Computational Complexity of the Ground State Energy Density Problem

James D. Watson
Toby S. Cubitt
ucapjdj@ucl.ac.uk
t.cubitt@ucl.ac.uk
Department of Computer Science, University College London
London, UK

ABSTRACT

We study the complexity of finding the ground state energy density of a local Hamiltonian on a lattice in the thermodynamic limit of infinite lattice size. We formulate this rigorously as a function problem, in which we request an estimate of the ground state energy density to some specified precision; and as an equivalent promise problem, GSED, in which we ask whether the ground state energy density is above or below specified thresholds.

The ground state energy density problem is unusual, in that it concerns a single, fixed Hamiltonian in the thermodynamic limit, whose ground state energy density is just some fixed, real number. The only input to the computational problem is the precision to which to estimate this fixed real number, corresponding to the ground state energy density. Hardness of this problem for a complexity class therefore implies that the solutions to all problems in the class are encoded in this single number (analogous to Chaitin's constant in computability theory).

This captures computationally the type of question most commonly encountered in condensed matter physics, which is typically concerned with the physical properties of a single Hamiltonian in the thermodynamic limit. We show that for classical, translationally invariant, nearest neighbour Hamiltonians on a 2D square lattice, $\mathsf{P}^{\mathsf{NEEXP}} \subseteq \mathsf{EXP}^{\mathsf{GSED}} \subseteq \mathsf{EXP}^{\mathsf{NEXP}}$, and for quantum Hamiltonians $\mathsf{P}^{\mathsf{NEEXP}} \subseteq \mathsf{EXP}^{\mathsf{GSED}} \subseteq \mathsf{EXP}^{\mathsf{QMA}_{\mathsf{EXP}}}$. With some technical caveats on the oracle definitions, the EXP in some of these results can be strengthened to PSPACE. We also give analogous complexity bounds for the function version of GSED.

CCS CONCEPTS

• **Theory of computation → Problems, reductions and completeness**; **Complexity classes**; *Quantum complexity theory*.

KEYWORDS

complexity, Hamiltonian complexity, physics, energy, quantum, precision, condensed matter physics

STOC '22, June 20–24, 2022, Rome, Italy

ACM ISBN 978-1-4503-9264-8/22/06…$15.00
https://doi.org/10.1145/3519935.3520052

ACM Reference Format:
James D. Watson and Toby S. Cubitt. 2022. Computational Complexity of the Ground State Energy Density Problem. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing (STOC '22), June 20–24, 2022, Rome, Italy.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3519935.3520052

1 INTRODUCTION

The connection between computational complexity theory and many-body physics dates back over 40 years. Barahona's [4] proof of NP-completeness of the ground state energy problem for classical many-body models with local interactions[1] — or "local Hamiltonians" for short — on a finite number of particles (spins), established the ground state energy as one of the canonical physical quantities for which computational complexity yields insight.

The *Hamiltonian* is the function mapping a state of the particles to its corresponding energy. The ground state is then the minimum-energy state of the system, and the ground state energy the associate minimum-energy. The problem of estimating the ground state energy is often formulated as an equivalent (up to polynomial-time computation) decision problem known as the Local Hamiltonian problem: given a Hamiltonian and an energy threshold, decide whether the ground state energy is above or below that threshold.

Nearly 20 years later, Kitaev [25] proved QMA-completeness (the quantum analogue of NP-completeness) for quantum local Hamiltonians on a finite number of quantum particles. There has been a plethora of papers following — too many to comprehensively list here — building on Barahona and Kitaev's seminal results. These have extended hardness of the ground state energy problem to ever more restrictive classes of Hamiltonian, with specific, physically-motivated types of local interaction, and with restricted patterns of local interaction. In particular, amongst many other related results, we now know that the classical and quantum ground state energy problems remain NP- and QMA-complete when restricted to nearest-neighbour interactions on a finite 2D square lattice and a finite 1D chain, respectively [1, 4]. Properties beyond the ground state energy have been studied, including density of states [9], expectation values on low energy subspaces [3], the energy of excited states [23], detecting energy barriers [18], determining whether a system is frustrated, and many others [16].

The input to all of the above problems is a description of a local Hamiltonian on a finite number of particles, and the complexity-theoretic hardness is a function of varying the Hamiltonian. However, many-body and condensed matter physicists are more often

[1] Namely, the 2D Ising model with fields.

interested in properties of a many-body system in the *thermodynamic limit* of infinitely many particles. Most many-body physics properties, such as phase transitions, phase diagrams, spectral gaps, etc., are only well-defined theoretically in this limit. Moreover, in experimental physics, these models often arise as idealisations of physical materials, where a typical sample will contain such a large number of atoms that the properties of the material are well-approximated by the infinite limit.

Furthermore, they are typically interested in computing the physical properties of a *single* Hamiltonian – or a family of Hamiltonians parametrised by a small, constant number of parameters. Often, the local interactions have some regular structure, such as translational invariance where all the local interactions take the same form. The standard formulation of the ground state energy problem does not capture this type of question.

1.1 Related Work

There are a small number of results proving hardness of estimating the ground state energy for a translationally invariant Hamiltonian where the local interaction is fixed, and the only input to the problem is the lattice size. Here, since a lattice of size L can be specified in $\lceil \log(L) \rceil$ bits, the natural complexity class is NEXP (or $\mathsf{QMA}_{\mathsf{EXP}}$ in the quantum case), rather than NP. The Wang tiling completion problem is known to be NEXP-complete [20, 30], which can trivially be translated to the ground state energy problem for a single, fixed, translationally invariant, nearest-neighbour, classical Hamiltonian on a 2D square lattice, where the state at some of the boundaries is fixed (fixed boundary conditions). As the interaction is fixed, the only remaining problem input is the size of the lattice. Remarkably, this provides a degree of freedom in which one can encode computational problems, and so suffices for the hardness result. Gottesman and Irani [20] also extended these results to more natural types of boundary condition. They went on to prove the analogous $\mathsf{QMA}_{\mathsf{EXP}}$-completeness result for quantum Hamiltonians on a 1D chain. However, these results still concern Hamiltonians on finite numbers of particles; indeed, the problem input is the number of particles the Hamiltonian acts on.

In the thermodynamic limit, the ground state energy is no longer a meaningful quantity; it typically has infinite magnitude, and is not physically measurable. In this setting, the more relevant quantity is the ground state energy *density*: the minimum energy *per particle*. Just as the ground state energy is a key starting point for studying the physics of finite many-body systems, the ground state energy density (GSED) is a key starting point for physics in the thermodynamic limit. Methods of approximating the ground state energy density in condensed matter systems have been the subject of much study in the physics literature [22, 31].

Less is known about the computational complexity of the ground state energy density problem, than for the ground state energy. Gottesman and Irani [20] proved that the ground state energy density problem for translationally invariant, nearest-neighbour, quantum Hamiltonians on a 1D chain with a $\Omega(1/2^n)$ promise-gap is NEXP-complete. Here, the input is a description of the local interactions of the system, and the complexity is a function of varying the Hamiltonian. Meanwhile, as a stepping stone to their undecidability result for the spectral gap, [12, 13] proved that deciding whether the ground state energy density is 0 or strictly positive, with no promise gap, is undecidable. Their result holds for quantum, translationally invariant, nearest neighbour Hamiltonians on a 2D square lattice with a fixed local dimension. [5] later extended this undecidability result to 1D chains (again as a stepping stone to the spectral gap problem) and [7] extends to 2D systems for which the local interactions are analytic in the input parameter.

However, as with most ground state energy complexity results, these results still have as input the description of the Hamiltonian, and the hardness is a result of varying the Hamiltonian.

1.2 The Ground State Energy Density Problem

If we restrict to a single, fixed Hamiltonian in the thermodynamic limit, it may seem that there are no input parameters left, and complexity theory can have nothing to say! However, this is not quite the case. We can still ask about the complexity of estimating the ground state energy density to a given precision, where the only input is the precision required. (See section 2 for precise problem definitions.) Arguably, this is the problem formulation closest to that often encountered in condensed matter physics.

If we learn the ground state energy density to precision 2^{-n}, then we can hope to learn the first n bits of its binary representation. An n bit string can encode the solutions to at most n different decision problems. But an index into this bit string, specifying the index of the decision problem we are interested in, requires only $\log n$ bits. Therefore, the natural complexity class for GSED is NEEXP, or related doubly-exponential time complexity classes. (At least for hardness results.)

In this work, drawing on techniques developed in [13], we prove upper and lower bounds on the complexity of this Ground State Energy Density (GSED) problem: we show that GSED is NEEXP-hard under exponential-time Turing reductions, and contained in $\mathsf{EXP}^{\mathsf{NEXP}}$. In fact, we prove the following slightly stronger results for the natural promise-problem formulation of GSED, for a fixed, classical, translationally invariant, nearest-neighbour Hamiltonian on a 2D square lattice:

$$\mathsf{P}^{\mathsf{NEEXP}} \subseteq \mathsf{EXP}^{\mathsf{GSED}} \subseteq \mathsf{EXP}^{\mathsf{NEXP}}$$

The natural promise-problem formulation of GSED takes as input two energy density thresholds α and β with $\beta - \alpha = \Omega(2^{-n})$, and outputs whether the ground state energy density is above β or below α.

The analogous complexity bounds for the function problem formulation of GSED readily follow from this. For the function problem formulation, the input is the precision ϵ, and the output is an estimate of the ground state energy density to precision ϵ.

For quantum Hamiltonians, a very similar argument to the classical case establishes the analogous upper bound of $\mathsf{EXP}^{\mathsf{QMA}_{\mathsf{EXP}}}$ for the quantum GSED problem. The same lower bound as above follows trivially from the fact that classical Hamiltonians are a special case of quantum. However, we are not able to prove $\mathsf{QMA}_{\mathsf{EEXP}}$-hardness of the quantum problem. (We comment on this briefly in section 6.)

