

The Rotation-Invariant Hamiltonian Problem Is QMA_{EXP} -Complete

Jon Nelson   

Joint Center for Quantum Information and Computer Science (QuICS),
Department of Computer Science, University of Maryland, College Park, MD, USA

Daniel Gottesman   

Joint Center for Quantum Information and Computer Science (QuICS),
Department of Computer Science, University of Maryland, College Park, MD, USA

Abstract

In this work we study a variant of the local Hamiltonian problem where we restrict to Hamiltonians that live on a lattice and are invariant under translations and rotations of the lattice. In the one-dimensional case this problem is known to be QMA_{EXP} -complete. On the other hand, if we fix the lattice length then in the high-dimensional limit the ground state becomes unentangled due to arguments from mean-field theory. We take steps towards understanding this complexity spectrum by studying a problem that is intermediate between these two extremes. Namely, we consider the regime where the lattice dimension is arbitrary but fixed and the lattice length is scaled. We prove that this rotation-invariant Hamiltonian problem is QMA_{EXP} -complete answering an open question of [6]. This characterizes a broad parameter range in which these rotation-invariant Hamiltonians have high computational complexity.

2012 ACM Subject Classification Theory of computation \rightarrow Quantum complexity theory

Keywords and phrases Hamiltonian complexity, Local Hamiltonian problem, Monogamy of entanglement

Digital Object Identifier 10.4230/LIPIcs.TQC.2025.12

Funding *Jon Nelson*: supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE 2236417.

1 Introduction

In order to understand the behavior of a quantum many-body system, it is crucial to study its Hamiltonian. The Hamiltonian operator not only governs the system's dynamics through the Schrödinger equation but also encodes its low-energy states and energy spectrum. It is thus important to understand which Hamiltonians are tractable to analyze. In this work, we study the computational complexity of estimating the ground-state energy of a Hamiltonian with only short-range interactions, which is known as the local Hamiltonian problem.

Often the goal is to show that a specific variant of this problem is QMA -complete, which implies that for certain Hamiltonians not even a quantum computer can be expected to find its ground state energy. On the one hand, this is a negative result for being able to calculate ground state energies. On the other hand, these results lead to constructions of highly complex quantum systems that are interesting objects of study in their own right.

Kitaev initiated the study of local Hamiltonian problems in his landmark result proving that this problem in its most general form is QMA -complete [7]. However, Kitaev's result only applies for a worst-case family of Hamiltonians, which are not physically natural. In order to study the complexity of more physically relevant cases, subsequent work has extended Kitaev's result to apply under additional constraints that capture what it means to be “natural”. This has included restricting the local dimension of each particle as well as



© Jon Nelson and Daniel Gottesman;
licensed under Creative Commons License CC-BY 4.0

20th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2025).

Editor: Bill Fefferman; Article No. 12; pp. 12:1–12:18



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

constraining the geometry of the interactions. For example, [10] extends Kitaev's result to qubits on a 2D lattice while [2] further restricts to particles on a line with constant local dimension.

Additional follow-up work has also emphasized symmetry constraints. This is motivated by the observation that many systems in nature are highly symmetric. For instance, the laws of gravity, electromagnetism, etc., do not change depending on where you are or how you are oriented; thus, these laws are translation and rotation-invariant. For translation-invariant one-dimensional spin chains, [6] showed that this local Hamiltonian problem is QMA_{EXP} -complete.

Although much work is now known about the complexity of translation-invariant systems [6, 4, 3, 8, 11], there have been very few results for the rotation-invariant case. In fact, it was posed as an open question of [6] whether their results can be extended to rotation-invariant Hamiltonians in higher dimensions. In this work, we solve this question and hope similar techniques can be used to lift other translation-invariant results to the rotation-invariant case.

Translation invariance with reflection symmetry and rotation invariance coincide in 1D and so the main challenge is to extend [6] to Hamiltonians on higher dimensional lattices. It is trivial to extend their result to higher dimensional translation-invariant lattices simply by ignoring all but one dimension. However, this breaks the rotation symmetry, which requires that the Hamiltonian terms act identically in all directions. In this case, the key challenge is to handle the increasingly high degree of interaction without increasing the number of parameters in the Hamiltonian. This presents issues, for example, when attempting to encode computation into the Hamiltonian's ground state, which is an essential step for proving hardness. Controlling this computation requires the ability to track time, which can be accomplished in the 1D setting by moving a clock pointer along the spin chain [6]. However, this same idea cannot be used in higher dimensions since the paths can branch in many directions throughout the lattice. In order to pick out a specific time direction, we must engineer a family of Hamiltonians that spontaneously breaks the rotation symmetry.

Technical difficulties aside, it may seem intuitive that increasing the lattice dimension only makes the local Hamiltonian problem more difficult, and so one might assume that the complexity for higher dimensional cases follows from the one-dimensional case. However, due to the rotation symmetry, the increase in lattice dimension does not correspond to more Hamiltonian parameters, and so it is unclear how the complexity actually compares. In fact, standard condensed matter arguments imply that increasing the dimension can instead make the problem easier. This follows from the observation that for higher-dimensional lattices, mean-field theory (which uses a product-state ansatz to approximate the ground state) becomes more and more accurate [12]. In the quantum setting, this can be explained by an effect called monogamy of entanglement [14], which states that a particle cannot be highly entangled with many other particles. Thus, for high lattice dimension, each particle has many neighbors and so on average they must be nearly unentangled. Due to this effect, the product state becomes a good approximation of the ground state, suggesting that this problem could now be more tractable than the lower dimensional cases.

This has been formalized in [5] which shows that for lattice dimension r there is a product state that approximates the ground-state energy by an average error of $O(r^{-1/3})$ per term of the Hamiltonian. Furthermore, [9] rigorously show that in the limit as $r \rightarrow \infty$ the ground state is exactly a product state when the Hamiltonian is translation and rotation invariant. These results suggest that if the lattice dimension is high enough, the problem loses its quantum hardness since the low-energy states become unentangled. Another result that

captures this phenomenon is [1], who show that a commuting version of the local Hamiltonian problem becomes easier as the interaction graphs become more expanding, which intuitively corresponds to more interaction.

In this work, we consider a lattice dimension that is in an intermediate regime between one-dimensional spin chains, which are hard and spin chains with $r \rightarrow \infty$, which are easy. In particular, we consider an arbitrary but fixed lattice dimension and show that this rotation-invariant Hamiltonian problem is quantumly hard as you scale the lattice length.

1.1 Results

We informally describe the rotation-invariant Hamiltonian problem as follows.

► **Definition 1** (Rotation-invariant Hamiltonian problem (Informal)). *Consider the Hamiltonian where a single two-body term is applied to each neighboring pair of qudits on an r -dimensional lattice of side length n . Is the ground-state energy below a or greater than b ?*

Our main result is that the rotation-invariant Hamiltonian problem is QMA_{EXP} -complete, where QMA_{EXP} is the same as QMA except the witness and verification circuit are allowed to be exponentially large in the input size. The reason we consider QMA_{EXP} rather than QMA is that the input size to our problem is actually very small in comparison to the size of the Hamiltonian. To see this, notice that our Hamiltonian can be completely described by 1) the two-body term, 2) the dimension r , and 3) the lattice length n . The first two of these only require a constant number of bits to specify, while n requires $\log n$ bits to specify. Therefore, the Hamiltonian description length is exponentially smaller than the total number of qudits. However, we still would like to allow an “efficient” algorithm to run in polynomial time with respect to the number of qudits, which in turn is exponential in the input size. To accommodate this technicality, we must prove quantum hardness even when the quantum computer is allowed an exponential amount of computation time.

