho) \text{ accepts}) \geq 2/3$.
- If $x \notin L$, $\forall \rho$, $\text{Prob}(V(x, \rho) \text{ accepts}) \leq 1/3$.

Here is the definition of the analogous fermionic class:

Definition 10 (FermPCP). A promise problem $L = L_{\text{yes}} \cup L_{\text{no}} \in \text{FermPCP}[q]$ if there exists a quantum polynomial fermionic verifier and a polynomial $p(\cdot)$ with the following properties. The verifier receives as input a classical string x and a $p(x)$ -fermionic state ρ_f . The verifier randomly picks q fermionic modes (i.e., $2q$ operators) out of $p(x)$ according to some scheme using her random bits, and then runs a polynomial-sized quantum circuit $V(x, \rho_f)$, possibly using $\text{poly}(n)$ ancillary qubits, and the circuit V accesses only the chosen q fermionic modes as inputs. The accept/reject output of V is obtained by a measurement in the Z -basis of one of the ancilla qubits. Then $L \in \text{QPCP}[q]$, when

- If $x \in L_{\text{yes}}$, $\exists \rho_f$ such that $\text{Prob}(V(x, \rho_f) \text{ accepts}) \geq 2/3$.
- If $x \in L_{\text{no}}$, $\forall \rho_f$, $\text{Prob}(V(x, \rho_f) \text{ accepts}) \leq 1/3$.

We note that the fermionic verifier and its access to a fermionic proof could be mapped to a qubit verifier and its proof by a standard fermion-to-qubit mapping such as the Jordan-Wigner transformation, but such a mapping clearly does not preserve the limited access structure of QPCP. Before we prove two statements about the relations between these classes, let us review what problem is complete for these classes and define the local Hamiltonian density problem:

Definition 11 ((Fermionic) Local Hamiltonian Density Problem (LHD)). Let $H = \frac{1}{m} \sum_{i=1}^m H_i$ be a local n -qubit (resp. n -fermion) Hamiltonian with $H_i \geq 0$, $\|H_i\| = O(1)$. Either $\lambda(H) \leq a$ or $\lambda(H) \geq b$, for constants $a > 0, b > 0$ and a constant $\epsilon = b - a > 0$. The qubit (resp. fermionic) local Hamiltonian density (LHD) problem is to decide which is the case.

It is known that $\text{LHD} \in \text{QPCP}[q = O(1)]$ by the proof of [KSV02] and identical arguments can be made to show that fermionic $\text{LHD} \in \text{FermPCP}[q = O(1)]$. The idea is that the expectation value of any (fermionic) Hamiltonian term H_i can be estimated using qubit ancillas. Here, one considers a spectral decomposition of $H_i = \sum_j \omega_{j,i} \sigma_{j,i}$ with $1 \geq \omega_{j,i} \geq 0$, and one defines a unitary which uses an additional ancilla qubit in the state $|0\rangle$ and maps

$$\sigma_{j,i} \otimes |0\rangle \langle 0| \rightarrow \sigma_{j,i} \otimes (\sqrt{\omega_{j,i}} |0\rangle + \sqrt{1 - \omega_{j,i}} |1\rangle) (\sqrt{\omega_{j,i}} \langle 0| + \sqrt{1 - \omega_{j,i}} \langle 1|), \quad (27)$$

(see [KSV02]). The desired expectation can be estimated by estimating the probability of obtaining the outcome ‘1’ when the qubit is measured.

The opposite, namely that qubit LHD is hard (and thus complete) for $\text{QPCP}[O(1)]$, has been informally proven in [AAV13]. For completeness, we prove it here and we extend it to a fermionic version:

Proposition 12. The qubit LHD problem is $\text{QPCP}[O(1)]$ -hard by a polynomial-time quantum reduction. Similarly, the fermionic LHD problem is $\text{FermPCP}[O(1)]$ -hard by a polynomial-time quantum reduction.