The ground state energy density of the specific Hamiltonian we construct is a single, real number $\mathcal{E}_\rho$. Our hardness results imply the solutions to *all* instances of NEEXP-complete problem are encoded in the digits of this single number, with successive digits

of $\mathcal{E}_\rho$ giving the solution to successive instances of a canonical NEEXP-complete problem. In this sense, the ground state energy density of this Hamiltonian is somewhat reminiscent of Chaitin's constant [11], but encoding solutions to problems in a certain complexity class, rather than the Halting problem.

In a similar vein to our result, [14] considered the complexity of computing the n^{th} bit of an algebraic number and determined some containment results, but non-trivial bounds are not currently known. Finally we note that similar but complimentary results were shown simultaneously with this work by [2].

2 MAIN RESULTS

Define the energy density of the finite lattice as

Definition 2.1 (Ground State Energy Density). Consider a translationally invariant Hamiltonian defined on a rectangular lattice with length L, width W. We denote this Hamiltonian as $H^{\Lambda(L\times W)}$. The *ground state energy density* is defined as

$$\mathcal{E}_\rho(L,W) := \frac{\lambda_0(H^{\Lambda(L\times W)})}{LW}. \tag{1}$$

The *thermodynamic limit* of the ground state energy density is defined as the limiting value as the lattice width and height are taken to infinity:

$$\mathcal{E}_\rho := \lim_{L,W\to\infty} \mathcal{E}_\rho(L,W). \tag{2}$$

If the ground state energy density is referred to without qualification, then it is referring to the thermodynamic limit case.

This limit is well defined [13]. We now consider some useful definitions for the computational problems. For all these definitions we will be referring to the infinite lattice case.

We can cast the problem of finding $\mathcal{E}_\rho$ as a computational promise problem similar in spirit to the local Hamiltonian problem:

*Definition 2.2 (Ground State Energy Density (*GSED*) promise problem).*
Problem Parameters: A fixed, translationally invariant, nearest-neighbour Hamiltonian on a $2D$ infinite square lattice of d-dimensional spins.
Input: Two real numbers β and α, such that $\beta - \alpha = \Omega(2^{-p(n)})$, for some integer n and polynomial $p(n)$.
Output: Determine whether $\mathcal{E}_\rho > \beta$ (No instance) or $\mathcal{E}_\rho < \alpha$ (YES instance).
Promise: The ground state energy density does not lie between in the interval $[\alpha, \beta]$.

This is perhaps more naturally thought of in terms of the corresponding function problem:

*Definition 2.3 (Ground State Energy Density (*FGSED*) function problem).*
Problem Parameters: A fixed, translationally invariant, nearest-neighbour Hamiltonian acting on a $2D$ infinite lattice of d-level spins.
Input: An error bound ϵ, specified in binary.
Output: An approximation to the ground state energy density, $\tilde{\mathcal{E}_\rho}$ such that $|\mathcal{E}_\rho - \tilde{\mathcal{E}_\rho}| \leq \epsilon$.

The promise and function problems are equivalent up to log-space computation, by standard binary search arguments.

We will often restrict GSED in definition 2.2 to classical Hamiltonians, rather than general (quantum) Hamiltonians. When we wish to highlight this distinction, we refer to these as *classical* GSED and *quantum* GSED, respectively.

The main results of this work are as follows:

Theorem 2.4. $\mathsf{P}^{\mathsf{NEEXP}} \subseteq \mathsf{EXP}^{\mathsf{GSED}} \subseteq \mathsf{EXP}^{\mathsf{NEXP}}$ *for classical* GSED.

Here NEEXP is defined analogously with NP, but the verifying TM is allowed doubly exponential time to run and the witness can be doubly exponentially long. We expect that the $\mathsf{EXP}^{\mathsf{NEXP}}$ upper bound presented here is tight and there is potentially room to improve the lower bound. The above theorem implies:

Corollary 2.5. GSED *is* NEEXP*-hard under exponential time Turing reductions, for a classical, translationally invariant, nearest-neighbour Hamiltonian.*

We also prove:

Theorem 2.6. *Classical* GSED $\in$ NEXP.

Corollary 2.5 and theorem 2.6 are not in conflict with each other. Allowing exponential-time Turing reductions (as opposed to the polytime Turing reductions usually considered) allows exponentially harder problems to be solved.

The fact we are considering $\mathsf{EXP}^{\mathsf{GSED}}$ rather than GSED with polytime reductions is fundamental to the problem being about estimating the ground state energy density for *a particular Hamiltonian*, where the problem instances differ only in the precision to which that same ground state energy density should be computed (rather than each problem instance corresponding to a different Hamiltonian). We show that, using our hardness construction, one should not expect $\mathsf{NP} \subseteq \mathsf{P}^{\mathsf{GSED}}$ unless the polynomial hierarchy collapses to Σ_2^p.

We can also consider the case of quantum Hamiltonians:

Theorem 2.7. $\mathsf{P}^{\mathsf{NEEXP}} \subseteq \mathsf{EXP}^{\mathsf{GSED}} \subseteq \mathsf{EXP}^{\mathsf{QMA}_{\mathsf{EXP}}}$ *for quantum* GSED.

For the function problem, one readily obtains the corresponding complexity bounds:

Theorem 2.8. FGSED $\in \mathsf{FP}^{\mathsf{NEXP}}$ *for classical* FGSED.

We also get the bound

Lemma 2.9. $\mathsf{FP}^{\mathsf{NEEXP}} \subseteq \mathsf{FEXP}^{\mathsf{FGSED}} \subseteq \mathsf{FEXP}^{\mathsf{NEXP}}$*, for* FGSED *for a fixed classical, translationally invariant, nearest neighbour Hamiltonian.*

3 PRELIMINARIES

Let $\mathcal{B}(\mathcal{H})$ be the space of bounded linear operators on a complex Hilbert space $\mathcal{H}$. Define $\Lambda(L\times W) := \{1,\ldots,L\}\times\{1,\ldots,W\}$ to be the rectangular lattice of length L, width W, with $L, W \in \mathbb{N}$. We attach to each site $i \in \Lambda(L\times W)$ in the lattice a Hilbert space $\mathcal{H}_i \cong \mathbb{C}^d$. Given a string $x \in \{0,1\}^n$, then $|x| = n$ will denote the binary length of the string. For a given Hamiltonian H, we will denote its eigenvalues as $\lambda_i(H)$, such that $\lambda_0(H) \leq \lambda_1(H) \leq \lambda_2(H) \leq \ldots$.

Given a lattice $\Lambda(L\times W)$, a Hamiltonian $H = \sum_i h_i$ is *nearest-neighbour* if $h_i \in \mathcal{B}(\mathbb{C}^d \otimes \mathbb{C}^d)$ such that each h_i acts non-trivially

only on neighbouring pairs of lattice sites. We write the interaction between neighbouring sites as $h_{\langle i,j\rangle}$. Furthermore, *translational invariance* implies $h_{\langle i,j\rangle} = h \in \mathcal{B}(\mathcal{H})$ for any i, j. We allow the vertical and horizontal interactions of translationally invariant Hamiltonians to be different, and this is needed for our construction. By a *classical Hamiltonian*, we mean a Hamiltonian which is diagonal in the standard basis. To distinguish general Hamiltonians from classical Hamiltonians we will often call them *quantum Hamiltonians*.

We now define the relevant complexity classes:

Definition 3.1. NEXP *or* NEXPTIME

A language L is in NEXP if there exists a positive constant k and a deterministic Turing Machine M such that for each instance x and a classical witness w such that $|w| = O(2^{|x|^k})$, on input (x, w), M halts in $O(2^{|x|^k})$ steps and

- if $x \in L$, $\exists w$ such that M accepts (x, w) with probability 1.
- if $x \notin L$ then $\forall w$, M accepts (x, w) with probability 0.

We will also be concerned with the doubly exponential version of NEXP, defined as:

Definition 3.2. NEEXP *or* N2EXP

A language L is in NEEXP if there exists a positive constant k and a deterministic Turing Machine M such that for each instance x and a classical witness w such that $|w| = O(2^{2^{|x|^k}})$, on input (x, w), M halts in $O(2^{2^{|x|^k}})$ steps and

- if $x \in L$, $\exists w$ such that M accepts (x, w) with probability 1.
- if $x \notin L$ then $\forall w$, M accepts (x, w) with probability 0.

We also define $\mathrm{QMA}_{\mathrm{EXP}}$ and $\mathrm{QMA}_{\mathrm{EEXP}}$ the same way as the class QMA, but allowing for an exponentially and doubly-exponentially long witness and circuit runtime respectively.

4 TILING PRELIMINARIES

Wang tilings will play a central role in this work.

Definition 4.1 (Wang Tiles). *Wang tiles* are unit length square tiles with markings on each of the four edges. For a given set of Wang tiles $\{t_i\}_{i=1}^n$, the markings define horizontal matching rules $\mathcal{R}_{Horz}$ (respectively, vertical matching rules $\mathcal{R}_{Vert}$) such that two tiles t_i, t_j can only be placed next to each other horizontally (vertically) if $(t_i, t_j) \in \mathcal{R}_{Horz}$ ($(t_i, t_j) \in \mathcal{R}_{Vert}$).

We now consider specific sets of Wang tiles that we will employ throughout this work.

4.1 Robinson Tiles

Robinson's tiling [32] is based on a set of tiles with the rule that one tile can be placed next to another only if the arrow heads on the first tile correctly join with the arrow tails on the adjacent tile. I.e. the tiling rules enforce the condition that *arrow heads on one tile must meet arrow tails of the same type on its neighbour in the appropriate direction.* We will use a set of tiles augmented with certain additional markings, described in [13, 32], and which are shown in figure 1.

The lefter-most tile in figure 1 has arrows on all sides of the tile and is known as a *cross* and in this depiction is said to face up and to the right. The other 6 tiles are known as *arms*. Each of the arms has a principal arrow across the centre of the tile and which indicates its direction (all the tiles depicted in figure 1 are facing downwards). Arrow markings can be either red or green. On a given arm the horizontal and vertical arrows must have different colours and on cross tiles we force all arrow markings to have the same colour. The Robinson tile set includes all rotations and reflections of these basic tiles.