To prove QMA_{EXP} -completeness, we use the standard method of reducing an instance x of an arbitrary QMA_{EXP} problem to an instance $R(x)$ of the rotation-invariant Hamiltonian problem. It turns out that in this reduction only n depends on the original problem instance x , so we take everything else (such as the two-body term and the lattice dimension) to be parameters rather than inputs to the problem. This is described in the more technical definition of the problem in Section 2. This differs from the standard QMA -completeness result, where the Hamiltonian terms themselves are given as input. We argue that this is a more natural setting since often one is studying a particular Hamiltonian, and so it is more suitable to consider the hardness of a given Hamiltonian for increasingly large system sizes. A desirable feature of our reduction is that the Hamiltonian we construct has no dependence on the system size or lattice dimension.

The QMA_{EXP} -completeness of this problem has a number of interesting implications. First, it suggests that not even a quantum computer can find the ground-state energy of certain rotation-invariant Hamiltonians. Since nature can be viewed as a quantum computer this means that the system itself cannot find its own ground state either, suggesting the emergence of spin glass behavior at low temperatures. Next, our result implies (assuming $\text{QCMA}_{\text{EXP}} \neq \text{QMA}_{\text{EXP}}$) that the ground state of these Hamiltonians cannot have an efficient classical description and thus cannot be well approximated by a product state. In fact, our result directly implies a lower bound of $\Omega(n^{-r}r^{-1})$ for how close the average ground-state energy per term can be approximated by a product state. This complements the result in [5] by providing a corresponding lower bound for the product-state approximation error.

1.2 Techniques

Our main approach is to carve out one-dimensional chains within our higher dimensional lattice so that we can apply [6]’s 1D construction to these chains. To do this, we take inspiration from the closely related classical problem of tiling. In this problem, imagine fitting together square tiles to cover an entire floor, where we are given a penalty for placing certain color tiles next to each other. The task is now to design a set of tiling rules where the least penalized configuration is a pattern of stripes. If such a set of tiling rules exists, we can use a portion of each qudit’s Hilbert space to represent a tile and encode the tiling rules in our two-body Hamiltonian term. This enforces that the tiling of the ground state will have this striped pattern. Our construction then proceeds by applying the 1D construction onto neighboring qudits with same-colored tiles, which are now effectively spin chains.

Unfortunately, such a set of tiling rules does not actually exist, and so we will have to modify this classical technique to incorporate some quantum phenomena. To see this, consider the 3D case with periodic boundary conditions. No matter what set of rules are given, the optimal tiling will always take the following form. Start by tiling the first column of the first 2D slice with the optimal 1D configuration. Then, for each subsequent column in this slice, tile it by offsetting this sequence by exactly one. Finally for each subsequent 2D slice, tile by shifting the entire previous 2D configuration by one. With some thought, one can convince themselves that this is the correct tiling. The issue is that this configuration cuts the 3D lattice into diagonal planes which is not the desired 1D structure. Additionally, this argument also shows that these rotation-invariant tiling problems are in P whereas their translation-invariant (but not rotation-invariant) counterparts are NEXP-complete for dimension 2 and higher [6]. This further shows how the additional symmetry constraints can potentially simplify the complexity.

It is possible to still achieve our tiling goal by combining it with some purely quantum effects. In particular, we introduce a new technique that uses the monogamy of entanglement to enforce an effective 1D geometry. This is performed by first appending two qubits to the Hilbert space of each particle. The key idea is to enforce that same-colored neighbors share an EPR pair among their qubits. Since each site only has two qubits, it can only share an EPR pair with two neighbors by the monogamy of entanglement. Thus, it can only have two same-colored neighbors. This is already close to our goal, since now our lattice must be colored by disjoint same-colored loops. It remains to make sure that these loops do not have any turns but instead cut straight across the lattice. This can be handled by imposing some further classical tiling constraints, which we discuss in more detail in Section 3. We hope that this EPR pair technique can also find use in other Hamiltonian complexity problems that benefit from embedding lower dimensional geometries into higher ones.

One last technicality to resolve is that [6]’s 1D Hamiltonian is frustrated and has an energy of at least $1/2$. This results in an overall energy of $n^{r-1}/2$ when this Hamiltonian is embedded into n^{r-1} 1D lines in the lattice. In order to balance this energy penalty with the rest of the Hamiltonian terms, it is necessary to normalize this contribution with a coefficient that depends on n . However, such a system size dependence is unnatural and is preferably avoided. To fix this, we first embed a 2D translation-invariant Hamiltonian into the lattice by encoding stripes in two different directions as opposed to just one. In this case, the same result can be achieved as in the 1D case except the construction can now be made to be nearly frustration-free, removing the need for a system-size dependent normalization term.

1.3 Outline

We begin by describing our notation and briefly state the technical version of our result in Section 2. Next, we define the Hamiltonian construction in Section 3. Finally, the proof of our main theorem is presented in Section 4 where it is shown that our construction satisfies both completeness and soundness.

2 Notation and technical result

Define $\Lambda_r(n) := \mathbb{Z}^r / n\mathbb{Z}^r$ to be a periodic lattice. In other words, the lattice is r -dimensional, where each dimension has length n and each site is denoted by integer coordinates. For a lattice point $u \in \Lambda_r(n)$ we denote the i th coordinate as u_i . To define distance between points in the lattice while respecting the periodic boundary conditions, we use the Lee metric, which is defined as follows:

► **Definition 2** (Lee metric). *Let $x, y \in \Lambda_r(n)$.*

$$d(x, y) = \sum_{i=1}^r \min(|x_i - y_i|, n - |x_i - y_i|)$$

The set of nearest-neighbor pairs is defined as $E_{\Lambda_r(n)} = \{\{x, y\} : x, y \in \Lambda_r(n), d(x, y) = 1\}$. If h is a 2-local Hamiltonian term and $u, v \in \Lambda_r(n)$ then $h^{u,v}$ denotes the Hamiltonian term h applied to sites u and v . For an operator H , we denote the lowest eigenvalue of H by $E_0(H)$.

► **Definition 3** (QMA_{EXP}). *A language L is in QMA_{EXP} if there exists a quantum verifier V such that on input x , V has runtime $O(2^{|x|^k})$ for some k . In addition, if $x \in L$ then there exists a state $|\psi\rangle \in \mathbb{C}^{O(2^{|x|^k})}$ such that $V(x, |\psi\rangle)$ accepts with probability at least $2/3$. If $x \notin L$, then $V(x, |\psi\rangle)$ accepts with probability at most $1/3$.*

With this notation in hand, the formal definition of the rotation-invariant Hamiltonian problem can be stated as follows:

► **Definition 4** (r -DIM-RIH (Rotationally-Invariant Hamiltonian)).

