Undecidability of the Spectral Gap: Supplementary Discussion

Toby Cubitt*[1], David Perez-Garcia†[3,4], and Michael M. Wolf‡[5]

[1]*Department of Computer Science, University College London, Gower Street, London WC1E 6BT, United Kingdom*
[2]*DAMTP, University of Cambridge, Centre for Mathematical Sciences, Wilberforce Road, Cambridge CB3 0WA, United Kingdom*
[3]*Departamento de Análisis Matemático and IMI, Facultad de CC Matemáticas, Universidad Complutense de Madrid, Plaza de Ciencias 3, 28040 Madrid, Spain*
[4]*ICMAT, C/Nicolás Cabrera, Campus de Cantoblanco, 28049 Madrid*
[5]*Department of Mathematics, Technische Universität München, 85748 Garching, Germany*

November 25, 2015

The aim of this supplementary material is to give a detailed overview and discussion of how the proof of our spectral gap undecidability result works. To facilitate the readability of this supplementary material, it is written in a self-contained way, expanding the short explanation given in the main text. (Reference numbers refer to the reference list in main text.)

The most important step in proving undecidability of the spectral gap is to prove undecidability of another relevant quantity: the ground state energy density. Once we have this, it is relatively easy to "lift" it to undecidability of the spectral gap. (More precisely, we give a reduction from the ground state energy density problem to the spectral gap problem.) This this is explained in the main text; technical proofs can be found in [1].

In fact, undecidability of the ground state energy density is stronger than we really need to prove undecidability of the spectral gap. It is sufficient to prove undecidability of the ground state energy *with constant promise gap*, i.e. with a promise that the ground state energy is either $\leq a$ or $\geq b$ for some constant $b - a$.

*t.cubitt@ucl.ac.uk
†dperezga@ucm.es
‡m.wolf@tum.de

Undecidability of the ground state energy density implies that this holds even with a promise gap diverging to infinity.

We start by fixing some notation, most of which was already introduced in the main text. Let $\Lambda(L) := \{1, \dots, L\}^2$ be the set of sites (or vertices) of a square lattice of size $L \in \mathbb{N}$, which we assume to be at least 2. By $\mathcal{E} \subset \Lambda(L) \times \Lambda(L)$ we denote the set of edges of the square lattice, directed such that $(i, j) \in \mathcal{E}$ implies that j lies north or east of i. We assign a Hilbert space $\mathcal{H}^{(i)} \simeq \mathbb{C}^d$ to each site $i \in \Lambda(L)$ and the tensor product $\bigotimes_{i \in S} \mathcal{H}^{(i)}$ to any subset $S \subseteq \Lambda(L)$. To every neighbouring pair $(i, j) \in \mathcal{E}$, we assign a Hermitian operator $h^{(i,j)} \in \mathcal{B}(\mathcal{H}^{(i)} \otimes \mathcal{H}^{(j)})$ describing the interaction between the sites. In addition, we may assign an on-site Hamiltonian given by a Hermitian matrix $h_1^{(k)} \in \mathcal{B}(\mathcal{H}^{(k)})$ to every site $k \in \Lambda(L)$.

Throughout, we consider Hamiltonians that are built up from such nearest-neighbour and on-site terms in a translationally invariant way. That is, when identifying Hilbert spaces, $h_1^{(k)} = h_1^{(l)}$ for all $k, l \in \Lambda(L)$ and $h^{(i',j')} = h^{(i,j)}$ if there is a $v \in \mathbb{Z}^2$ so that $(i', j') = (i + v, j + v)$. The total Hamiltonian

$$H^{\Lambda(L)} := \sum_{(i,j) \in \mathcal{E}} h^{(i,j)} + \sum_{k \in \Lambda(L)} h_1^{(k)} \tag{1}$$

can thus be specified by three Hermitian matrices: a $d \times d$ matrix h_1 and two $d^2 \times d^2$ matrices h_{row} and h_{col}, which describe the interactions between neighbouring sites within any row and column respectively. Hence, it may alternatively be written as

$$H^{\Lambda(L)} = \sum_{\text{rows}} \sum_{c} h_{\text{row}}^{(c,c+1)} + \sum_{\text{columns}} \sum_{r} h_{\text{col}}^{(r,r+1)} + \sum_{i \in \Lambda(L)} h_1^{(i)}. \tag{2}$$

$\max\{\|h_{\text{row}}\|, \|h_{\text{col}}\|, \|h_1\|\}$ is called the local interaction strength of the Hamiltonian, which we normalise to be 1.

Let $\operatorname{spec} H^{\Lambda(L)} := \{\lambda_0, \lambda_1, \dots\}$ denote the spectrum, i.e. the set of eigenvalues of $H^{\Lambda(L)}$ listed in increasing order $\lambda_0 \leq \lambda_1 \leq \dots$. For clarity we will sometimes write the Hamiltonian in question as an argument of the eigenvalues. $\lambda_0(H^{\Lambda(L)})$ will be called the *ground state energy*, and the corresponding eigenvector the *ground state*. A Hamiltonian $H^{\Lambda(L)}$ *frustration-free* if its ground state energy is zero whilst all $h^{(i,j)}, h_1^{(k)}$ are positive semi-definite. That is, a ground state of a frustration-free Hamiltonian minimises the energy of each interaction term individually.