Proof. We show that any QPCP[$O(1)$] proof system can be mapped onto a LHD problem. The verifier uses $O(\log n)$ bits to draw from some probability distribution p_i and each i corresponds to picking a certain subset of q qubits and the number of such sets is $m = \text{poly}(n)$. Let the acceptance qubit be labelled ancilla qubit number 1. The probability for acceptance $p_{\text{accept}} = \sum_i p_i p_{\text{accept}}^i$ with $p_{\text{accept}}^i = \frac{1}{2}(1 + \text{Tr}(O_i \rho))$ where O_i is a traceless q -qubit observable acting only non-trivially on the chosen q qubits of ρ . Then we construct H as $H = \frac{1}{m} \sum_i H_i$ with $H_i = \frac{1}{2} p_i (\mathbb{I} - O_i)$, obeying $\|H_i\| \leq 1$ and $H_i \geq 0$. Hence, if $x \in L_{\text{yes}}$, there exists a ρ such that $\lambda(H) \leq \frac{1}{3}$. When $x \in L_{\text{no}}$, we have $\forall \rho, \lambda(H) \geq \frac{2}{3}$. To construct O_i , one needs to run the verifier's quantum circuit where we replace the $q = O(1)$ input qubits by all possible eigenstates of the Pauli operators on q qubits and apply process tomography on the superoperator $\mathcal{S}_i(\rho) = \text{Tr}_{\text{anc. but 1}} V_i \rho \otimes |00 \dots 0\rangle \langle 00 \dots 0| V_i^\dagger$, where $|00 \dots 0\rangle$ is the initial state of all ancillary qubits and V_i is the circuit V using the Pauli qubit operators in the chosen subset i . Using process tomography one constructs the q to $q + 1$ -qubit TPCP map $\mathcal{S}_i(\cdot)$ ¹, and thus $\mathcal{S}_i^\dagger(\cdot)$ (with negligible error $1/\text{poly}(n)$), and applies it to Z_1 to construct O_i , i.e. $O_i = \mathcal{S}_i^\dagger(Z_1)$.

In the fermionic case, all goes through similarly. Let ancilla qubit 1 be the acceptance qubit. The verifier measures $Z_1 = \pm 1$ or $p_{\text{accept}}^i = \frac{1}{2}(1 + \text{Tr}(Z_1 V_i \rho \otimes |00 \dots 0\rangle \langle 00 \dots 0| V_i^\dagger)) = \frac{1}{2}(1 + \text{Tr}(Z_1 \mathcal{S}_i(\rho))) = \frac{1}{2}(1 + \text{Tr}(O_i \rho))$. Here V_i only uses those fermionic modes in the chosen subset i and the state $|00 \dots 0\rangle$ of all ancillary qubits. To construct \mathcal{S}_i which maps q fermion modes to q fermions and 1 ancilla qubit, we can construct its corresponding Choi-Jamiolkowski state [Bra05]. Let $c_{i[1]}, \dots, c_{i[2q]}$ be the chosen set i of Majorana operators and take an additional set of reference operators $c_{r,1}, \dots, c_{r,2q}$, then $\rho_{\text{Choi}} = \mathcal{S}_i(\frac{1}{2^{2q}} \prod_{j=1}^{2q} (\mathbb{I} + i c_{i[j]} c_{r,j}))$. From ρ_{Choi} , a set of $O(1)$ Kraus operators $A_{k,i}$, –which can be expressed as even polynomials in the Majorana operators $c_{i[1]}, \dots, c_{i[2q]}$ and single ancilla qubit Paulis–, can be obtained [Bra05], and hence O_i can be constructed and thus H can be constructed via a quantum reduction. \square

We note that the qubit LHD problem is not clearly hard for the class FermPCP[$O(1)$]. In the following Lemmas we establish some relations between QPCP and FermPCP:

Lemma 13. QPCP[$O(1)$] \subseteq FermPCP[$O(1)$].

Proof. Let $L \in \text{QPCP}$, and so we seek to construct a fermionic proof system for L . First, if $p(x)$ is odd, add an extra qubit in $|0\rangle$ to the proof so that $p(x)$ is even. In addition, wlog we assume that q , the number of access qubits is even. We use the unitary qubit-assimilation mapping \mathcal{M} of Lemma 15 in Appendix A to map the $k = p(x)$ -qubit state ρ onto a state ρ_f on $3k$ Majorana operators ($3k/2$ fermionic modes), that is, if we expand the witness state ρ in terms of Pauli X_j, Z_j and we replace those by Eqs. (28), (29). Similarly, each occurrence of X_j or Z_j where j is one of the chosen input qubits in the circuit of the verifier is replaced using Eqs. (28), (29). The ancillary qubits remain as is. From the qubit proof, the verifier selects q qubits, and this thus corresponds to selecting $3q$ Majorana operators ($3q/2 = O(1)$ fermionic modes). Thus we can map the qubit proof for $x \in L_{\text{yes}}$ directly onto a fermionic proof and the acceptance probability is the same. Note that the additional fermionic Hilbert space in Lemma 15 is simply not used in this conversion.