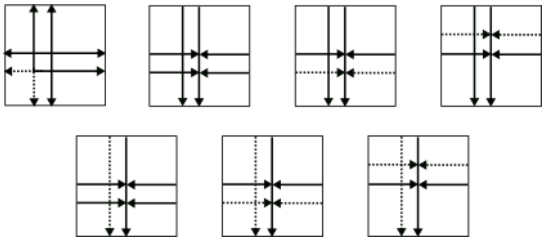

Figure 1: The modified Robinson tiles we will use.

When these tiles are used to tile a grid, the tiling rules force a pattern of interlocking, nested squares to form in any valid tiling of the plane (see figure 2(c)). The series of squares have side lengths $3, 5, 9, 17, 33, \ldots, 2^n + 1$, for $n \in \mathbb{N}$ (see figure 3). Robinson adds additional coloured markings to the tiles, such that for odd n the borders formed by the double-arrow tile markings running along the edges of the squares are green, and for even n they are red. We direct the reader to [32] and [13] for more detailed discussions of the tiling pattern and how it is formed.

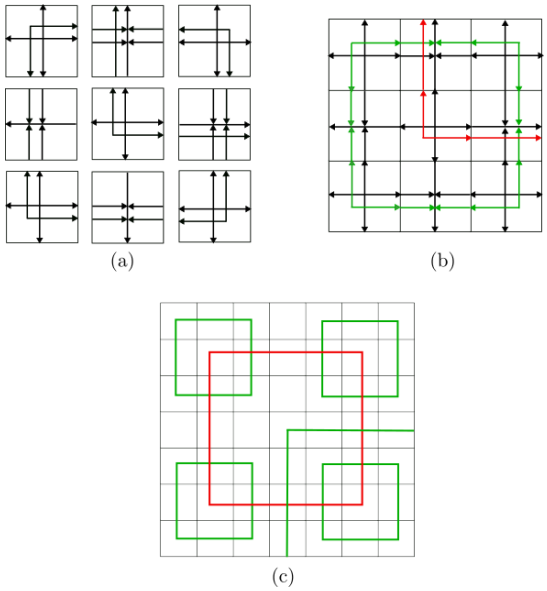

Figure 2: Various configurations of Robinson tiles: the dashed markings seen in figure 1 have been removed for clarity. (a) shows a possible tiling arrangement to create the lowest level of green squares in the Robinson tiling. (b) shows the same square once the coloured arrows have been introduced. (c) shows only the coloured markings which form a 1-square.

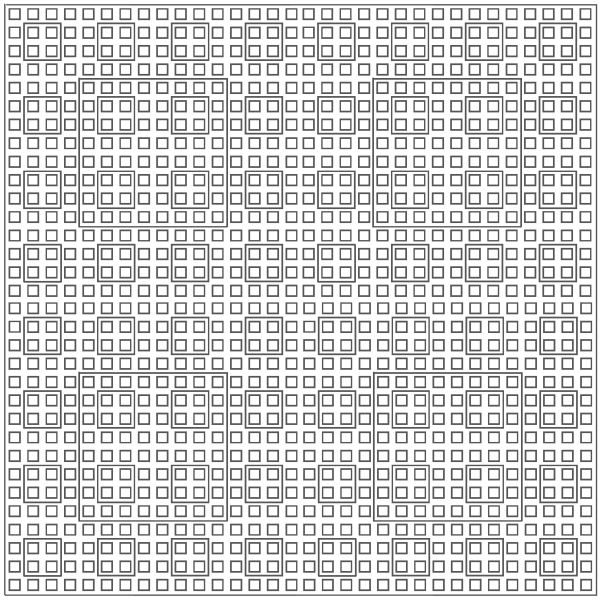

Figure 3: A Robinson tiling pattern showing only markings associated with red borders.

For our purposes we will focus on red borders, and refer to these as just *borders*. The interior of the border is referred to as a *square*. We refer to a red border of side length $4^n + 1$ as an *n-border*, and an n-border plus its interior as an n-square.

4.2 Encoding Turing Machines with Tiles

It is well known that the evolution of a classical Turing Machine can be encoded as a set of Wang tiles [8, 32]. To see this, consider a particular TM. The TM tape at a particular time step is a set of tape cells with symbols written in them, where one particular cell has the TM head over it. The TM will then evolve deterministically according to its transition rules.

Now consider an $L \times L$ tiling grid. It is possible to construct a set of Wang tiles such that the tiling pattern simulates the TM's evolution for L steps. The tile set is chosen to be tiles with all possible combinations of Turing Machine tape cell markings, plus TM head and state markings. The evolution of the TM can then be encoded as a tiling of a square lattice, where rows of tiles represent the configuration of the TM tape, together with the head location and current internal state, at a particular time step. Adjacent rows encode the TM configuration at successive time steps. The correct TM evolution is then enforced by tiling rules. (See figure 4 for an example of such an encoding, and see [6, 8, 13, 20, 32] for some further detailed discussions on this topic.)

4.3 Encoding Turing Machines in the Robinson Tiling

In this section we review how the tiling-encoding of TMs can be combined with the Robinson tiling to create a new set of tiles which, when the plane is tiled according to the tiling rules, encodes the evolution of a separate TM within each n-square in the Robinson

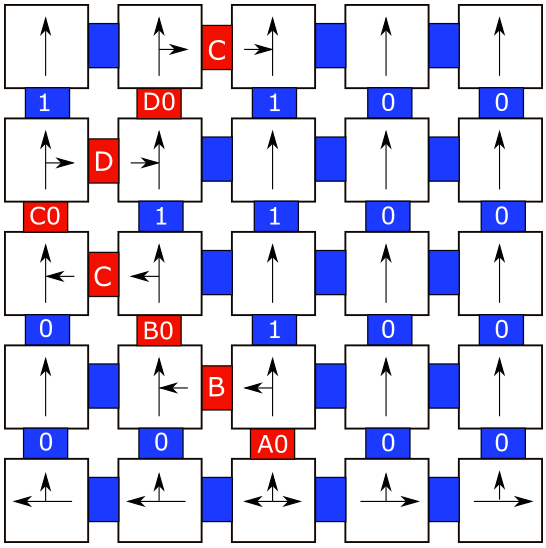

Figure 4: The evolution of a classical TM can be represented by Wang tiles, where colours of adjacent tiles have to match, and arrow heads have to meet arrow tails. Here the evolution runs from the bottom of the square to the top. Blue labels represent the entry to the corresponding cell on the TM tape. The red labels between adjacent rows represent the position and state of the TM head in addition to the cell's entry, and the red labels between adjacent columns represent movement of the TM head after it has acted on the cell.

tiling pattern. This construction was introduced in [32] to prove undecidability of the tiling of a 2D plane.

Encoding the evolution of a TM directly within the interior of a n-border is not possible as the Robinson tiling pattern is composed of m-squares nested within other n-squares, $m < n$. Thus TMs associated the m-squares and n-square would overlap with each other. [32] circumvents this problem by identifying a sub-grid within each Robinson n-border which allows a TM to be encoded without overlapping with the smaller m-squares, $m < n$, nested within.

Definition 4.2 (Free Rows/Columns and Free Squares, [32]). A *free row/column* of square is a row/column in a Robinson n-border that stretches across the border's interior uninterrupted by any of the m-borders with $m < n$.

A *free square or tile* is a square in the grid that is both in a free row and a free column (see figure 5 for an example). Within an n-square there are exactly $2^n + 1$ free rows/columns.

Lemma 4.3 (Encoding TM in Robinson Tiling, [32]). *Consider any classical Turing Machine which can have its evolution be encoded in a $(2^n + 1) \times (2^n + 1)$ grid of Wang tiles. Then the evolution of this TM can be encoded in the free rows and columns of an n-square in a Robinson Tiling.*

Following [32], to demarcate where the free tiles are within an Robinson n-square, so that we can encode a Turing Machine in them, a new kind of marking called an 'obstruction signal' is introduced. These signals are designed so they are emitted and

absorbed from the outside of a red border and while also being absorbed by the inside of a border, as seen in figure 5. In terms of tiles, these markings are formed by adding an additional set of markings such that Robinson borders "emit" the obstruction signals from one side and "absorb" them on both sides. Tiles that do not emit or absorb obstruction signals force them to propagate in the same direction. A *free tile* is one which does not have an obstruction signal going across it in either direction. In our new tile set, we only encode the Turing Machine tape, head and state symbols in the free tiles.

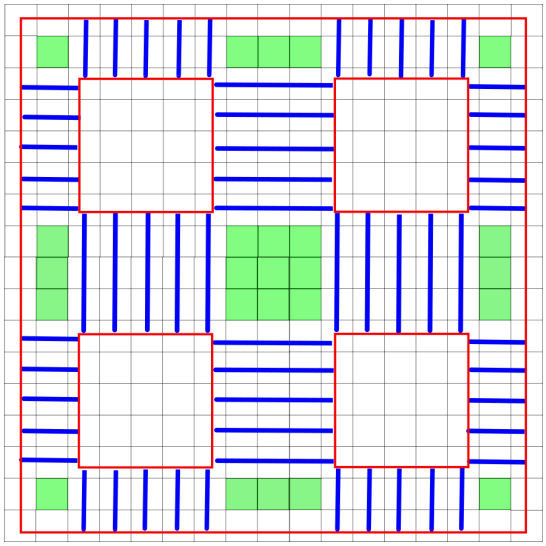

Figure 5: The obstruction signals for a red 2-square are shown in blue and its free tiles are highlighted in green. Each of the tiles within the 1-squares emits a signal outwards. The free rows are the rows in which there are no obstruction signals running horizontally (for example the central row). The free columns are the columns in which there are no obstruction signals running vertically (for example the central column). The free tiles have no obstruction signals on them in either direction.