Problem Parameter: *The geometric dimension of the lattice r . A permutation-invariant Hermitian operator h . Two polynomials p and q .*

Input: *Integer n specified in binary.*

Promise: *Let $N = |\Lambda_r(n)| = n^r$. Consider the Hamiltonian $H = \sum_{\{u,v\} \in E_{\Lambda_r(n)}} h^{u,v}$. The ground state energy of H is either at most $p(n)$ or at least $p(n) + 1/q(n)$.*

Output: *Determine whether the ground-state energy of H is at most $p(n)$ or at least $p(n) + 1/q(n)$.*

In particular, notice that h does not depend on the system size n or the lattice dimension r in this definition. Our main result is the following theorem.

► **Theorem 5.** *r -DIM-RIH is QMA_{EXP} -complete for $q(n) = 1$.*

3 Hamiltonian construction

Our strategy will be first to embed a two-dimensional translation-invariant Hamiltonian without reflection symmetry into our *rotationally*-invariant Hamiltonian. In the case where our lattice has dimension r , we will break up the lattice into n^{r-2} 2D slices and embed

the Hamiltonian into each of these slices. Then we can utilize the extra parameters of the 2D translation-invariant Hamiltonian to embed a one-dimensional hard Hamiltonian that is nearly frustration-free.

To accomplish this, it is crucial to have a mechanism to break the rotation symmetry by selecting a direction. We will do this by embedding two sets of directed stripes to indicate the two directions of the 2D grid.

3.1 Embedding directed stripes

In our construction, we will attach the following Hilbert spaces to each site in the lattice:

$$\mathcal{H}_{T_1} \otimes \mathcal{H}_{T_2}$$

Our Hamiltonian will act diagonally in these subspaces and, therefore, it will only enforce classical constraints with respect to a given set of basis states, which we denote by the sets T_1 and T_2 , respectively. We refer to these basis states as “tiles” that we can assign to each site.

We define $T_1 := \{\text{red, yellow, blue}\}$ and $T_2 := \{0, 1, 2\}$, which associate a color and number with each tile, respectively. It can be enforced that two tiles t_1 and t_2 cannot be placed next to each other by including the term $|t_1, t_2\rangle\langle t_1, t_2|$ in the Hamiltonian. In this way, we incorporate the following rules for which tiles are allowed to be placed next to each other:

1. If two neighboring tiles have the same color then they must have different numbers.
2. If two neighboring tiles have different colors then they must have the same number.

To write down the Hamiltonian terms associated with these rules more explicitly, let V represent the set of illegal neighboring tiles. Then we include the following Hamiltonian term:

$$h_{\text{tile}} = 8 \sum_{(s,t) \in V} |s, t\rangle\langle s, t|$$

The energy cost of 8 is carefully chosen to balance out other competing terms introduced later.

3.1.1 EPR projections


Next, we would like to enforce that the qubits of same-colored neighbors form EPR pairs with each other. We can do this by attaching the following two additional Hilbert spaces to each site:

$$\mathcal{H}_{\sigma_1} \otimes \mathcal{H}_{\sigma_2}$$


Since each site only has two qubits it can only form two EPR pairs and therefore can only have two same-colored neighbors. Thus, this accomplishes our goal of forcing the same-colored tiles to form one-dimensional chains.

Given two same-colored neighbors u and v , it remains to determine which of their qubits must form EPR pairs. To do this, we first define directed edges between each basis state of T_2 such that $0 \rightarrow 1$, $1 \rightarrow 2$ and $2 \rightarrow 0$. Due to the constraints in the previous section, u and v must both have different numbers. Without loss of generality, if the directed edge between these numbers points towards v 's tiles then we enforce that u 's σ_2 qubit forms an EPR pair with v 's σ_1 qubit. We denote this EPR pair state as $|\Phi^+\rangle_{u\sigma_2 v\sigma_1} = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)_{u\sigma_2 v\sigma_1}$. More formally we add the following term acting on sites u and v :

$$A^{u,v} = 16 \sum_{i \in T_2} |i, i+1 \bmod 3\rangle\langle i, i+1 \bmod 3| \otimes \left(\frac{\mathbb{I} - |\Phi^+\rangle\langle\Phi^+|}{2} \right)_{u\sigma_2 v\sigma_1}$$



■ **Figure 1** When a chain of sites is numbered sequentially around the cycle \mathbb{Z}_3 , each qubit is matched with exactly one other qubit to form an EPR pair and so the EPR constraint can easily be satisfied.



■ **Figure 2** When there are three consecutive same-color neighbors that are not numbered monotonically around \mathbb{Z}_3 (for instance 0, 1, 0) then two different qubits are matched with the same qubit to form an EPR pair. Due to the monogamy of entanglement this constraint cannot be satisfied and incurs an energy penalty.

Our convention throughout is to separate sites of the lattice by commas within the bracket notation and to separate subspaces within each site by tensor product symbols. Notice that if u and v have different numbers then this implies they also have the same color and so this does not need to be additionally conditioned on in this Hamiltonian term. In order to preserve rotation invariance, we add this term for both orderings of the particles u and v :

$$h_{\text{EPR}} = A^{u,v} + A^{v,u} \quad (1)$$


In addition to enforcing that contiguous regions of same-colored sites form 1D chains, these EPR constraints also require that each same-colored chain is numbered as the periodic sequence: 0, 1, 2, 0, 1, 2, ... either in the forwards or backwards direction (see Figure 1). This is because there can never be a site where the directed edges incident to it are both pointing away or both pointing towards it. In either case, this requires one of that site's qubits to be in two different EPR pairs, which is not allowed by the monogamy of entanglement. This scenario is depicted in Figure 2. This sequential numbering, is very helpful because it defines a direction to each chain. Notice that such a numbering is only possible when n is a multiple of 3, so we will later incorporate this restriction into our hardness reduction.

3.1.2 1D chain boundary conditions

Given the current Hamiltonian terms, the same-colored sites can form chains with either open or periodic boundary conditions (lines or loops). It will be convenient later that the boundary conditions are fixed and so we add another term to enforce periodic boundary conditions. In particular, we define the following term acting on sites u and v :

$$h_{\text{loop}} = 2 \sum_{c,d \in T_1, c \neq d} |c, d\rangle\langle c, d|,$$

which penalizes neighboring sites with different colors. This adds a penalty of $2r - 2$ for every site in the middle of the chain. This is because each site has $2r$ neighbors, where all but exactly two are colored differently. This results in a penalty of $2(2r - 2)/2$ where we have divided by two to fix double-counting. Using similar reasoning, a penalty of $2r - 1$ is



■ **Figure 3** An example of how a 2D lattice can be tiled to optimize the Hamiltonian terms. Specifically, each copy forms a striped pattern where each stripe is numbered in a cyclic sequence. In addition, the two copies of tiles must have stripes pointing in different directions. Finally, the rows/columns that do not hold stripes must be numbered with the same number. Now a translation-invariant Hamiltonian can be simulated by using these tilings as a guideline for which sites are above, below, left, or right of each other. As our convention, we take the first copy to denote the horizontal direction and the second copy to denote the vertical direction.

incurred for every site at the endpoint of an open chain. Assuming that the classical tiling and EPR constraints are satisfied, this term is optimized when all 1D chains form loops so that each site always neighbors two other sites of the same color. The scaling of 2 is chosen so that the energy savings of coloring neighboring sites the same will never outweigh the energy penalty of violating the EPR constraints. This tradeoff will be worked out in detail in Section 4.2.