When $x \in L_{\text{no}}$, the verifier proceeds identically using the mapping \mathcal{M} , but the prover may provide an arbitrary fermionic state ρ_f of $3k/2$ fermionic modes as input, not obeying the mapping \mathcal{M} . So we need to argue that this state can always be mapped to some qubit state ρ_q , such that if the fermionic verifier is fooled in accepting by this state, then so will the qubit verifier with input ρ_q , which was excluded by definition. The state ρ_f will be an even polynomial in the operators $c_{y,j}, c_{z,j}$ and $c_{x,j}$. Using the \tilde{c}_j operators in Eq. (30), these operators can be mapped to qubit Pauli operators X_j, Z_j and \tilde{c}_j , hence obtaining a state ρ_{qf} . Then we take $\tilde{\rho}_q = \text{Tr}_f(\rho_{qf})$, the partial trace over the fermionic system to get a n -qubit state ρ_q (effectively this

¹Note that it is q to $q + 1$ since one ancilla qubit is fixed as $|0\rangle$ on input but is needed/used as acceptance output qubit.

means omitting any term which involves the operators \tilde{c}_j). Note that since the fermionic verifier uses \mathcal{M} the fermionic gates and final measurement of the fermionic verifier never use the operators \tilde{c}_j and hence tracing out has no effect on the action of the verifier and the selection of $3q/2$ fermionic modes corresponds precisely to the selection qubits of q qubits in ρ_q . Thus if ρ_f would lead the fermionic verifier V to accept with probability $> 1/3$, then the qubit verifier would accept with probability $> 1/3$ on the state ρ_q , in contradiction with the qubit proof system. \square

Lemma 14. $\text{FermPCP}[O(1)] \subseteq \text{QPCP}[O(\log n)]$.

Proof. Let $L \in \text{FermPCP}$, and so we seek to construct a qubit proof system for L . Using the Bravyi-Kitaev transformation [BK02], we map the $k = p(n)$ -fermionic mode state ρ_f onto a k -qubit state and the chosen $2q$ Majorana operators c_j are replaced by products of $O(\log(p(n)))$ single-qubit Pauli operators, hence requiring access to $O(\log n)$ qubits. Thus we can map the fermionic qubit proof for $x \in L_{\text{yes}}$ directly onto a qubit proof, albeit with logarithmic access. When $x \in L_{\text{no}}$, the prover may provide an arbitrary k -qubit state ρ_q as input, not necessarily obeying the Bravyi-Kitaev transformation, while the verifier applies the circuit obtained through the mapping. Applying the inverse of the Bravyi-Kitaev transformation to the state ρ_q leads to fermionic state ρ_f of k fermionic modes. If the qubit verifier accepted on ρ_q with probability $> 1/3$, then applying the inverse would have lead to the fermionic verifier accepting with probability $> 1/3$, which was excluded. \square

Note that the qubit (resp. fermionic) LHD problem which is complete for $\text{QPCP}[O(1)]$ (respectively $\text{FermPCP}[O(1)]$) is not necessarily sparse. Previous results (Theorem 13 in [BH13]) give a polynomial-time classical algorithm for the *dense* qubit case, so that problem certainly cannot be QMA-hard. It is an open question whether the *dense* fermionic LHD problem can be QMA-hard. Finally, whether $\text{FermPCP}[O(1)] = \text{QPCP}[O(1)]$ is an open question, which relates to (1) whether the *sparse* fermionic LHD problem is $\text{FermPCP}[O(1)]$ -hard (note that Proposition 12 merely shows that the general fermionic LHD problem is hard), and (2) the existence of fermion-to-qubit mappings for sparse fermionic interactions which do not introduce non-local constraints. The latter topic is discussed in Section 7.

6 Acknowledgements

We thank Joel Klassen and Sergei Bravyi for discussions on fermion-to-qubit mappings. Y.H. and B.M.T. acknowledge support by QuTech NWO funding 2020-2024 – Part I “Fundamental Research”, project number 601.QT.001-1, financed by the Dutch Research Council (NWO). Y.H and J.H. acknowledge support from the Quantum Software Consortium (NWO Zwaartekracht). AA acknowledges support through the NSF CAREER Award No. 2238836 and NSF award QCIS-FF: Quantum Computing & Information Science Faculty Fellow at Harvard University (NSF 2013303).