Transmitting Signals between Free Tiles. Thus we are able to encode the evolution of a Turing Machine in these free tiles, effectively creating a $(2^n+1)\times(2^n+1)$ square for it to run in. There is a problem in that the free tiles are not spatially close to each other. To solve this, [32] implicitly introduces a new set of tile markings: Turing Machine signals. These signals can be emitted and absorbed by free tiles and run along free rows and columns. Otherwise they are absorbed by tiles with double arrowed red markings. Tiles which are not free tiles, and do not absorb the TM signals, force the TM signals to propagate across them. These signal markings allow the tiling to transmit the necessary conditions between spatially distant free tiles. A similar set of markings can be used to ensure the correct initialisation of the encoded Turing Machine.

5 PROOF OUTLINE

5.1 Classical Hardness for $\mathsf{P}^{\mathsf{NEEXP}}$

In this section we set out to outline the proof of the following theorem:

Theorem 5.1. $\mathsf{P}^{\mathsf{NEEXP}} \subseteq \mathsf{EXP}^{\mathsf{GSED}}$, *for* GSED *as defined in definition 2.2, for a classical, nearest-neighbour, translationally invariant Hamiltonian.*

To prove this result, we will show that it is possible to encode the outputs of a doubly-exponential time nondeterministic TM in the ground state energy density of a particular, fixed, classical Hamiltonian.

To start, we want to enumerate over all input strings for a TM deciding some language, encode these using tiles, and arrange for the TMs running on different inputs to be encoded within Robinson borders of different sizes.

As per lemma 4.3, we are able to encode the evolution of a TM in the $(2^n+1)\times(2^n+1)$ grid of free tiles of Robinson n-squares. We then choose the encoded TMs to first run a counter TM which computes the square size it is encoded in (similar to the constructions in [20, 30]). The counter works by outputting all binary strings in lexicographic order, such that at one end of the grid it outputs the number of free rows/columns within the n-square in binary: $L = 2^n + 1$. We then follow this counter TM with second TM which takes L as input and computes $\log_2(L-1) - n_0$ in binary, where because it is guaranteed that $L = 2^n + 1$ for $n \in \mathbb{N}$, the output is a binary string expressing the $(n-n_0)^{th}$ integer. Here n_0 is defined as the smallest integer such that the above Turing Machine can run correctly within a Robinson n_0-square[2]. We label the output of the above TMs when within an n-square as x_n.

We then run a fixed non-deterministic Turing Machine M which takes as input x_n and runs for time $2^{2^{c|x_n|}}$, $c \geq 1$. To allow for non-deterministic TMs, we slightly modify the construction seen in section 4 so that the tiling rules allow for multiple possible transitions at a given point by having multiple tiling rules which are permissible. Thus a valid tiling corresponds to one possible path that the non-deterministic TM takes. This gives the following lemma:

Lemma 5.2 (TMs in Robinson Squares). *Let $x_n \in \{0,1\}^*$ be the $(n-n_0)^{th}$ string in lexicographic order where n_0 is a fixed integer, and let M be a non-deterministic TM. It is possible to construct a tile set such that all valid tilings of an $L \times L$ lattice consist of the pattern of nested squares formed by the Robinson tiling, such that within each complete n-border, $\forall n \geq n_0$, the tiles encode a valid computational evolution of $M(x_n)$ for time $2^{2^{c|x_n|}}$, $c \geq 1$.*

Note that, at this point, the tiling here can encode any computational path (even those which reject when there is an accepting path) of the nondeterministic TM M as we have not constrained the output in any way.

5.1.1 Mapping Tiles to Hamiltonians. So far we have presented the problem in terms of a tiling problem and need to map this to a classical Local Hamiltonian problem. This is a standard technique (see [20, Section 3] or the appendix of [6] for a summary). Consider a set of Wang tiles $\mathcal{T}$ rules with horizontal constraints $\mathcal{R}_{Horz} \subseteq$

[2]The output of these TMs for n-squares $n < n_0$ will be irrelevant for our purposes.

$\mathcal{T} \times \mathcal{T}$ such that if t_i is placed to the left of t_j, then it must be the case that $(t_i, t_j) \in \mathcal{R}_{Horz}$ and likewise for the vertical tiling rules $\mathcal{R}_{Vert}$.

Map every tile type $t_i \in \mathcal{T}$ to a spin state of a classical particle $|t_i\rangle$. We then impose a Hamiltonian over the lattice such that if the tiling pair $(t_i, t_j) \notin \mathcal{R}_{Horz}$ (or $(t_i, t_j) \notin \mathcal{R}_{Vert}$ depending on the orientation), then we introduce the term $|t_i t_j\rangle \langle t_i t_j|$ for all forbidden pairings (t_i, t_j) over all points in the lattice. Adding this term to the Hamiltonian applies an energy penalty if the states $|t_i\rangle |t_j\rangle$ appear directly adjacent to each other in the lattice.

Thus we end up with a Hamiltonian composed of local interactions of the form

$$h^{col}_{k,k+1} = \sum_{(t_i,t_j)\notin\mathcal{R}_{Horz}} |t_i t_j\rangle \langle t_i t_j|_{k,k+1} \tag{3}$$

$$h^{row}_{k,k+1} = \sum_{(t_i,t_j)\notin\mathcal{R}_{Vert}} |t_i t_j\rangle \langle t_i t_j|_{k,k+1}, \tag{4}$$

We now map the tiling rules produced by lemma 5.2 to a Hamiltonian to get a nearest-neighbour, translationally invariant Hamiltonian. We further add a term penalising rejecting instances of the verification computation; Π_{NO} is an additional term which assigns an energy penalty of 1 to Robinson squares which contain No problem instances.

We encapsulate the definition of the Hamiltonian in the following:

Definition 5.3 (Robinson + Computation Hamiltonian). Let $h^{col,Rob}, h^{row,Rob} \in \mathcal{B}(\mathbb{C}^R \otimes \mathbb{C}^R)$ be the local terms which encode the local matching rules for the Robinson tiling, obstruction rules and TM rules from lemma 5.2. Let $(\Pi_{NO})_{j,j+1}$ be a projector onto the reject state of the encoded TM, M, on a site in row j, and a Robinson border tile on the adjacent site in row $j+1$. Then the overall local terms are:

$$h^{row}_{i,i+1} = \Lambda h^{row,Rob}_{i,i+1} \tag{5}$$

$$h^{col}_{j,j+1} = \Lambda h^{col,Rob}_{j,j+1} + (\Pi_{NO})_{j,j+1} \tag{6}$$

where $\Lambda \in \mathbb{N}$ is a parameter that we will fix later.

Π_{NO} is constructed such that the energy penalty is only applied at the edge of a Robinson border where a TM has halted in the No state (i.e. once the TM has stopped running). Λ characterises the energy penalty for breaking the Robinson tiling, the obstruction signals, or the TM signals. We will need to choose Λ to be a sufficiently large constant to make it energetically unfavourable to break the Robinson tiling in the ground state.

If we take this Hamiltonian and restrict to the subspace corresponding to a correctly tiled n-border, we will see it is possible to encode the outcome of a doubly exponential time TM in the energy of this region. By lemma 5.2 valid tilings encode the evolution of a non-deterministic TM $M(x)$, where x is the $(n - n_0)^{th}$ string in lexicographic order.

If x is a Yes instance, then $M(x)$ must have an accepting computational path, and so there must be a set of states that encode the correct evolution which finishes in an accepting state. Hence there is no energy penalty and the ground state is 0. If x is a No instance, then there is no accepting path. Any correct evolution of $M(x)$ therefore enters the rejecting state, and the tile marking the rejecting state of the TM picks up an energy penalty of 1 from the term $(\Pi_{NO})_{k,k+1}$. Provided $\Lambda > 1$, then the ground state encodes one of these rejecting paths. This gives the following lemma:

Lemma 5.4. *Define $H(4^n)|_P$ to be the Hamiltonian on a $(4^n + 1) \times (4^n + 1)$ region described by the local terms given in equations 5 6, restricted to the subspace P corresponding to correct tilings of the region that contain a correctly tiled Robinson n-border. Let $x \in \{0,1\}^*$ be the $(n - n_0)^{th}$ string in lexicographic order and let M be a non-deterministic Turing Machine running for time $2^{2^{cm}}$ on inputs of length m, $c \geq 1$.*

Then for $n \geq n_0$, the ground state energy of $H(4^n)|_P$ is

$$\lambda_0(H(4^n)|_P) = i_n := \begin{cases} 0 & M(x) \textit{outputs} \text{ Yes} \\ 1 & M(x) \textit{outputs} \text{ No}. \end{cases} \tag{7}$$

5.1.2 Robustness of the Ground State. We now want to find the ground state energy density of the lattice with the Hamiltonian from definition 5.3. For the purposes of proving the main theorem, we wish to show that the ground state is composed of the Robinson squares + Turing Machines encoded within. However, it may in fact be the case that it is energetically favourable to introduce tiling defects — points at which the tiling rules are not obeyed — to disrupt the encoded Turing Machines which pick up an energy penalty. We note [13, 28] proved similar bounds, but are not strong enough for our purposes.

Consider a set of defects on a tiling grid, which we label D. If these defects are able to disrupt the tiling such that the number of Robinson squares disrupted relative to a Robinson tiling without defects is $\geq |D|^{1+\epsilon}$ for any $\epsilon > 0$, it will be energetically favourable to introduce defects into the tiling. The result is that the ground state will not be the Turing Machine + tiling construction we wish. Hence, to prove the ground state is the TM + tiling construction we desire, we need to show that a defect disrupting the tiling can disrupt at most $O(|D|)$ Robinson squares. We give an outline of the argument here. For full mathematical details, see [34].

To start we will consider the number of Robinson n-borders which can be destroyed by allowing tiling defects (points where the tiling rules are not satisfied) at a fixed set of points relative to a tiling grid where the tiling rules are correct everywhere.

Definition 5.5 (Border deficit). The *total border deficit*, deficit(T), of a tile configuration T is the difference between the total number of complete borders in T and the number of complete borders in a Robinson tiling of the same region, maximised over Robinson tilings.