3.2 Adding the second dimension

Now that we have embedded stripes in one direction of the lattice, we must repeat this process to embed stripes in another direction. To do this, we can simply make a copy of each site's Hilbert space and apply the same Hamiltonian terms to the copy. It remains to ensure that both copies do not have stripes oriented in the same direction. To do this, it is sufficient to simply disallow two neighboring tiles from having matching colors on both copies. In other words, we add the following term where we let \mathcal{H}_{T_1} denote the T_1 subspace of the first copy and \mathcal{H}'_{T_1} denote that of the second copy.

$$h_{\text{copy}} = \left(\sum_{c \in T_1} |c, c\rangle\langle c, c| \right)_{\mathcal{H}_{T_1}} \otimes \left(\sum_{d \in T_1} |d, d\rangle\langle d, d| \right)_{\mathcal{H}'_{T_1}}$$

3.3 Embedding the translation-invariant Hamiltonian

An arbitrary two-dimensional translation-invariant Hamiltonian H_{TI} can now be embedded into our rotation-invariant Hamiltonian by using the directed stripes as guidelines. Our convention will be to let the first copy of each site represent the horizontal stripes and the second copy represent the vertical stripes. These tiling patterns are depicted in Figure 3. Now, the horizontal Hamiltonian term is applied only when the first copy tiles have the same color (i.e., different numbers). In addition, the orientation of which site is on the left and which is on the right will be decided by the directed edge in between the two tile's numbers where each arrow points from left to right. The vertical Hamiltonian term is applied similarly with respect to the second copy of each site's tiles.

To make this more concrete, we first attach to each site of our lattice the Hilbert space of a site in H_{TI} which we denote by \mathcal{H}_{2D} . Denote the horizontal 2-body term of H_{TI} by h_{TI} . To incorporate this into our construction we add the following term:

$$h_{\text{RI}} = \sum_{i \in T_2} |i, i+1 \bmod 3\rangle\langle i, i+1 \bmod 3| \otimes h_{\text{TI}} \quad (2)$$

$$+ \sum_{i \in T_2} |i, i-1 \bmod 3\rangle\langle i, i-1 \bmod 3| \otimes S h_{\text{TI}} S \quad (3)$$

where S is the swap operator on the two sites and the term acts on the first copy of \mathcal{H}_{T_2} . Note that while h_{TI} does not have reflection symmetry, h_{RI} does. The vertical term is implemented in the same way but acts on the second copy. We denote the vertical term by v_{RI} .

It remains to now describe the 2D translation-invariant Hamiltonian that encodes the computational hardness. We will combine techniques from [6] in order to encode the desired ground state energy in the yes and no cases. The details are deferred to Section A but the key result is outlined below.

► **Theorem 6.** *Let L be a QMA_{EXP} -complete language. There exists an efficiently computable function $f : \{0, 1\}^* \rightarrow \mathbb{Z}$ and 2-local positive semidefinite Hamiltonian terms h_{TI} and v_{TI} with the following properties:*

1. *$f(x)$ is a multiple of 3 and $f(x)/3$ is prime. Furthermore, $\log f(x) = O(\text{poly}(|x|))$ and f is computable in $O(\text{poly}(|x|))$ time.*
2. *Let $f(x) = n \geq n_0$ for some constant n_0 and a given problem instance $x \in \{0, 1\}^*$. For an $n \times n$ 2D lattice with periodic boundary conditions, let E_h be the set of ordered pairs of horizontal neighbors and let E_v be the set of ordered pairs of vertical neighbors. Consider the Hamiltonian $H_{\text{TI}} = \sum_{(u,w) \in E_h} h_{\text{TI}}^{u,w} + \sum_{(x,y) \in E_v} v_{\text{TI}}^{x,y}$.*
 - a. *If $x \in L$ then $E_0(H_{\text{TI}}) \leq O(n^{-k})$ for an arbitrarily large constant k .*
 - b. *If $x \notin L$ then $E_0(H_{\text{TI}}) \geq \Omega(1/n^3)$*

Using the prime symbol to denote terms acting on the second copy of a site's Hilbert space, we can now write the entire 2-body Hamiltonian term of our construction as follows:

$$h = h_{\text{tile}} + h_{\text{EPR}} + h_{\text{loop}} + h'_{\text{tile}} + h'_{\text{EPR}} + h'_{\text{loop}} + h_{\text{copy}} + h_{\text{RI}} + v_{\text{RI}} \quad (4)$$

We can now define the reduction from any QMA_{EXP} problem instance to an r-DIM-RIH problem instance.

► **Definition 7** (Reduction from QMA_{EXP} to r-DIM-RIH). *Let L be a language in QMA_{EXP} and let $x \in \{0, 1\}^*$ be a problem instance. Define the function $R(L, x, r) = \sum_{\{u,v\} \in E_{\Lambda_r(f(x))}} h^{u,v}$ where $f : \{0, 1\}^* \rightarrow \mathbb{Z}$ is constructed as in Theorem 6 and h as in Eq. 4.*

Next, the detailed analysis of this construction is provided.

4 Analysis

In this section we prove the main theorem that the rotation-invariant Hamiltonian problem is QMA_{EXP} -complete.

4.1 Completeness

We begin by first considering the case where $x \in L$. In this case, we present a ground state that achieves energy below $p(n) = 4n^r(r-1) + 1/g(n)$ for an arbitrarily large polynomial $g(n)$.

► **Lemma 8.** *Let L be a language in QMA_{EXP} and let $x \in \{0,1\}^*$ be a problem instance. Let $H = R(L, x, r)$. If $x \in L$, then $E_0(H) \leq 4n^r(r-1) + 1/g(n)$ for any polynomial $g(n)$.*

Proof. Consider the following tiling, which generalizes the tiling of the 2D lattice depicted in Figure 6. We first construct a 3-coloring of the $(r-1)$ -dimensional lattice. This always exists because an $(r-1)$ -dimensional lattice can be decomposed as a Cartesian product of cycle graphs. Since each cycle graph is 3-colorable their Cartesian product is also 3-colorable by a result by Sabidussi [13]. To color a site in the full r -dimensional lattice we simply drop the 1st coordinate and assign the color from the 3-coloring of the remaining $r-1$ coordinates. This enforces that any two neighboring sites that have the same 1st coordinate are colored by different colors and any two neighbors that only differ in the 1st coordinate are tiled by the same colors. This has the effect of coloring 1D chains of sites that travel in a straight line along the 1st coordinate dimension.

For this coloring, it is easy to satisfy the numbering constraints. This can be accomplished by tiling all particles u with the number $u_1 \bmod 3$. This simultaneously tiles all 1D chains with the ordering $0, 1, 2, 0, 1, 2, \dots$ which eventually wraps around since n is a multiple of 3. In addition, all neighboring tiles with different color tiles are tiled with the same number since this only occurs when the two particles have the same 1st coordinate.