7 Discussion

In this paper we constructed fermionic NLTS Hamiltonians from qubit Hamiltonians with well-spread low-energy states. An interesting open question is to consider the reverse construction. It is known that one cannot map all $O(1)$ -local Majorana operators to qubit operators of weight less than $\Omega(\log(n))$ (see e.g. the argument in [HSHT22]), i.e. the Bravyi-Kitaev construction used in Lemma 14 has the optimal scaling in n . It must be noted though, that for *sparse* fermionic Hamiltonians there exists the Bravyi-Kitaev superfast encoding [BK02], which unitarily maps local fermionic Hamiltonians terms onto $O(1)$ -local qubit interactions. However this mapping is only valid in a subspace specified by a set of non-local stabilizer generators. When mapping a fermionic NLTS Hamiltonian to a qubit NLTS Hamiltonian, these generators would have to

be included in the qubit Hamiltonian. In general these generators are not $O(1)$ -local, nor would the resulting Hamiltonian necessarily be sparse. Adapting the BK superfast encoding to avoid these two properties is the topic of active research in quantum simulation (see [CSBM⁺23] and references therein). In Appendix B we describe a construction that achieves both of these objectives, i.e. it encodes a sparse local fermionic Hamiltonian into a sparse local qubit Hamiltonian, through an adaptation of the BK superfast encoding. However it requires $\Theta(n^2)$ qubits to encode n fermionic modes, making it useless when constructing NLTS Hamiltonians. This leaves open the problem of finding a way to encode an n -mode sparse $O(1)$ -local fermionic Hamiltonian into an $O(n)$ -qubit sparse $O(1)$ -local qubit Hamiltonian. Given our repeated failures at finding such a construction we believe there might be some fundamental obstruction. If $\text{FermPCP}[O(1)]$ is strictly larger than $\text{QPCP}[O(1)]$, pursuing a fermionic PCP theorem might be a fruitful endeavor. Finally, we note that PCP versions of QCMA [WFC23] are equivalent when formulated in terms of qubits or fermions.

This paper does not make progress towards understanding the quantum PCP conjecture, but shows that one can construct fermionic Hamiltonians with no trivial low energy states. Extending such a result beyond Majorana fermion stabilizer Hamiltonians would be of some interest. On the more practical side, a natural next step would be to give more realistic families of fermionic Hamiltonians without low-energy trivial states. Developing numerical tools in order to diagnose the fermionic NLTS property in realistic Hamiltonians is another direction of practical interest.

A Qubit assimilation mapping

We detail the mapping used in Theorem 7. See Ref. [BGKT19] for an earlier use of this mapping.

Lemma 15. *For even n , consider the Hilbert space $\mathcal{H}^{(q-f)}$ of n qubits (with Paulis X_j, Y_j and Z_j for $j \in [n]$) and $n/2$ fermions (operators \tilde{c}_j for $j \in [n]$), and the Hilbert space $\mathcal{H}^{(f)}$ of $3n/2$ fermions ($c_{x,j}, c_{y,j}, c_{z,j}$ for $j \in [n]$). There is a unitary map from $\mathcal{H}^{(qf)}$ to $\mathcal{H}^{(f)}$, defined by its action on the generating operators*

$$X_j \mapsto ic_{y,j}c_{z,j}, \quad (28)$$

$$Z_j \mapsto ic_{x,j}c_{y,j}, \quad (29)$$

$$\tilde{c}_j \mapsto ic_{x,j}c_{y,j}c_{z,j}. \quad (30)$$

Proof. Consider the bases of Hermitian operators in $\mathcal{H}^{(f)}$ and $\mathcal{H}^{(qf)}$,

$$\mathcal{P}^{(f)}(\mathbf{s} = \{s_{\alpha,j}, \alpha \in \{x, y, z\}\}) = i^{\frac{|\mathbf{s}|(|\mathbf{s}|-1)}{2}} \prod_j c_{x,j}^{s_{x,j}} c_{y,j}^{s_{y,j}} c_{z,j}^{s_{z,j}}, \quad (31)$$

$$\mathcal{P}^{(qf)}(\mathbf{s}^x = \{s_j^x\}, \mathbf{s}^z = \{s_j^z\}, \mathbf{s}^c = \{s_j^c\}) = i^{\mathbf{s}^x \cdot \mathbf{s}^z + \frac{|\mathbf{s}^c|(|\mathbf{s}^c|-1)}{2}} \prod_j X_j^{s_j^x} Z_j^{s_j^z} \tilde{c}_j^{s_j^c}. \quad (32)$$