To consider where the tiling is forced to be correct, we take the points at which defects occur and use them as vertices to form a complete graph, which we call the defect graph $G = (D, E)$. To bound the border deficit, we split the lattice into regions known as n-domains and n-undomains: where an n-undomain is a maximal connected region of the lattice such that any m-border with $m \geq n$ that overlaps U necessarily either intersects a defect, or intersects an edge in of length $\leq 4^n$. An n-domain to be a maximal connected region of the lattice that does not overlap any n-undomain.

It can then be shown that a lattice cell contained in an n-undomain has an edge from E of length $\leq 4^n$ within distance $\leq 4^n$. Using

these definitions, it is possible to show that within n-domains, the Robinson tiling is "correct" up to a certain level:

Proposition 5.6. *Consider a tile configuration, a set of defects D, and $G = (D, E)$ its defect graph. Within any n-domain $\mathcal{D}$, the tile configuration contains the same periodic pattern of m-borders for all $m \leq n$ as a normal Robinson tiling of $\mathcal{D}$, except where an m-border would intersect a point in D.*

Roughly, one can show that, relative to some perfect reference Robinson tiling, a n-border is guaranteed to be intact (i.e. not destroyed) if there is no edge in the defect graph which intersects one of its edges in the reference tiling. This fact is not quite strong enough to prove the necessary bound; the defect graph has $\propto |D|^2$ edges, and so we only get an upper bound $\leq |D|^2$ on the number of borders destroyed.

To improve on this, we take the Delaunay triangulation of all the defects; a Delaunay triangulation is planar graph such that no vertex is inside the circumcircle of any triangle the graph [27]. Importantly, we are able to show that if an n-border is intersected by an edge of length $\leq 4^n$ in the defect graph G, then there is an extended region around the n-border which must be cut by an edge in the Delaunay triangulation. We call this area the n-frame.

Since the Delaunay graph is planar, by Euler's formula for graphs it must have $\leq 3|D| - 6$ edges, where $|D|$ is the number of vertices. By carefully counting across domains and undomains, it can be shown that an edge in the Delaunay triangulation cuts at most a constant number n-frames. This allows us to prove:

Theorem 5.7. *Let T be a tile configuration of grid with perimeter of length L. Let D denote its defect set. The border deficit of T is bounded by*

$$\text{deficit}(T) \leq 399|D| + L. \tag{8}$$

This isn't quite sufficient for our purposes; it needs to be proven that the number of borders which are prevented from being able to encode Turing Machines in their interiors. There are several ways this can happen: (a) the smaller Robinson squares nested within an n-border are destroyed by defects or otherwise misaligned (b) the obstruction signals along free rows or columns are disrupted (c) the Turing Machine signals along free rows and columns are disrupted. However, it can be shown that if an n-border is intact, then its interior tiling must be correct unless one of the smaller m-borders, $m < n$, which are not contained in any k-borders, $m < k < n$, is destroyed. Or alternatively if there is a defect in its interior which is not contained in any smaller m-border.

Using these facts gives:

Lemma 5.8. *Let T be a tile configuration of a finite subregion of $\mathbb{Z}_2$ with perimeter of length L, D its defect set.*

Define the total deficit *of T*, $\text{total\_deficit}(T)$, *to be the difference between the total number of complete Robinson squares in T with a correct internal Turing Machine tiling, and the number of these in a Robinson tiling of the same region, maximised over Robinson tilings.*

The total deficit of T is bounded by

$$\text{total\_deficit}(T) \leq 801|D| + 2L. \tag{9}$$

With this in hand, we will use this to find the ground state energy density of the Robinson tiles + TM construction.

5.1.3 Applying the Defect Bounds. In the following, we use the square deficit bounds established in section 5.1.2 to show that it is energetically unfavourable to have too many tiling defects, regardless of how many No instances might be encoded in n-squares.

To do this we must calculate the energy of different tiling configurations. We first realise from geometry that the number of Robinson squares of size $4^n + 1$ on an $L \times L$ lattice must be in the interval $\left[\left(\left\lfloor \frac{L}{2^{2n+1}} \right\rfloor - 1 \right)^2, \left(\left\lfloor \frac{L}{2^{2n+1}} \right\rfloor + 1 \right)^2 \right]$. Then applying lemma 5.8, we see that one can destroy at most $\leq 801|D| + 2L$ squares containing Turing Machines with the corresponding defect set D. Since each of these Robinson squares has maximum energy 1, this reduces the energy by at most an amount $O(|D|) + O(L)$. However, each tiling defect will incur an energy penalty Λ due to the mismatched tiling rules. Thus, if D is the set of defects, then the total energy change due to the defects will be $\Lambda|D| - k_1|D| - k_2 L$. Thus by setting Λ to be a sufficiently large constant we can make it energetically unfavourable for $|D|$ to increase (i.e. for defects to occur).

Accounting for the fact that all intact Robinson borders will otherwise contribute a certain energy $\lambda_0(H(4^n)|_P) \in \{0, 1\}$ depending on whether their internal nondeterministic Turing Machine accepts or rejects, we get the following expression for the grounds state energy on an $L \times L$ lattice.

Lemma 5.9. *Let $h^{row}, h^{col} \in \mathcal{B}(\mathbb{C}^R \otimes \mathbb{C}^R)$ be the local interactions that encode the tiling rules given by equations 5 and 6. Let $H^{\Lambda(L\times L)}$ be the Hamiltonian with these local interactions on $\Lambda(L \times L)$.*

Then for sufficiently large L, the ground state energy $\lambda_0(H^{\Lambda(L\times L)})$ is contained in the interval

$$\left[\sum_{n=n_0}^{\lfloor \log_4(L/2) \rfloor} \left(\left\lfloor \frac{L}{2^{2n+1}} \right\rfloor - 1 \right)^2 \lambda_0(H(4^n)|_P) + \Lambda|D| - k_1|D| - k_2 L, \right.$$
$$\left. \sum_{n=n_0}^{\lfloor \log_4(L/2) \rfloor} \left(\left\lfloor \frac{L}{2^{2n+1}} \right\rfloor + 1 \right)^2 \lambda_0(H(4^n)|_P) + \Lambda|D| - k_1|D| - k_2 L \right] \tag{10}$$

for some constants Λ, k_1 and k_2 such that $\Lambda \gg k_1 + k_2$, and $|D| = O(L)$.

To calculate the ground state energy density, we simply need to divide by L^2 take the limit as $L \to \infty$ to get the following:

Lemma 5.10. *Consider an $L \times L$ lattice with a local Hamiltonian interactions given by equations 5 and 6, and let $H(4^n)|_P$ and i_n be defined as in lemma 5.4. In the limit of $L \to \infty$, the ground state energy density is*

$$\mathcal{E}_\rho = \frac{1}{4} \sum_{n=n_0}^{\infty} \frac{\lambda_0(H(4^n)|_P)}{16^n} = \frac{1}{4} \sum_{n=n_0}^{\infty} \frac{i_n}{16^n}. \tag{11}$$

We now prove part of the main theorem, which we restate here for convenience.

Theorem 5.11 ($\mathsf{P}^{\mathsf{NEEXP}} \subseteq \mathsf{EXP}^{\mathsf{GSED}}$). *$\mathsf{P}^{\mathsf{NEEXP}} \subseteq \mathsf{EXP}^{\mathsf{GSED}}$, for* GSED *as defined in definition 2.2, for a classical, translationally invariant, nearest-neighbour Hamiltonian.*

Proof. Consider any polytime bounded TM M_1. We will show we can simulate M_1^{NEEXP} with M_2^{GSED} where M_2 is another exptime TM. If M_1^{NEEXP} takes an n-bit input, it can then make $O(\text{poly}(n))$

adaptive queries. Denote these queries by $\{q_i\}_{i=1}^{O(\text{poly}(n))}$. Each individual query must have length $|q_i| = O(\text{poly}(n))$. The M_1 machine then runs for an $O(\text{poly}(n))$ time and produces some output.

To simulate this, M_2 will take the n-bit input and calculate the first query q_1 that M_1 makes to the NEEXP oracle. M_2 takes the query q_1 and reduces it to an instance of determining the output of a doubly-exponentially time non-deterministic TM, M, on input y_1. This reduction can be computed in polynomial time, as the problem of determining the output of double-exponential-time non-deterministic TMs is manifestly NEEXP-hard. (Note by using padding arguments we can reduce any language in NEEXP to $\text{NTIME}(2^{2^{cn}})$ for some $c > 1$ [30]). For now we assume that M_2 can query the GSED oracle using a subroutine to obtain the oracle answer — we show this in the second part of the proof. The M_2 machine then uses the output of the query to compute the next query that M_1 makes, and computes q_2.

The process then repeats: assuming M_2 has calculated the first $\{q_i\}_{i=1}^{k-1}$, it then uses the output from these simulated queries to calculate q_k. As before, it reduces this to an instance of determining the output of a doubly-exponentially time non-deterministic TM, M, on input y_k. It then uses the GSED subroutine to extract the corresponding answer from the oracle. M_2 does this for all queries M_1 makes, $\{q_i\}_{i=1}^{O(\text{poly}(n))}$, and obtains the answer to the queries from the GSED oracle.

Determining Queries using a GSED Oracle. We will use the GSED oracle for the Hamiltonian of definition 5.3 to perform a binary search in order to obtain a sufficiently precise approximation to the ground state energy density $\mathcal{E}_\rho$, such that we can extract the result of computing M on all inputs up to y_j. To do this, we need to query the GSED oracle on all the instances before it in lexicographic order, of which there are $k = O(2^{\text{poly}(n)})$ many.