The EPR constraint can easily be satisfied since the particles have been tiled as disjoint 1D chains with the appropriate numbering. In addition, all chains are loops and so the constraint on having periodic boundary conditions is also satisfied. This ensures that only a penalty of $n^r(2r-2) = 2n^r(r-1)$ is introduced by the h_{loop} term. We can repeat this for the second copy of tiles but now directing the 1D chains along the second coordinate. This introduces another penalty of $2n^r(r-1)$.

With this choice of tiles the lattice has effectively been broken up into n^{r-2} 2D slices where we can now apply the 2D translation-invariant Hamiltonian construction. By Theorem 6, since $x \in L$ we have that each 2D slice contributes an energy of at most $O(n^{-k})$. The total energy is thus $O(n^{r-2-k}) = O(n^{-k'})$ where we have chosen $k = r-2+k'$. This results in a final energy upper bound of $4n^r(r-1) + O(n^{-k'})$. ◀

4.2 Soundness

In this section we will prove the following lemma:

► **Lemma 9.** *Let L be a language in QMA_{EXP} and let $x \in \{0,1\}^*$ be a problem instance. Let $H = R(L, x, r)$. If $x \notin L$, then $E_0(H) \geq 4n^r(r-1) + 1$.*

Our general strategy will be to first lower bound the energy of any state with a given classical tiling. We call these “tile states” and define them as follows:

► **Definition 10 (Tile state).** *A **tile state** is a state $|\psi_{c,c'}\rangle = |c\rangle \otimes |c'\rangle \otimes |\phi\rangle$ where $c, c' \in T_1^N \times T_2^N$ and $|\phi\rangle \in (\mathcal{H}_{\sigma_1} \otimes \mathcal{H}_{\sigma_2} \otimes \mathcal{H}'_{\sigma_1} \otimes \mathcal{H}'_{\sigma_2} \otimes \mathcal{H}_{2D})^{\otimes N}$. The notation $|\psi_c\rangle$ is used whenever only one of the tilings is relevant.*

Lower bounding the energy of an arbitrary state follows straightforwardly, since each of the Hamiltonian terms is diagonal with respect to the tile Hilbert spaces.

We start by first establishing the fact that a qubit cannot be in two EPR pairs at once.

► **Fact 11.** Consider a Hamiltonian on three qubits u, v and w defined by $(\frac{\mathbb{I} - |\Phi^+\rangle\langle\Phi^+|}{2})_{u,v} + (\frac{\mathbb{I} - |\Phi^+\rangle\langle\Phi^+|}{2})_{v,w}$. By direct computation this has a minimum eigenvalue of $1/4$.

We next focus on one copy of the tilings at a time and define the following notation.

► **Definition 12.** For a given site u , let n_u be the number of u 's neighbors that are tiled with the same color as u .

With these in hand, we can now lower bound the energy of any tile state in terms of the number of same-colored neighbors of each site.

▷ **Claim 13.**

$$\langle \psi_c | \sum_{\{u,v\} \in E_{\Lambda_r(n)}} (h_{\text{tile}} + h_{\text{EPR}} + h_{\text{loop}})^{u,v} | \psi_c \rangle \geq 2n^r r - \sum_{u \in \Lambda_r(n)} n_u + 4 \sum_{u \in \Lambda_r(n)} \lfloor n_u/3 \rfloor \quad (5)$$

Proof. Recall that h_{loop} gives an energy penalty of 2 for neighboring particles that are tiled by opposite colors. Every particle has $2r$ neighbors and so if every particle is tiled with a different color than all of its neighbors then the total penalty due to this Hamiltonian term would be $2 \frac{n^r(2r)}{2} = 2n^r r$. The total number of neighboring pairs with the same color tiling is $\sum_{u \in \Lambda_r(n)} \frac{n_u}{2}$. Since each of these pairs saves an energy of 2, the total energy with respect to h_{loop} applied to each edge is $2n^r r - 2 \sum_{u \in \Lambda_r(n)} \frac{n_u}{2} = 2n^r r - \sum_{u \in \Lambda_r(n)} n_u$.

Now for every particle u we can group its same-colored neighbors into groups of three until a full group of three cannot be formed. There will be $\lfloor n_u/3 \rfloor$ such groups. For a given group of three neighbors, label the particles v, w, z . If h_{tile} applied to edge $\{u, v\}$, $\{u, w\}$, or $\{u, z\}$ is violated then this incurs a penalty of 4 per particle involved (8 in total for the edge).

If none are violated then this means at least two of the three particles, v, w, z , must be tiled with the same number. This is because h_{tile} enforces that same-colored neighbors of u are tiled with a different number than u and there are only two such numbers. Without loss of generality assume that v and w are tiled with the same number and it is exactly one higher (mod 3) than u 's number. Then we have the following bound

$$E_0(h_{\text{EPR}}^{u,v} + h_{\text{EPR}}^{u,w}) \geq E_0(A^{u,v} + A^{u,w}) \quad (6)$$

$$\geq 16E_0((\frac{\mathbb{I} - |\Phi^+\rangle\langle\Phi^+|}{2})_{u\sigma_2 v\sigma_1} + (\frac{\mathbb{I} - |\Phi^+\rangle\langle\Phi^+|}{2})_{u\sigma_2 w\sigma_1}) \quad (7)$$

$$\geq 4 \quad \text{By Fact 11} \quad (8)$$

Repeating this argument for each site u in the lattice does not double count energy penalties since we only used the $A^{u,v}$ term of h_{EPR} , and so when considering v we would use the $A^{v,u}$ term instead. Therefore, either h_{tile} or h_{EPR} must be violated and either way a penalty of at least 4 is added to the energy. This argument can be repeated for each group of three same-colored neighbors and so this contributes at least $4 \sum_{u \in \Lambda_r(n)} \lfloor n_u/3 \rfloor$ to the energy. ◁

It follows immediately from Equation (5) that each particle must have exactly two same colored-neighbors. Otherwise this induces an energy penalty of at least one which is enough to imply the bound in Lemma 9. This is helpful since we have now shown that the tiling pattern forms loops. It still remains to show that these loops are straight. Before moving on to this, we first make the above statements rigorous as follows:

► **Definition 14.** A classical tiling $c \in T_1^N \times T_2^N$ is **looped** if $n_u = 2 \forall u \in \Lambda_r(n)$

12:12 The Rotation-Invariant Hamiltonian Problem Is QMA_{EXP}-Complete

▷ **Claim 15.** If c or c' is not looped, then $\langle \psi_{c,c'} | H | \psi_{c,c'} \rangle \geq 4n^r(r-1) + 1$.

Proof. Without loss of generality assume that c is not looped and that c' is any tiling. The inequality in Eq. 5 is minimized when $n_u = 2$ for all $u \in \Lambda_r(n)$. This is because $-n_u + 4\lfloor n_u/3 \rfloor = -2$ for $n_u = 2$ and $-n_u + 4\lfloor n_u/3 \rfloor \geq -1$ for $n_u \neq 2$ where n_u is a nonnegative integer. The right-hand side of Eq. 5 is equal to $2n^r r - 2n^r = 2n^r(r-1)$ when $\forall u, n_u = 2$. Therefore, when $n_u \neq 2$ for some $u \in \Lambda_r(n)$, $\langle \psi_c | \sum_{\{u,v\} \in E_{\Lambda_r(n)}} (h_{\text{tile}} + h_{\text{EPR}} + h_{\text{loop}})^{u,v} | \psi_c \rangle \geq 2n^r(r-1) + 1$. The terms h'_{tile} , h'_{EPR} , and h'_{loop} add a penalty of at least $2n^r(r-1)$ as we have just argued. Finally, each term in H is positive semidefinite and so the energy with respect to the entire Hamiltonian is lower bounded by the energy with respect to a subset of the terms. The claim then follows. ◁

We now must show that these loops are in fact straight and do not contain any turns.