Both $\mathcal{H}^{(q-f)}$ and $\mathcal{H}^{(f)}$ are isomorphic to the $3n/2$ -qubit Hilbert space (denote it as $\mathcal{H}^{(q)}$) via the Jordan-Wigner transformation. Under this isomorphism, the Pauli basis of Hermitian operators in $\mathcal{H}^{(q)}$ is equivalent to the Hermitian bases in Eqs. (31)-(32). Sets $\{c_{x,j}, c_{y,j}, c_{z,j}\}$ and $\{X_j, Z_j, \tilde{c}_j\}$ are two alternative sets of generators of the Pauli group, i.e., independent Pauli strings in $\mathcal{H}^{(q)}$. The mapping (denote it as \mathcal{M}) between these generators defines the mapping on the whole Pauli group,

$$\mathcal{M} \left(\prod_j X_j^{s_j^x} Z_j^{s_j^z} \tilde{c}_j^{s_j^c} \right) = \prod_j \mathcal{M} \left(X_j^{s_j^x} \right) \mathcal{M} \left(Z_j^{s_j^z} \right) \mathcal{M} \left(\tilde{c}_j^{s_j^c} \right). \quad (33)$$

Observe that for two elements \mathcal{P}_A and \mathcal{P}_B of the Pauli group defined in terms of $\{X_j, Z_j, \tilde{c}_j\}$ as in Eq. (32), the map \mathcal{M} has the property

$$\mathcal{M}(\mathcal{P}_A)\mathcal{M}(\mathcal{P}_B) = \mathcal{M}(\mathcal{P}_A\mathcal{P}_B). \quad (34)$$

This property follows from Eq. (33) and the fact that $\mathcal{M}(X_j)$, $\mathcal{M}(Z_j)$ and $\mathcal{M}(\tilde{c}_j)$ mutually commute or anticommute in the same way as X_j , Z_j and \tilde{c}_j (among themselves and across different j). The property in Eq. (34) implies that the unitary map claimed in the Lemma exists, defined as a Clifford transformation on $\mathcal{H}^{(q)}$ [CRSS97]. \square

Since the map of Lemma 15 is unitary, it admits an inverse from $\mathcal{H}^{(f)}$ to $\mathcal{H}^{(qf)}$, which acts on the fermionic operators as

$$c_{x,j} \mapsto X_j \tilde{c}_j, \quad c_{y,j} \mapsto -Y_j \tilde{c}_j, \quad c_{z,j} \mapsto Z_j \tilde{c}_j. \quad (35)$$

A curious observation (which we do not use in this work) is that the mapping of $H^{(f)}$ onto $H^{(qf)}$ using Eqs. (35) can be repeated until all fermions are replaced with qubits. This repetition procedure can be done in a variety of ways, each giving a full fermion-to-qubit mapping. In the worst case, the scaling features of the resulting mapping resemble those of Bravyi-Kitaev transformation [BK02]. In practical applications, one may pick a variation of the procedure which leads to the best locality for the resulting qubit Hamiltonian.

B Sparse and local superfast encoding

The Bravyi-Kitaev superfast encoding is a method for mapping sparse and local fermionic Hamiltonians to sparse and local qubit Hamiltonians (up to the enforcement of certain stabilizer constraints). Here we give a version of this encoding which makes sure these stabilizer constraints are themselves local and sparse. Briefly, the BK encoding proceeds as follows. Given an n -fermion Hamiltonian H we construct a graph G with n vertices (one for each fermionic mode) and place an edge if a term in the Hamiltonian involves both modes. More precisely, one can define vertex and edge operators as

$$V_i = c_{2i-1}c_{2i} \quad i \in V(G) = [n], \quad (36)$$

$$E_{i,j} = c_{2i}c_{2j} \quad (i,j) \in E(G). \quad (37)$$

The Hamiltonian H can then be recovered as an element of the ‘‘graph algebra’’ generated by the above operators. We can represent this graph algebra on qubits by placing a qubit on each edge of the graph and defining a map \mathcal{B} taking the E and V operators to Pauli operators which are local with respect to the graph G (see e.g. [SBMW19] for an explicit description). However the graph algebra comes with the non-trivial constraint that the product of the edge operators $E_{i,j}$ along any cycle of G is equal to the identity, as is clear from Eq. (37). For \mathcal{B} to be a proper algebra homomorphism we must thus restrict its image to the subspace where this is true also for the operators $\mathcal{B}(E_{i,j})$. It turns out these constraints all commute and thus form a stabilizer group generated by the products along a cycle basis of G .