By lemma 5.10, outputs i_n to the queries $\{y_i\}_i$ are encoded as

$$\mathcal{E}_\rho = \frac{1}{4}\sum_{n=n_0}^{\infty}\frac{i_n}{16^n}. \tag{12}$$

We extract the i_k iteratively as follows. Assume for simplicity that $n_0 = 1$. (If this is not the case, n can straightforwardly be adjusted appropriately.) To determine the i_1, note that if $i_1 = 0$, then the maximum $\mathcal{E}_\rho$ can be is

$$\mathcal{E}_\rho = \frac{1}{4}\sum_{n=2}^{\infty}\frac{1}{16^n} = \frac{1}{960} \tag{13}$$

and otherwise the minimum it can be is $1/64$. Hence M_2 asks the GSED oracle whether $\mathcal{E}_\rho \geq \beta_1 = 1/64$ or $\mathcal{E}_\rho \leq \alpha_1 = 1/960$. Thus

$$i_1 = \begin{cases} 0 & \text{if } \mathcal{E}_\rho < 1/960 \\ 1 & \text{if } \mathcal{E}_\rho > 1/64. \end{cases}$$

M_2 then performs a similar process for all i_m, $1 \leq m < k$, assuming it has previously extracted $i_1, i_2, \ldots, i_{m-1}$. When extracting the m^{th} instance, we have that either $\mathcal{E}_\rho \leq \alpha_m$ or $\mathcal{E}_\rho \geq \beta_m$, where

$$\beta_m = \frac{1}{4}\left(\frac{1}{16^m} + \sum_{n=1}^{m-1}\frac{i_n}{16^n}\right)$$
$$\alpha_m = \frac{1}{4}\left(\sum_{n=1}^{m-1}\frac{i_n}{16^n} + \sum_{n=m+1}^{\infty}\frac{1}{16^n}\right). \tag{14}$$

Since y_j is the k^{th} string in lexicographic order, $k = O(2^{\text{poly}(n)})$, the maximum precision we need to go to is $\Omega(2^{-2^{\text{poly}(n)}})$, which is possible provided α_m, β_m can have binary length $|\alpha_m|, |\beta_m| = O(2^{\text{poly}(n)})$. Since M_2 is an exponential time machine, it has time and space to write these strings to the oracle tape. Furthermore, M_2 only needs to make $O(2^{\text{poly}(n)})$ queries. Thus M_2^{GSED} is able extract all the answers to the queries made by M_1^{NEEXP}, and hence after making these queries and performing the relevant post-processing, output the solution. □

5.2 Classical Containment in EXP^{NEXP}

We now need to show that for classical GSED, as defined in definition 2.2, $\text{EXP}^{\text{GSED}} \subseteq \text{EXP}^{\text{NEXP}}$. The first step is to show that the ground state energy density of a finite $L \times L$ part of the lattice is a good estimate for the energy density of the full lattice [13]:

Lemma 5.12. *Consider a translationally invariant, nearest-neighbour Hamiltonian on $\Lambda(L \times L)$ lattice defined by local terms $h^{row}_{i,i+1}, h^{col}_{j,j+1}$. Let $\mathcal{E}_\rho(L)$ be the energy density of the Hamiltonian on this lattice, and let $\mathcal{E}_\rho$ be the energy density in the $L \to \infty$ limit. Then*

$$|\,\mathcal{E}_\rho(L) - \mathcal{E}_\rho\,| = \frac{4\max\left\{\left\|h^{row}_{i,i+1}\right\|, \left\|h^{col}_{i,i+1}\right\|\right\}}{L}. \tag{15}$$

Proof. Let $H(L)$ be the Hamiltonian defined on $\Lambda(L \times L)$ and let $t \in \mathbb{N}$. Let $H_{grid}(L,t)$ be the Hamiltonian with the same local terms, but with the terms $h^{row}_{i,i+1}, h^{col}_{j,j+1}$ removed for $i, j \in t\mathbb{N}$. Then:

$$H_{grid}(L,t) = H(tL) - \sum_{i \bmod t=0} h^{row}_{i,i+1} - \sum_{j \bmod t=0} h^{row}_{j,j+1}. \tag{16}$$

The interaction graph of $H_{grid}(L,t)$ is a set of t^2 squares of size $L \times L$. Hence equation 16 gives

$$\left\|H_{grid}(L,t) - H(tL)\right\| \leq 4t^2 L \max\left\{\left\|h^{row}_{i,i+1}\right\|, \left\|h^{col}_{i,i+1}\right\|\right\}.$$

It is straightforward to see that $\lambda_0(H_{grid}(L,t)) = t^2\lambda_0(H(L))$. Combining these gives

$$|t^2\lambda_0(H(L)) - \lambda_0(H(tL))| \leq 4Lt^2 \max\left\{\left\|h^{row}_{i,i+1}\right\|, \left\|h^{col}_{i,i+1}\right\|\right\}.$$

Dividing through by t^2L^2 to get energy densities gives

$$|\,\mathcal{E}_\rho(L) - \mathcal{E}_\rho\,| \leq \frac{4\max\left\{\left\|h^{row}_{i,i+1}\right\|, \left\|h^{col}_{i,i+1}\right\|\right\}}{L}. \tag{17}$$

□

Lemma 5.13. GSED ∈ NEXP *for any classical, nearest-neighbour, translationally invariant Hamiltonian, for* GSED *as defined in definition 2.2.*

Proof. (α, β) is the input of the problem, $\beta - \alpha = \Omega(2^{-q(n)})$. We show an EXP machine will be able calculate $\mathcal{E}_\rho(L)$ (using the notation of lemma 5.12) using a classical witness for $L = 2^{p(n)}$, for a polynomial p.

First compute the ground state energy of an $L \times L$ square of the lattice. Take as the witness the ground state of the Hamiltonian restricted to an $L \times L$ region of the lattice: $|\psi\rangle = |\phi_1\rangle \otimes |\phi_2\rangle \otimes \ldots |\phi_{L^2}\rangle$, where $|\phi_i\rangle \in \mathbb{C}^{|\mathcal{S}|}$ is the state of the spin at lattice site i. Now,

$$\mathcal{E}_\rho(L) = \frac{1}{L^2} \sum_{\langle i,j \rangle} \langle \phi_i | \langle \phi_j | h_{i,j} | \phi_i \rangle | \phi_j \rangle ,$$

where $\langle i, j \rangle$ denotes pairs of nearest-neighbours. $\langle \phi_i | \langle \phi_j | h_{i,j} | \phi_i \rangle | \phi_j \rangle$ can be computed in $O(1)$ time, and there are $O(L^2)$ such terms. Since $L = 2^{p(n)}$, the estimate $\mathcal{E}_\rho(L)$ can be computed in $O(L^2) = O(2^{2p(n)})$ time. By lemma 5.12, $|\mathcal{E}_\rho(L) - \mathcal{E}_\rho| = O(L^{-1})$, hence provided we choose $p(n)$ to be sufficiently large relative to $q(n)$, the approximation $\mathcal{E}_\rho(L)$ allows us to determine $\mathcal{E}_\rho > \beta$ or $\mathcal{E}_\rho < \alpha$ for $\beta - \alpha = \Omega(2^{-q(n)})$. □

An immediate corollary of this is:

Corollary 5.14 (EXP$^{\text{NEXP}}$ Containment). EXP$^{\text{GSED}} \subseteq$ EXP$^{\text{NEXP}}$, *for* GSED *as defined in definition 2.2, for a fixed, classical Hamiltonian.*

Why not Polytime Turing Reductions, P$^{\text{GSED}}$*?* Naturally a question arises as to why we consider EXP$^{\text{GSED}}$ here, rather than P$^{\text{GSED}}$. Here we show that using our hardness construction, one cannot even hope to prove NP $\subseteq$ P$^{\text{GSED}}$ unless the polynomial hierarchy collapses to Σ_2^P.

Theorem 5.15. *Let* P$^{\text{GSED}_h}$ *be the class of languages decided by a polynomial time oracle machine with access to a* GSED *oracle for the Hamiltonian of definition 5.3 only. If* NP $\subseteq$ P$^{\text{GSED}_h}$*, then the polynomial hierarchy collapses to* Σ_2^P.

Proof. We will first show that P$^{\text{GSED}_h} \subseteq P/\text{poly}$. Let $M_1^{\text{GSED}_h}$ be a polytime TM with oracle access to a GSED oracle for the Hamiltonian defined in definition 5.3 only. $M_1^{\text{GSED}_h}$ can make at most $O(\text{poly}(n))$ length queries to the oracle, corresponding to α, β queries such that $\beta - \alpha = \Omega(2^{-p(n)})$ for some polynomial p. After making at most poly(n) queries, it performs some post-processing and finally outputs an answer.

However, we note that a polytime TM M_2 with access to an advice string of $\tilde{\mathcal{E}_\rho}$, such that $\tilde{\mathcal{E}_\rho}$ has $O(\text{poly}(n))$ bits and satisfies $|\tilde{\mathcal{E}_\rho} - \mathcal{E}_\rho| \leq 2^{-q(n)}$, where q is a polynomial $q \gg p$. Thus M_2 is able to replicate any queries that M_1 makes to the GSED oracle as it has been given the relevant ground state energy density to sufficient precision to replicate the queries for all α, β satisfying $\beta - \alpha = \Omega(2^{-p(n)})$.

Thus it is the case that P$^{\text{GSED}_h} \subseteq P/\text{poly}$. However, it is known that if NP $\subseteq P/\text{poly}$, then the polynomial hierarchy collapses to Σ_2^P [24].

□

This provides strong evidence that our hardness construction is not NP-hard under polytime Turing reductions.

5.2.1 Improving the Hardness Result. We can improve our containment and hardness results by using a PSPACE oracle machine. There is, however, some controversy as to how a PSPACE oracle machine should have access to its oracle; in particular whether the input tape to the oracle has a polynomial space bound or not [10, 15, 21]. Here we consider both of these definitions and show how they can be used to tighten our complexity bounds on GSED.

Definition 5.16 (1st PSPACE *Oracle Machine Definition).* A PSPACE^O oracle machine is a PSPACE machine with access to an oracle input tape, for which it can make *polynomial length* queries to the oracle.

For this definition we get:

Theorem 5.17. PSPACE$^{\text{NEEXP}} \subseteq$ EXP$^{\text{GSED}}$.

Proof. Identical to the proof for theorem 5.11 except M_1 is now a PSPACE machine which needs to be simulated by the EXP$^{\text{GSED}}$ oracle machine. □

A potentially more interesting result occurs when we use the following definition:

Definition 5.18 (2nd PSPACE *Oracle Machine Definition).* A PSPACE^O oracle machine is a PSPACE machine with access to a write only oracle input tape, for which it can make *exponential length* queries to the oracle.