► **Definition 16.** Let $g : \Lambda_r(n)^2 \rightarrow \mathbb{N}$ be a function that outputs the number of coordinates that differ between two lattice sites.

► **Definition 17.** A classical tiling $c \in T_1^N \times T_2^N$ has a **turn** if there exists three particles u , v , and w where the following is true: u and w are neighbors of v (i.e. $d(u,v) = d(v,w) = 1$), u , v , and w are all tiled by the same color, and $g(u,w) = 2$.


We can now penalize turns by using the tiling rule that different color neighbors must have the same number. This is because if there is a turn then there exists some neighboring site outside of the loop that borders two sites right where the turn occurs. This causes a penalty because both sites in the loop must have the same number as the site outside of the loop, which contradicts the sequential numbering of the loop. This is argued in more detail below.

▷ **Claim 18.** If c or c' is looped but has a turn, then

$$\langle \psi_{c,c'} | H | \psi_{c,c'} \rangle \geq 4n^r(r-1) + 4 \quad (9)$$

Proof. Once again without loss of generality, assume that c is looped and has a turn and that c' is any tiling. If c has a turn then there exists u , v , and w that are all tiled by the same color and $g(u,w) = 2$. Without loss of generality, let $u_1 = 0$, $u_2 = 0$, $v_1 = 0$, $v_2 = 1$, $w_1 = 1$, $w_2 = 1$. This simplification is for clarity but note that the following argument works for any coordinates such that $d(u,v) = d(v,w) = 1$ and $g(u,w) = 2$. Consider a fourth site z at coordinates $z_1 = 1$, $z_2 = 0$. Note that $d(z,u) = d(z,w) = 1$ and so it neighbors u and w . This situation is depicted in Figure 4. If z is tiled by the same color as u , v , and w then it would form a loop of length four. This violates h_{tile} since the loop is not a multiple of three and so the numbers can not sequentially wrap around \mathbb{Z}_3 (see Figure 4a). If z is tiled with a different color than u , v , and w then it must be tiled with the same number as u and w since h_{tile} enforces that different colored neighbors must have the same number. However, as we have argued previously, a site cannot have two same-colored neighbors that are tiled with the same number as each other since this causes the EPR constraint to be violated (see Figure 4b). In either case, there is an energy penalty of at least 4. In addition, the h_{loop} and h'_{loop} terms together add a penalty of $4n^r(r-1)$. Finally, noting the positive semidefiniteness of each Hamiltonian term concludes the proof. ◁

So far we have given a sufficient lower bound for any tile states that do not consist of only straight loops. It will be helpful to make a few more observations on the structure of these 1D loops.



■ **Figure 4** Two examples of how energy penalties can arise when there is a turn in the loop. Here, the loop consists partially of u , v , and w which contains a turn since u and w differ in more than one coordinate. In (a), z is colored the same as the rest but this results in a penalty since a loop of size 4 cannot be numbered cyclically around \mathbb{Z}_3 . This results in the illegal configuration of u and v having the same color and the same number. In (b), z is colored differently but this leads to part of the loop being numbered as 1, 0, 1, which is also illegal as depicted in Figure 2.



■ **Figure 5** This configuration always arises if the tiling is not uniformly directed since otherwise each loop is always pointed in the same direction. This causes a rule violation since the blue and yellow tiles must have the same number but are forced to hold different numbers due to the red tiles.

► **Definition 19.** A classical tiling $c \in T_1^N \times T_2^N$ is **uniformly directed** if it is looped without turns and each 1D loop is oriented along the same dimension.

▷ **Claim 20.** If c or c' is looped without turns but not uniformly directed then

$$\langle \psi_{c,c'} | H | \psi_{c,c'} \rangle \geq 4n^r(r-1) + 8 \quad (10)$$

Proof. If the tiling is not uniformly directed then the following situation will necessarily arise, which is illustrated in Figure 5. Without loss of generality, consider a square of sites within the lattice such that the top two sites are the same color and thus a part of the same 1D chain, but the bottom two sites have two different colors from the top and from each other. Since different color tiles must have the same numberings, then all four squares would be required to have the same number. However, since the top two have the same color, they are required to have different numbers and so it is impossible to assign numbers that do not incur any penalties. ◁

One final property we will need is that each 1D loop is numbered consistently.

► **Definition 21.** A classical tiling $c \in T_1^N \times T_2^N$ is **numbered consistently** if it is uniformly directed towards a given dimension d and every site with the same d th coordinate value has the same number.

▷ **Claim 22.** If c or c' is uniformly directed, but not numbered consistently then

$$\langle \psi_{c,c'} | H | \psi_{c,c'} \rangle \geq 4n^r(r-1) + 8 \quad (11)$$

Proof. Consider the $r-1$ dimensional sublattice of sites that each have the same d th coordinate value. Each pair of neighbors in this lattice must have different colors. Otherwise, the coloring would not be uniformly directed towards the d th dimension since this means there is some 1D loop that points in a different direction. Since each pair of numbers has different colors, they all must have the same number, i.e. the tiles must be numbered consistently. Otherwise, a penalty of at least 8 is incurred. ◁

12:14 The Rotation-Invariant Hamiltonian Problem Is QMA_{EXP} -Complete

All that remains is to now lower bound the ground-state energy when both tilings are numbered consistently.

▷ **Claim 23.** Let c and c' be numbered consistently. Let L be a language in QMA_{EXP} and let $x \in \{0,1\}^*$ be a problem instance. Let $H = R(L, x, r)$. If $x \notin L$, then $\langle \psi_{c,c'} | H | \psi_{c,c'} \rangle \geq 4n^r(r-1) + \Omega(n^{r-5})$.

Proof. First, we must handle the case where c and c' both have 1D chains pointing in the same direction. This would incur a penalty of $n^r r$ from the h_{copy} term alone, which would clearly imply the desired lower bound. Next, we focus on the case where they point in different directions. We let the 1D chains in the first tiling represent the horizontal rows of each 2D slice and those of the second tiling represent vertical rows. In addition, we let the order of the numberings define the left, right, up and down directions of the slices. In this way we can embed n^{r-2} 2D translation-invariant Hamiltonians. Since $x \notin L$, by Theorem 6, these terms contribute an energy penalty of $\Omega(n^{r-5})$. ◁

To complete the proof it remains to deal with the non-tile states but these can easily be handled since the Hamiltonian is diagonal in the tile Hilbert spaces.