The interesting question is whether these stabilizer generators can be made sparse and local. This corresponds to choosing a cycle basis for the graph G that contains only cycles of constant length, and where each edge only participates in a constant number of basis cycles. It is clear that there are graphs of bounded degree for which no such basis exists. Consider for instance the family of n -vertex bounded-degree expander graphs given in [Mar82], which have girth $\Omega(\log(n))$, and hence have no cycles of constant length. Moreover the total length of any basis must then be $\Omega(n \log(n))$ which, from a pigeonhole argument, means




Figure 1: **(a)**: Illustration of the stacking process described in the proof of Lemma 16 with a graph on five vertices. The blue edges are the new edges added in between copies of the graph. **(b)**: Illustration of the sewing process in the proof of Lemma 16. A single large cycle is shown inside a graph, as well as its sewn up version. The blue edges are again added in the process.

that there is at least one edge that is present in $\Omega(\log(n))$ basis cycles.

Here we give a construction that solves both of these problems, by constructing from the graph G a larger graph \hat{G} that has a cycle basis of short cycles that use every edge only a constant number of times, and has G as an induced subgraph. Note that this means that the graph algebra of G is a subalgebra of the graph algebra of \hat{G} (by considering a subset of the generators), and hence this provides a valid mapping of the original fermionic Hamiltonian to a qubit Hamiltonian. The downside of this construction is that the graph \hat{G} is of size $\Theta(n^2)$, which makes it difficult to use it for NLTS-style arguments where the scale of the system (n vs n^2) is important. However the construction might be of use in quantum simulation, and as a starting point for more sophisticated constructions with better parameters.

Lemma 16. *Consider a bounded-degree connected graph G on n vertices. There exists a (polynomial time constructible) connected graph \hat{G} on $O(n^2)$ vertices with G as an induced subgraph that has a cycle basis consisting of cycles of length at most 4 which uses no edge in \hat{G} more than 4 times.*

Proof. We will explicitly construct \hat{G} . First, compute a minimum length cycle basis C for G (for instance through Horton’s algorithm, which takes $O(n^4)$ time). Since G has bounded degree and is connected, the cycle basis has $E(\hat{G}) - V(\hat{G}) + 1 = O(n)$ elements. Order the cycles in C in some arbitrary way. We now construct the graph \hat{G} as follows. For each element of C we make a copy of the graph G . We take these graph copies and “stack” them on top of each other, connecting each vertex in a graph copy to the corresponding vertex in the copies directly above and below (see Fig. 1(a) for an illustration). This creates $E(G)(|C| - 1)$ “vertical” cycles of length 4. It is easy to see that the vertical cycles form an independent set, since every cycle contains an edge that is not used by any other vertical cycle. Furthermore, the set of cycles C is still an independent set in \hat{G} . Furthermore, the union of these two sets is also independent. This union is in fact a basis for \hat{G} , which one can see (by direct calculation) that the dimension of the cycle space of \hat{G} (i.e., $E(\hat{G}) - V(\hat{G}) + 1$) precisely matches the number of vertical cycles plus the dimension of the cycle space of G .

Continuing our construction, consider for each cycle in the set C the associated copy of G . In this copy, “sew” up the cycle by adding edges across the cycle, in the manner illustrated in Fig. 1(b). For each cycle c in C this creates a number of cycles of length 3 or 4. Note that we add $\lceil |c|/2 \rceil$ edges to \hat{G} . Since the total cycle length of a minimum length cycle basis is $O(n \log(n))$ [KLM⁺09, Theorem 4.4], we end up adding

at most $O(n \log(n))$ edges. This completes the construction of \hat{G} . Note also that the degree of \hat{G} is at most three higher than the degree of G .

We now propose the following basis for the cycle space of the graph \hat{G} . We take all vertical cycles, and all the short cycles created by the sewing procedure for each cycle in C . And by our earlier argument the union of the vertical cycles and the cycle basis of G was a basis for \hat{G} before the sewing procedure. Since the sewing procedure adds an independent cycle for each edge it adds, the resulting set is a basis for \hat{G} .