This is the preferred definition of several authors [15, 26]. For this definition of oracle machine, we realise that one can do the binary search protocol used in the proof of theorem 5.11 to get:

Theorem 5.19. P$^{\text{NEEXP}} \subseteq$ PSPACE$^{\text{GSED}}$.

Proof. The proof will be similar to the proof for theorem 5.11, except now the PSPACE machine will have to make exponentially long oracle calls to the GSED oracle to extract the query results while using only polynomial space everywhere else.

Let M^{GSED} be a PSPACE machine with (for convenience) two work tapes[3] (bounded by polynomial space) and one unbounded oracle tape which is read only. Let the GSED oracle be the one for the Hamiltonian of definition 5.3. Let M^{GSED} have made $(k-1)$ queries to the oracle machine with outputs $i_1, i_2 \ldots i_{k-1}$, for i_j as defined in lemma 5.4, such that it now needs to make a k^{th} query. To do so, it needs to calculate a pair (α_k, β_k) which will allow it to extract i_k. Assume M has the string $i_1 i_2 \ldots i_{k-1}$ stored on one of the two work tapes. We need to write out the numbers α_k, β_k in binary as given in equation 14.

Without loss of generality, assume the oracle input tape is initially in the all 0 state. To write out β_k on the input tape, M take a query outcome i_j, then moves $4j+2$ down the tape and places i_j in the $(4j+2)^{th}$ cell (corresponding to value $\frac{1}{4}\frac{i_j}{16^j}$). Finally in the $(4k+2)^{th}$ cell it places a 1. To determine where the head is on the oracle input tape, we let M have a binary counter on its second work tape. M can determine where the head is on the input tape by increment/decrementing the binary counter whenever the head moves right/left.

[3] This can be reduced to a single work tape by standard arguments.

M cannot write out α_k exactly, as it does not have a finite binary expansion. Instead, upper bound it by a number $a_k > \alpha_k$, $\beta - a_k = \Omega(2^{-\text{poly}(k)})$ which does have a finite expansion

$$a_k = \frac{1}{4}\left(\sum_{n=1}^{k-1}\frac{i_n}{16^n} + \frac{2}{16^{k+1}}\right) > \alpha_k. \tag{18}$$

To write this out, M also places i_j in the $(4j+2)^{th}$ cell, for $j \leq k-1$. We then place a 1 in the $(4k+3)^{th}$ cell (which is the contribution from the $2 \times 16^{-k-1}$ term). Hence querying the oracle for (a_k, β_k) gives the same answer as querying with (α_k, β_k).

M then continues with the computation until all the necessary queries have been extracted. Since only poly(n) many queries are made, the PSPACE machine is capable of storing them all on its work tape. It can then post-process the queries and output the answer to the relevant $\mathsf{P}^{\mathsf{NEEXP}}$ computation.

Since M only needs to record the number of queries $k = O(\text{poly}(n))$ and the binary counter it uses to keep track of the TM head on the input string — which counts up to $16^{O(\text{poly}(k))}$ — can be expressed in poly(k) = poly(n) bits, we have that M only uses poly(n) space on its two work tapes, as required. □

The results from this section immediately give:

Corollary 5.20. $\mathsf{P}^{\mathsf{NEEXP}} \subseteq \mathsf{PSPACE}^{\mathsf{GSED}} \subseteq \mathsf{PSPACE}^{\mathsf{NEXP}}$

5.2.2 Complexity Results for FGSED. We show containment of the function problem version FGSED:

Theorem 5.21. $\mathsf{FGSED} \in \mathsf{FP}^{\mathsf{GSED}} \subseteq \mathsf{FP}^{\mathsf{NEXP}}$ *for classical* FGSED.

Proof. Let ϵ be the input to FGSED, such that $|\epsilon| = n$. Let M^{GSED} be a polytime TM with oracle access to GSED. Then using poly(n) many (α, β) queries to GSED, for $\beta - \alpha = \Omega(2^{-\text{poly}(n)})$, we can use a binary search procedure to find an estimate $\tilde{\mathcal{E}_\rho}$ such that $|\tilde{\mathcal{E}_\rho} - \mathcal{E}_\rho| = O(2^{-\text{poly}(n)}) < \epsilon$. Thus a M^{GSED} machine can compute FGSED. Since GSED $\in$ NEXP, this implies $\mathsf{FGSED} \in \mathsf{FP}^{\mathsf{GSED}} \subseteq \mathsf{FP}^{\mathsf{NEXP}}$. □

Lemma 5.22. $\mathsf{FP}^{\mathsf{NEEXP}} \subseteq \mathsf{FEXP}^{\mathsf{FGSED}} \subseteq \mathsf{FEXP}^{\mathsf{NEXP}}$ *for classical* FGSED.

Proof. To show $\mathsf{FEXP}^{\mathsf{FGSED}} \subseteq \mathsf{FEXP}^{\mathsf{NEXP}}$, consider two exponential time oracle machines M_1^{FGSED} and M_2^{NEXP}. Let M_1 make $O(\exp(n))$ oracle calls to FGSED, and then do some exponential time post-processing. M_2 can simulate these oracle calls by, for each oracle call M_1 makes, estimating using the NEXP oracle $\exp(n)$ to estimate the ground state energy density produced by FGSED. Since M_1 makes $\exp(n)$ queries, M_2 needs to make $O(\exp(n)) \times O(\exp(n)) = O(\exp(n))$ queries. It can then perform the same post-processing as M_1. Thus $\mathsf{FEXP}^{\mathsf{FGSED}} \subseteq \mathsf{FEXP}^{\mathsf{NEXP}}$.

To show $\mathsf{FP}^{\mathsf{NEEXP}} \subseteq \mathsf{FEXP}^{\mathsf{FGSED}}$, consider a polytime oracle machine M_3^{NEEXP} and an exptime oracle machine M_4^{FGSED}. M_3 can make at most $O(\text{poly}(n))$ queries to the NEEXP oracle of at most $O(\text{poly}(n))$ length, and then do post-processing to output the relevant function. M_4 can simulate all of these queries by asking the FGSED oracle for an estimate for ϵ such that $|\epsilon| = O(\exp(n))$, from which it can extract all the NEEXP queries. It can then do the relevant post-processing and output the same function as M_3.

□

5.3 Quantum Containment in $\mathsf{EXP}^{\mathsf{QMA}_{\mathsf{EXP}}}$

In this section we show containment of GSED for quantum Hamiltonians.

Lemma 5.23. GSED $\in \mathsf{QMA}_{\mathsf{EXP}}$ *for any quantum, nearest-neighbour, translationally invariant Hamiltonian, for* GSED *as defined in definition 2.2.*

Proof. (α, β) is the input of the problem for $\beta - \alpha = \Omega(2^{-p(n)})$. Let $|\psi\rangle$ be the ground state an $L \times L$ section of the lattice, for $L = 2^{q(n)}$, which our $\mathsf{QMA}_{\mathsf{EXP}}$ machine will take as a witness. Perform quantum phase estimation of $e^{iH^{\Lambda(L)}}$ to $q(n)$ bits of precision, which gives an estimate $\tilde{\lambda_0}$ of $\lambda_0(H^{\Lambda(L)})$ such that $|\tilde{\lambda_0} - \lambda_0(H^{\Lambda(L)})| \leq 2^{-p(n)}$, and takes time $O(2^{q(n)})$ [29].

Since $\mathcal{E}_\rho(L) = \tilde{\lambda}_0$, and by lemma 5.12 that $|\mathcal{E}_\rho(L) - \mathcal{E}_\rho| = O(2^{-p(n)})$, choosing $q(n)$ to be sufficiently larger than $p(n)$ allows us to verify whether $\mathcal{E}_\rho > \beta$ or $\mathcal{E}_\rho < \alpha$. □

This immediately gives:

Corollary 5.24. $\mathsf{EXP}^{\mathsf{GSED}} \subseteq \mathsf{EXP}^{\mathsf{QMA}_{\mathsf{EXP}}}$ *for a fixed, nearest-neighbour, translationally invariant* ***quantum*** *Hamiltonian.*

Since classical Hamiltonians are a subset of quantum Hamiltonians, then $\mathsf{P}^{\mathsf{NEEXP}} \subseteq \mathsf{EXP}^{\mathsf{GSED}}$ is an immediate corollary of theorem 5.11 for quantum GSED.

6 DISCUSSION AND CONCLUSIONS

Quantum GSED. A natural question to ask is if tighter results can be found for GSED for quantum Hamiltonians. As we have seen, it follows straightforwardly that $\mathsf{EXP}^{\mathsf{GSED}} \subseteq \mathsf{EXP}^{\mathsf{QMA}_{\mathsf{EXP}}}$, but a tighter quantum lower bound does not follow easily.

Our proof of a $\mathsf{P}^{\mathsf{NEEXP}}$ lower bound works as we can enumerate over NEEXP-complete problems. Attempting to prove a similar quantum lower bound (e.g. $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$) runs into the problem that, since $\mathsf{QMA}_{\mathsf{EXP}}$ is a promise class, for a given $\mathsf{QMA}_{\mathsf{EXP}}$-complete problem there may be instances which do not satisfy the promise (so called "invalid queries"). This makes it impossible to enumerate over all instances of a given $\mathsf{QMA}_{\mathsf{EXP}}$-complete problem without potentially including instances which do not satisfy the promise. It is not currently known how to avoid these instances from occurring, although some techniques exist, such as [17, 19, 33].

Closing the Classical Upper and Lower Bounds. So far we have separate lower and upper bounds $\mathsf{P}^{\mathsf{NEEXP}}$ and $\mathsf{EXP}^{\mathsf{NEXP}}$. The containment protocol given here works via a natural binary search algorithm to determine $\mathcal{E}_\rho$, and as such we believe it is optimal. While it is not immediately clear how the lower bound might be improved, it is not clear whether the construction presented here should give a tight lower bound.