► **Lemma 24** (Restatement of Lemma 9). *Let L be a language in QMA_{EXP} and let $x \in \{0,1\}^*$ be a problem instance. Let $H = R(L, x, r)$. If $x \notin L$, then $E_0(H) \geq 4n^r(r-1) + 1$.*

Proof. First we can write an arbitrary state as a superposition of tile states: $|\zeta\rangle = \sum_i \alpha_i |c_i\rangle |\phi_i\rangle$ where $c_i \in T_1^{2N} \times T_2^{2N}$. Note that $\langle c_i | \langle \phi_i | H | c_j \rangle | \phi_j \rangle = 0$ for $c_i \neq c_j$. This is because all terms are diagonal on the Hilbert space of the classical tiles. Therefore, we have $\langle \zeta | H | \zeta \rangle = \sum_i |\alpha_i|^2 \langle c_i | \langle \phi_i | H | c_i \rangle | \phi_i \rangle$. This is an affine combination of tile state energies which we have already lower bounded by $4n^r(r-1) + 1$. Thus, the energy itself is also bounded from below by $4n^r(r-1) + 1$. ◀

5 Open boundary conditions

So far we have only considered the case where the Hamiltonian has periodic boundary conditions. It turns out that the same construction also works for open boundary conditions. Our method of embedding directed stripes still works in this case except now instead of closed loops the stripes form spin chains with open boundary conditions. This leaves the sites at the ends of the chain with one unpaired qubit that can still form an EPR pair with another site; however, this would require introducing a turn. The energy penalty for having a turn outweighs the energy bonus of having one more same-colored neighbor and so it is optimal to leave the qubit unpaired. Thus, the same construction can once again be used to embed a 2D translation-invariant non-reflection-invariant Hamiltonian with the exception that this Hamiltonian now has open boundary conditions. To complete the result, it remains to show the equivalent of Theorem 6 in the case of open boundary conditions. The proof of this statement is shown in Section A.2.

6 Conclusion

In this work, we have resolved the complexity for rotation-invariant Hamiltonians with constant lattice dimension, but it still remains interesting to better understand the complexity at even higher lattice dimensions. For instance, we know that as $r \rightarrow \infty$ the ground state becomes a product state, but how fast does it converge? Consider the problem where the lattice length is now fixed, and the lattice dimension is given as input. Is there a small

enough promise gap for which this problem is quantumly hard? Another direction to consider is to study more general permutation symmetries. In some sense, the rotation-invariant Hamiltonian problem interpolates between systems with comparatively low symmetry in the one-dimensional case to highly symmetric systems as the lattice dimension increases. It would be an interesting question to generalize this interpolation and probe whether there exists a complexity phase transition with respect to some symmetry parameter.


References

- 1 Dorit Aharonov and Lior Eldar. Commuting local hamiltonians on expanders, locally testable quantum codes, and the qpcp conjecture, 2013. [arXiv:1301.3407](#).
- 2 Dorit Aharonov, Daniel Gottesman, Sandy Irani, and Julia Kempe. The power of quantum systems on a line. *Communications in Mathematical Physics*, 287(1):41–65, January 2009. doi:10.1007/s00220-008-0710-3.
- 3 Johannes Bausch, Toby S. Cubitt, and James D. Watson. Uncomputability of phase diagrams. *Nature Communications*, 12(1), January 2021. doi:10.1038/s41467-020-20504-6.
- 4 Johannes Bausch and Stephen Piddock. The complexity of translationally invariant low-dimensional spin lattices in 3d. *Journal of Mathematical Physics*, 58(11), November 2017. doi:10.1063/1.5011338.
- 5 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, pages 871–880, New York, NY, USA, 2013. Association for Computing Machinery. doi:10.1145/2488608.2488719.
- 6 Daniel Gottesman and Sandy Irani. The quantum and classical complexity of translationally invariant tiling and hamiltonian problems. *Theory of Computing*, 9(2):31–116, 2013. doi:10.4086/toc.2013.v009a002.
- 7 A. Yu. Kitaev, A. H. Shen, and M. N. Vyalyi. *Classical and Quantum Computation*. American Mathematical Society, USA, 2002.
- 8 Tamara Kohler and Toby Cubitt. Translationally Invariant Universal Classical Hamiltonians. *Journal of Statistical Physics*, 176(1):228–261, July 2019. doi:10.1007/s10955-019-02295-3.
- 9 Christina V. Kraus, Maciej Lewenstein, and J. Ignacio Cirac. Ground states of fermionic lattice hamiltonians with permutation symmetry. *Physical Review A*, 88(2), August 2013. doi:10.1103/physreva.88.022335.
- 10 Roberto Oliveira and Barbara Terhal. The complexity of quantum spin systems on a two-dimensional square lattice. *Quantum information & computation*, 8, May 2005. doi:10.26421/QIC8.10-2.
- 11 Stephen Piddock and Johannes Bausch. Universal translationally-invariant hamiltonians, 2020. [arXiv:2001.08050](#).
- 12 Ch. Rickwardt, P. Nielaba, and K. Binder. A finite size scaling study of the five-dimensional ising model. *Annalen der Physik*, 506(6):483–493, 1994. doi:10.1002/andp.19945060606.
- 13 Gert Sabidussi. Graphs with given group and given graph-theoretical properties. *Canadian Journal of Mathematics*, 9:515–525, 1957. doi:10.4153/CJM-1957-060-7.
- 14 B. M. Terhal. Is entanglement monogamous? *IBM Journal of Research and Development*, 48(1):71–78, January 2004. doi:10.1147/rd.481.0071.

A Constructing the 2D translation-invariant Hamiltonian

A.1 Periodic boundary conditions

In this section we prove the following theorem:



■ **Figure 6** An example of an allowed configuration for a given set of translation-invariant tiling rules on a 2D grid defined in [6]. The key feature is that when the side length is a prime number, exactly one row can contain the H and V tiles and one site within this row can contain the HV tile. The above image is directly reproduced from [6] under CC-BY 3.0.

► **Theorem 25** (Restatement of Theorem 6). *Let L be a QMA_{EXP} -complete language. There exists an efficiently computable function $f : \{0, 1\}^* \rightarrow \mathbb{Z}$ and 2-local positive semidefinite Hamiltonian terms h_{TI} and v_{TI} with the following properties:*


1. *$f(x)$ is a multiple of 3 and $f(x)/3$ is prime. Furthermore, $\log f(x) = O(\text{poly}(|x|))$ and f is computable in $O(\text{poly}(|x|))$ time.*
2. *Let $f(x) = n \geq n_0$ for some constant n_0 and a given problem instance $x \in \{0, 1\}^*$. For an $n \times n$ 2D lattice with periodic boundary conditions, let E_h be the set of ordered pairs of horizontal neighbors and let E_v be the set of ordered pairs of vertical neighbors. Consider the Hamiltonian $H_{TI} = \sum_{(u,w) \in E_h} h_{TI}^{u,w} + \sum_{(x,y) \in E_v} v_{TI}^{x,y}$.*
 - a. *If $x \in L$ then $E_0(H_{TI}) \leq O(n^{-k})$ for an arbitrarily large constant k .*
 - b. *If $x \notin L$ then $E_0(H_{TI}) \geq \Omega(1/n^3)$*

Much of the proof will directly utilize techniques from [6]. When needed, a brief summary of these ideas is given but we direct the interested reader to [6] for the full details. The main idea of the construction is to use the 2D translation-invariant tiles to embed a 1D chain with a designated starting tile. This allows us to avoid the $1/2$ additive penalty that is required in the 1D construction when there is no designated starting tile (see section 6 “The quantum cycle” of [6]). Fortunately, the construction in tables 5 and 6 of section 4.1 of [6] accomplish exactly this. This construction is quite complicated and so we only present the end result. The last piece to handle is that this construction requires the grid length to be prime while our tiling rules require it to be a multiple of three. This can easily be remedied by simulating each tile in the 2D construction with a 3×3 grid of nine tiles. We explain each of these steps in more detail below.