Clearly every cycle in this set is of length no longer than 4. Moreover, every edge is used at most a bounded number of times by each basis cycle. For the vertical edges it is clear that their number of uses is bounded by the degree of G (as they only participate in “vertical” cycles), and the edges in each copy of the graph participate only in the 3 and 4 cycles making up the sewed cycle for that copy, as well as at most two vertical cycles. Hence the number of uses of each edge in the basis is also bounded. \square

References

- [AAG22] Anurag Anshu, Itai Arad, and David Gosset. An area law for 2d frustration-free spin systems. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing, 2022*. URL: <https://doi.org/10.1145%2F3519935.3519962>.
- [AALV09] Dorit Aharonov, Itai Arad, Zeph Landau, and Umesh Vazirani. The detectability lemma and quantum gap amplification. In *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing, STOC '09*, page 417–426, New York, NY, USA, 2009. Association for Computing Machinery. URL: <https://doi.org/10.1145/1536414.1536472>.
- [AAV13] Dorit Aharonov, Itai Arad, and Thomas Vidick. The quantum PCP conjecture. *CoRR*, abs/1309.7495, 2013. URL: <http://arxiv.org/abs/1309.7495>.
- [ABN22] A. Anshu, N.P. Breuckmann, and C. Nirkhe. Nlts Hamiltonians from good quantum codes. In *55th Annual ACM Symposium on Theory of Computing (STOC 2023)*, pages 1090–1096. ACM, 2022. URL: <https://dl.acm.org/doi/10.1145/3564246.3585114>.
- [AGK23] Eric R. Anschuetz, David Gamarnik, and Bobak Kiani. Combinatorial NLTS from the overlap gap property, 2023. URL: <https://doi.org/10.48550/arXiv.2304.00643>.
- [BCDZ99] H. Buhrman, R. Cleve, R. De Wolf, and C. Zalka. Bounds for small-error and zero-error quantum algorithms. In *Proc. of FOCS 99*, pages 358–368, 1999.
- [BG17] Sergey Bravyi and David Gosset. Complexity of quantum impurity problems. *Communications in Mathematical Physics*, 356(2):451–500, aug 2017. URL: <https://doi.org/10.1007%2Fs00220-017-2976-9>.
- [BGKT19] Sergey Bravyi, David Gosset, Robert König, and Kristan Temme. Approximation algorithms for quantum many-body problems. *Journal of Mathematical Physics*, 60(3):032203, 2019. URL: <https://doi.org/10.1063/1.5085428>.
- [BH13] Fernando G.S.L. Brandao and Aram W. Harrow. Product-state approximations to quantum ground states. In *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing, STOC '13*, page 871–880, New York, NY, USA, 2013. Association for Computing Machinery. URL: <https://doi.org/10.1145/2488608.2488719>.
- [BK02] Sergey Bravyi and Alexei Kitaev. Fermionic quantum computation. *Annals of Physics*, 298(1):210–226, 2002. URL: <https://doi.org/10.1006%2Faphy.2002.6254>.