Other Precision Problems. As far as the authors know, this is the first complexity hardness result about a theorem in which the only input parameter which is varied is the precision, but where the object of study is fixed. Furthermore, GSED can be viewed as a precision version of the Local Hamiltonian problem; can similar "precision based" problems be developed for other decision/promise problems? Is there a natural situation in which they occur?

We comment that one can introduce a similar problem to GSED, but for a density of local observables. One can think of this as a precision generalisation of the APX-SIM problem introduced by Ambainis in [3] and studied in [17, 19, 33], in the same way that GSED is a precision version of the local Hamiltonian problem. For example if a 1-local observable A_i, $\|A_i\| = 1$ which has an associated density, consider the following density:

$$A_\rho = \lim_{L\to\infty} \frac{1}{L^2} \sum_{i\in\Lambda(L\times L)} A_i.$$

Although one can prove identical hardness results to $\mathsf{EXP}^{\mathsf{GSED}}$ about the expectation of A_ρ on the ground state, containment does not follow in the same way as for GSED. In particular, [13] demonstrate that determining whether a particular local density expectation is 0 or 1 is undecidable.

ACKNOWLEDGMENTS

The authors would like to recognise useful discussions with J. Lockhart, S. Gharibian, and J. Sullivan. J.D.W. is supported by the EPSRC Centre for Doctoral Training in Delivering Quantum Technologies (grant EP/L015242/1). T.S.C. is supported by the Royal Society. This work has been supported in part by the EPSRC Prosperity Partnership in Quantum Software for Simulation and Modelling (grant EP/S005021/1), and by the UK Hub in Quantum Computing and Simulation, part of the UK National Quantum Technologies Programme with funding from UKRI EPSRC (grant EP/T001062/1).

REFERENCES

[1] Dorit Aharonov, Daniel Gottesman, Sandy Irani, and Julia Kempe. 2007. The Power of Quantum Systems on a Line. *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS07)* (2007). https://doi.org/10.1109/focs.2007.46

[2] Dorit Aharonov and Sandy Irani. 2021. Hamiltonian Complexity in the Thermodynamic Limit. *arXiv e-prints*, Article arXiv:2107.06201 (July 2021), arXiv:2107.06201 pages. https://doi.org/10.48550/arXiv.2107.06201 arXiv:2107.06201 [quant-ph]

[3] Andris Ambainis. 2014. On Physical Problems that are Slightly More Difficult than QMA. *2014 IEEE 29th Conference on Computational Complexity (CCC)* (2014). https://doi.org/10.1109/ccc.2014.12

[4] Francisco Barahona. 1982. On the computational complexity of Ising spin glass models. *Journal of Physics A: Mathematical and General* 15, 10 (1982), 3241.

[5] Johannes Bausch, Toby S. Cubitt, Angelo Lucia, and David Perez-Garcia. 2020. Undecidability of the spectral gap in one dimension. *Physical Review X* 10, 3 (2020). https://doi.org/10.1103/physrevx.10.031038

[6] Johannes Bausch, Toby S. Cubitt, Angelo Lucia, David Perez-Garcia, and Michael M. Wolf. 2018. Size-driven quantum phase transitions. *Proceedings of the National Academy of Sciences* 115, 1 (jan 2018), 19–23. https://doi.org/10.1073/pnas.1705042114 arXiv:1512.05687

[7] Johannes Bausch, Toby S. Cubitt, and James D. Watson. 2021. Uncomputability of phase diagrams. *Nature Communications* 12, 1 (2021). https://doi.org/10.1038/s41467-020-20504-6

[8] Robert Berger. 1966. The undecidability of the domino problem. *Memoirs of the American Mathematical Society* 66 (1966). https://doi.org/10.1090/memo/0066

[9] Brielin Brown, Steven T. Flammia, and Norbert Schuch. 2011. Computational Difficulty of Computing the Density of States. *Physical Review Letters* 107, 4 (2011). https://doi.org/10.1103/physrevlett.107.040501

[10] Jonathan F. Buss. 1988. Relativized alternation and space-bounded computation. *J. Comput. System Sci.* 36, 3 (1988), 351–378. https://doi.org/10.1016/0022-0000(88)90034-7

[11] Gregory J. Chaitin. 1975. A theory of program size formally identical to information theory. *J. ACM* 22, 3 (1975), 329–340. https://doi.org/10.1145/321892.321894

[12] Toby S. Cubitt, David Perez-Garcia, and Michael M. Wolf. 2015. Undecidability of the spectral gap. *Nature* 528, 7581 (12 2015), 207–211. https://doi.org/10.1038/nature16059 arXiv:1502.04573

[13] Toby S. Cubitt, David Perez-Garcia, and Michael M. Wolf. 2015. Undecidability of the spectral gap. (2015). https://doi.org/10.48550/arXiv.1502.04573

[14] Samir Datta and Rameshwar Pratap. 2011. Computing Bits of Algebraic Numbers. *arXiv e-prints*, Article arXiv:1112.4295 (Dec. 2011), arXiv:1112.4295 pages. https://doi.org/10.48550/arXiv.1112.4295 arXiv:1112.4295 [cs.CC]

[15] Lance Fortnow. 1994. The Role of Relativization in Complexity Theory. *Bulletin of the European Association for Theoretical Computer Science* 52 (1994), 52–229.

[16] Sevag Gharibian, Yichen Huang, Zeph Landau, and Seung Woo Shin. 2015. Quantum hamiltonian complexity. *Foundations and Trends® in Theoretical Computer Science* 10, 3 (2015), 159–282. https://doi.org/10.1561/0400000066

[17] Sevag Gharibian, Stephen Piddock, and Justin Yirka. 2020. Oracle Complexity Classes and Local Measurements on Physical Hamiltonians. In *37th International Symposium on Theoretical Aspects of Computer Science (STACS 2020) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 154)*, Christophe Paul and Markus Bläser (Eds.). Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 20:1–20:37. https://doi.org/10.4230/LIPIcs.STACS.2020.20

[18] Sevag Gharibian and Jamie Sikora. 2018. Ground State Connectivity of Local Hamiltonians. *ACM Trans. Comput. Theory* 10, 2, Article 8 (April 2018), 28 pages. https://doi.org/10.1145/3186587

[19] Sevag Gharibian and Justin Yirka. 2019. The complexity of simulating local measurements on quantum systems. *Quantum* 3 (Sept. 2019), 189. https://doi.org/10.22331/q-2019-09-30-189

[20] Daniel Gottesman and Sandy Irani. 2009. The quantum and classical complexity of translationally invariant tiling and Hamiltonian problems. In *Foundations of Computer Science, 2009. FOCS'09. 50th Annual IEEE Symposium on*. IEEE, 95–104. https://doi.org/10.1109/focs.2009.22

[21] Juris Hartmanis, Richard Chang, Jim Kadin, and Stephen G. Mitchell. 1993. Some Observations about Relativization of Space Bounded Computations. *Current Trends in Theoretical Computer Science* (1993), 423–434. https://doi.org/10.1142/9789812794499_0031

[22] Lloyd C. L. Hollenberg and Nicholas S. Witte. 1994. General nonperturbative estimate of the energy density of lattice Hamiltonians. *Physical Review D* 50, 5 (1994), 3382–3386. https://doi.org/10.1103/physrevd.50.3382

[23] Stephen P. Jordan, David Gosset, and Peter J. Love. 2010. Quantum-Merlin-Arthur–complete problems for stoquastic Hamiltonians and Markov matrices. *Physical Review A* 81, 3 (2010). https://doi.org/10.1103/physreva.81.032331

[24] Richard M. Karp and Richard J. Lipton. 1980. Some connections between nonuniform and uniform complexity classes. *Proceedings of the twelfth annual ACM symposium on Theory of computing - STOC 80* (1980). https://doi.org/10.1145/800141.804678

[25] Alexei Yu. Kitaev, Alexander Shen, and Mikhail N. Vyalyi. 2002. Classical and quantum computing. In *Quantum Information*. Springer New York, New York, NY, 203–217. https://doi.org/10.1007/978-0-387-36944-0_13

[26] Richard E. Ladner and Nancy A. Lynch. 1976. Relativization of questions about log space computability. *Mathematical Systems Theory* 10, 1 (1976), 19–32. https://doi.org/10.1007/bf01683260

[27] Der-Tsai Lee and Bruce J. Schachter. 1980. Two algorithms for constructing a Delaunay triangulation. *International Journal of Computer & Information Sciences* 9, 3 (1980), 219–242. https://doi.org/10.1007/bf00977785

[28] Jacek Miekisz. 1997. Stable Quasicrystalline Ground States. *Journal of Statistical Physics* 88, 3/4 (1997), 691–711. https://doi.org/10.1023/b:joss.0000015168.25151.22

[29] Michael A. Nielsen and Isaac L. Chuang. 2010. *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge. 676 pages. https://doi.org/10.1017/CBO9780511976667

[30] Christos H. Papadimitriou. 1994. *Computational complexity*. Addison-Wesley.

[31] John P. Perdew, J. A. Chevary, S. H. Vosko, Koblar A. Jackson, Mark R. Pederson, D. J. Singh, and Carlos Fiolhais. 1992. Atoms, molecules, solids, and surfaces: Applications of the generalized gradient approximation for exchange and correlation. *Physical Review B* 46, 11 (1992), 6671–6687. https://doi.org/10.1103/physrevb.46.6671

[32] Raphael M. Robinson. 1971. Undecidability and nonperiodicity for tilings of the plane. *Inventiones mathematicae* 12, 3 (09 1971), 177–209. https://doi.org/10.1007/BF01418780

[33] James D. Watson, Johannes Bausch, and Sevag Gharibian. 2020. The Complexity of Translationally Invariant Problems beyond Ground State Energies. *arXiv e-prints*, Article arXiv:2012.12717 (Dec. 2020), arXiv:2012.12717 pages. https://doi.org/10.48550/arXiv.2012.12717

[34] James D. Watson and Toby S. Cubitt. 2021. Computational Complexity of the Ground State Energy Density Problem. *arXiv e-prints*, Article arXiv:2107.05060 (July 2021), arXiv:2107.05060 pages. https://doi.org/10.48550/arXiv.2107.05060