We start by describing how to construct the function f referenced in the theorem statement. [6] show that given x there is a randomized algorithm a running in expected time $O(\text{poly } |x|)$ to find a prime number p such that the $1/3$ most significant digits represent x and $\log p = O(\text{poly}(|x|))$. Using this result, we simply define f as $f := 3 \cdot a(x)$.

Next, to construct the 2D translation-invariant Hamiltonian, we start by using the below result.

► **Lemma 26** (see section 4.1 of [6]). *There exists a set of translation-invariant, non-reflection-invariant tiling rules involving a constant number of total tile types on an $n \times n$ 2D grid with periodic boundary conditions such that when n is prime exactly one row must contain tiles of the form H and V and no other row can contain either of these tiles. In addition, exactly one site in this row must contain HV and all other sites in this row must contain H . This configuration is depicted in Figure 6.*



■ **Figure 7** This is the only allowed configuration for the subtile associated with each 3×3 grid. Each such grid represents a tile in the original tiling system.

Since this construction only works for prime n but our lattice length must be a multiple of three, it is necessary to replace each tile with a 3×3 grid of tiles that serve the same function as the original. First, for each tile in the original tiling, a corresponding center tile is defined as ■. Then, the following tiles and rules are added. ■ must be placed above ■. Similarly, ■ must be placed below ■, ■ placed to the left of it and ■ placed to the right of it. Now to fix the corners in place we enforce that ■ must be placed to the left of ■, ■ placed to the right of ■, ■ placed to the left of ■ and ■ placed to the right of ■. In other words, the center tile must always be surrounded by the border tiles. The reciprocal of each of these rules is also included. This enforces that the border tiles must always be accompanied by the center tile in the appropriate location. This results in Figure 7 being the only allowed configuration. The last thing to resolve is to ensure that these 3×3 grids are aligned with each other. To do this, we must regulate which border tiles can be placed next to those of a different 3×3 grid. We enforce that only ■ type tiles are allowed to the left of ■ type tiles. We incorporate the same rule for the other sides as well: only ■ type tiles are allowed to the right of ■ type tiles, only ■ type tiles are allowed above ■ type tiles and only ■ type tiles are allowed below ■ type tiles.


Now we can incorporate an original rule between tiles by applying it to their corresponding border tiles. This will exactly simulate the original but with each site replaced by a 3×3 grid. This results in an $n \times n$ grid where $n = 3p$ and p is a prime number.

With this tiling in hand, the 1D construction can now be embedded into the 3×3 grids associated with the ■ and ■ tiles. In particular, only the middle row of the 3×3 grids associated with each ■ and ■ tile is used for the chain (i.e. only the tiles ■, ■, and ■). In addition, the ■ tile associated with the ■ tile is used to mark the left endpoint of the 1D construction.

It remains to construct a 1D translation-invariant Hamiltonian on a $f(x)$ -length spin chain with the desired ground-state energy properties. To accomplish this [6] is directly used and is briefly summarized here. The main idea is to use the Hamiltonian terms to simulate a quantum Turing machine where each site in the length- $f(x)$ spin chain represents a different cell of the Turing machine tape except for two sites which are used to mark the boundaries. The goal of the first part of the Turing machine is to infer x from the length of the tape and write it on the tape. The second part of the Turing machine is quantum and uses x as the input along with a quantum witness to execute the QMA_{EXP} verification algorithm for L .

We now discuss the first part of the Turing machine, which we denote as M_{BC} . M_{BC} is a purely classical Turing machine implementing a binary counter. By incrementing a clock pointer from one side of the tape to the other, we can ensure that M_{BC} is run for exactly $f(x) - 3$ steps. Recall that $f(x) = 3p$ where $1/3$ of the most significant digits of p is x . Therefore, by using a sufficiently slow binary counter, it is possible to ensure that x is always written on the tape at the end of the $f(x) - 3$ steps.

Now that x is on the tape, we can run the quantum verification algorithm on x along with an arbitrary quantum witness. The verifier is also allowed a total of $f(x) - 3$ timesteps. Notice this is more strict than QMA_{EXP} , which allows the verifier $2^{\text{poly } |x|}$. QMA_{EXP} can be



■ **Figure 8** The only allowed configuration for the given set of translation-invariant tiling rules on a 2D grid with open boundary conditions.

reduced to this case by a standard padding argument where x is padded by zeros to have length $\text{poly}(|x|)$. Finally, an energy penalty is applied if the verifier does not accept. If the verifier accepts with probability $1 - \epsilon$ in the $x \in L$ case then the ground-state energy will be upper bounded by ϵ/n^2 . In order to drive $\epsilon \leq O(n^{-k})$ for an arbitrarily large constant, it is possible to use witness amplification. This incurs a $O(k \log n)$ overhead in the verifier's runtime [7], which can easily be accommodated by padding. In our construction, we would like to set $k = O(r)$ where r is the dimension, but would prefer not to have the verification algorithm depend on r . Therefore, we can instead set $k = \log n$ where $\log n \geq O(r)$ for some $n \geq n_0$ since r is a parameter of the problem and does not scale with n . Importantly, even though the number of rounds of witness amplification depends on n , our Hamiltonian term still does not depend on n since n is deduced from the length of the lattice and then given as input to the verification algorithm. For this construction [6] also show that in the $x \notin L$ case the ground-state energy is lower bounded by $\Omega(1/n^3)$.

Finally, each term in the construction is of one of two forms called Type I terms: $|ab\rangle\langle ab|$ and Type II terms:

$$\frac{1}{2}(|ab\rangle\langle ab| + |cd\rangle\langle cd| - |ab\rangle\langle cd| - |cd\rangle\langle ab|). \quad (12)$$

Both types are positive semidefinite and so the overall Hamiltonian term is also positive semidefinite.

A.2 Open boundary conditions

An equivalent theorem to Theorem 6 is also true in the case of open boundary conditions. The construction is also inspired by [6]. First, we define the tiles \boxtimes , \boxminus , \boxdot , and \square . The idea is then to introduce the following tiling rules: no tile is allowed to the left of or below \boxtimes , no tile is allowed below \boxminus and no tile is allowed to the right of or below \boxdot . Additionally, the only tile that is allowed to the right of \boxtimes is \boxminus and the only tile that is allowed to the left of \boxdot is also \boxminus . Finally, \square is not allowed to the left or right of \boxminus and the only tile allowed above \boxtimes , \boxminus and \boxdot is \square . This results in the configuration depicted in Figure 8. The 1D Hamiltonian can then be embedded into the bottom row of the 2D grid where the \boxtimes and \boxdot tiles denote the endpoints.