- [Bra05] Sergey Bravyi. Lagrangian representation for fermionic linear optics. *Quantum Info. Comput.*, 5(3):216–238, may 2005. URL: <https://arxiv.org/abs/quant-ph/0404180>.
- [Bra06] Sergey Bravyi. Universal quantum computation with the $\nu = 5/2$ fractional quantum hall state. *Phys. Rev. A*, 73:042313, Apr 2006. URL: <https://link.aps.org/doi/10.1103/PhysRevA.73.042313>.
- [CCNN23] Nolan J Coble, Matthew Coudron, Jon Nelson, and Seyed Sajjad Nezhadi. Local hamiltonians with no low-energy stabilizer states. *arXiv preprint arXiv:2302.14755*, 2023. URL: <https://doi.org/10.48550/arXiv.2302.14755>.
- [CRSS97] A. R. Calderbank, E. M. Rains, P. W. Shor, and N. J. A. Sloane. Quantum error correction and orthogonal geometry. *Phys. Rev. Lett.*, 78:405–408, Jan 1997. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.78.405>.
- [CSBM⁺23] Riley W Chien, Kanav Setia, Xavier Bonet-Monroig, Mark Steudtner, and James D Whitfield. Simulating quantum error mitigation in fermionic encodings. *arXiv preprint arXiv:2303.02270*, 2023. URL: <https://doi.org/10.48550/arXiv.2303.02270>.
- [FH13] M. H. Freedman and M. B. Hastings. Quantum systems on non- k -hyperfinite complexes: A generalization of classical statistical mechanics on expander graphs, 2013. URL: <https://doi.org/10.48550/arXiv.1301.1363>.
- [GG22] Sevag Gharibian and Francois Le Gall. Dequantizing the quantum singular value transformation: hardness and applications to quantum chemistry and the quantum PCP conjecture. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. ACM, jun 2022. URL: <https://doi.org/10.1145%2F3519935.3519991>.
- [HO22] Matthew B. Hastings and Ryan O’Donnell. Optimizing strongly interacting fermionic Hamiltonians. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2022, page 776–789, New York, NY, USA, 2022. Association for Computing Machinery. URL: <https://doi.org/10.1145/3519935.3519960>.
- [HSHT22] Yaroslav Herasymenko, Maarten Stroeks, Jonas Helsen, and Barbara Terhal. Optimizing sparse fermionic hamiltonians. *arXiv preprint arXiv:2211.16518*, 2022. URL: <https://doi.org/10.48550/arXiv.2211.16518>.
- [JSK⁺18] Zhang Jiang, Kevin J. Sung, Kostyantyn Kechedzhi, Vadim N. Smelyanskiy, and Sergio Boixo. Quantum algorithms to simulate many-body physics of correlated fermions. *Physical Review Applied*, 9(4), apr 2018. URL: <https://doi.org/10.1103%2Fphysrevapplied.9.044036>.
- [Kit06] Alexei Kitaev. Anyons in an exactly solved model and beyond. *Annals of Physics*, 321(1):2–111, 2006. URL: <https://doi.org/10.1016%2Fj.aop.2005.10.005>.
- [KLM⁺09] Telikepalli Kavitha, Christian Liebchen, Kurt Mehlhorn, Dimitrios Michail, Romeo Rizzi, Torsten Ueckerdt, and Katharina A Zweig. Cycle bases in graphs characterization, algorithms, complexity, and applications. *Computer Science Review*, 3(4):199–243, 2009. URL: <https://doi.org/10.1016/j.cosrev.2009.08.001>.
- [KLS96] Jeff Kahn, Nathan Linial, and Alex Samorodnitsky. Inclusion-exclusion: Exact and approximate. *Combinatorica*, 16(4):465–477, 1996. doi:10.1007/BF01271266.

- [KSV02] A. Yu. Kitaev, A.H. Shen, and M.N. Vyalyi. *Classical and Quantum Computation. Vol. 47 of Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2002.
- [LCV07] Yi-Kai Liu, Matthias Christandl, and Frank Verstraete. Quantum computational complexity of the n-representability problem: Qma complete. *Physical review letters*, 98(11):110503, 2007. URL: <https://doi.org/10.1103/PhysRevLett.98.110503>.
- [LZ22] Anthony Leverrier and Gilles Zémor. Quantum tanner codes. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 872–883. IEEE, 2022. URL: <https://doi.org/10.1109/FOCS54457.2022.00117>.
- [Mar82] Grigori A Margulis. Explicit constructions of graphs without short cycles and low density codes. *Combinatorica*, 2(1):71–78, 1982. URL: <https://doi.org/10.1007/BF02579283>.
- [ODMZ22] Michał Oszmaniec, Ninnat Dangniam, Mauro E.S. Morales, and Zoltán Zimborás. Fermion sampling: A robust quantum computational advantage scheme using fermionic linear optics and magic input states. *PRX Quantum*, 3:020328, May 2022. URL: <https://link.aps.org/doi/10.1103/PRXQuantum.3.020328>.
- [SBMW19] Kanav Setia, Sergey Bravyi, Antonio Mezzacapo, and James D Whitfield. Superfast encodings for fermionic quantum simulation. *Physical Review Research*, 1(3):033033, 2019. URL: <https://doi.org/10.1103/PhysRevResearch.1.033033>.
- [TD02] Barbara M. Terhal and David P. DiVincenzo. Classical simulation of noninteracting-fermion quantum circuits. *Physical Review A*, 65(3), mar 2002. URL: <https://doi.org/10.1103/PhysRevA.65.032325>.
- [WFC23] Jordi Weggemans, Marten Folkertsma, and Chris Cade. Guidable local hamiltonian problems with implications to heuristic ansätze state preparation and the quantum pcp conjecture. *arXiv preprint arXiv:2302.11578*, 2023. URL: <https://doi.org/10.48550/arXiv.2302.11578>.