1 Ground state energy density

Consider the square lattice $\Lambda(L)$ with edge length $L \in \mathbb{N}$ but in the general case of $\nu \in \mathbb{N}$ spatial dimensions, supporting a translationally-invariant, nearest-neighbour Hamiltonian

$$H^{\Lambda(L)} := \sum_{(i,j) \in \mathcal{E}} h^{(i,j)} + \sum_{k \in \Lambda(L)} h_1^{(k)}, \tag{3}$$

and let c be its local interaction strength. Assume open boundary conditions. The *ground state energy density* is defined as

$$E_\rho := \lim_{L\to\infty} E_\rho(L), \quad \text{where} \quad E_\rho(L) := L^{-\nu}\lambda_0(H^{\Lambda(L)}). \tag{4}$$

The following simple argument shows that this limit is indeed well defined. Consider two lattices of different sizes $L, L' \in \mathbb{N}$ such that $L = nL'$ for some $n \in \mathbb{N}$. Assume w.l.o.g. that the interaction terms in the Hamiltonian are all positive semi-definite. Then $H^{\Lambda(L)}$ is, as an operator, lower bounded by the sum of n^ν translates of $H^{\Lambda(L')}$. So we have that

$$\lambda_0(H^{\Lambda(L)}) \geq n^\nu \lambda_0(H^{\Lambda(L')}). \tag{5}$$

On the other hand, we can use a product of n^ν copies of the ground state of $H^{\Lambda(L')}$ in order to obtain an analogous upper bound on the ground state energy of $H^{\Lambda(L)}$ of the form

$$\lambda_0(H^{\Lambda(L)}) \leq n^\nu \lambda_0(H^{\Lambda(L')}) + 2\nu n^\nu L'^{(\nu-1)} c. \tag{6}$$

Dividing both inequalities by L^ν we are left with

$$E_\rho(L') \leq E_\rho(L) \leq E_\rho(L') + \frac{2\nu c}{L'}. \tag{7}$$

Hence, an interval of order $O(1/L')$ contains both lim inf and lim sup of $E_\rho(L)$ so that both must coincide, which proves that $\lim_{L\to\infty} E_\rho(L)$ is well defined.

The ground state energy density is an important physical quantity in its own right, as well as being our main stepping stone to the spectral gap results. It is therefore worth a brief digression to note that the above argument also shows that the ground state energy density can be computed to any constant precision $\delta > 0$ by exact diagonalisation of $H^{\Lambda(L')}$ for any $L' > 2\nu c/\delta$. This immediately implies that the ground state energy density problem is *decidable* if we provide a finite promise gap δ:

Proposition 1 (Decidability of g.s. energy density with promise gap) *Let $\delta > 0$ be a computable number and consider translationally-invariant, nearest-neighbour Hamiltonians on a $\nu-$dimensional square lattice with open boundary conditions, finite local Hilbert space dimensions and algebraic matrix entries. Then determining whether $E_\rho \leq 0$ or $E_\rho \geq \delta$ is decidable under the promise that $E_\rho \notin (0, \delta)$.*

(Since the real algebraic numbers form a computably ordered field whose cardinality is countably infinite, we can think of the input Hamiltonian as being encoded in a single natural number.)

This is in sharp contrast to the following, whose proof will be the aim of this supplementary material.

Theorem 2 (Undecidability of g.s. energy density) *Let $d \in \mathbb{N}$ be sufficiently large but fixed, and consider translationally-invariant, nearest-neighbour Hamiltonians on a 2D square lattice with open boundary conditions, local Hilbert space dimension d, algebraic matrix entries, and local interaction strength ≤ 1. For such Hamiltonians, determining whether $E_\rho = 0$ or $E_\rho > 0$ is an undecidable problem.*

In the next two sections, we will discuss two approaches to proving Theorem 2 that do *not* work. Section 2 describes a purely classical construction based on Wang tilings. This gives a Hamiltonian with the correct spectral properties, but it necessarily requires unbounded local Hilbert space dimension d. In Section 3, we briefly review the Feynman-Kitaev-style local Hamiltonian constructions used in recent QMA-hardness results. By a new and careful application of the quantum phase estimation technique (described in Section 4), and extending ideas from Gottesman and Irani[2] (Section 5), this approach can give a Hamiltonian with constant local dimension. But it necessarily fails to have the required spectral properties. Finally, in Section 6, we discuss how combining ideas from *both* these approaches allows us to achieve the required spectral properties whilst simultaneously keeping the local dimension constant, the details being given in Section 7.

2 Wang Tilings

The first approach one might consider to proving undecidability of the ground state energy is to note the close relationship between tilings and (classical) Hamiltonians, recalling Berger's classic result that the tiling problem is undecidable[3].

We will soon see that this approach is too weak to prove Theorem 2. Nonetheless, not only is it helpful to understand why this approach breaks down, the much more involved construction required to prove our main result will also make use of Wang tilings, albeit in a less direct way.

A unit square whose edges are coloured with colours chosen from a finite set is called a *Wang tile*. A finite set $\mathcal{T}$ of Wang tiles is said to *tile the plane* $\mathbb{Z}^2$ if there is an assignment $\mathbb{Z}^2 \to \mathcal{T}$ so that abutting edges of adjacent tiles have the same colour. The result we will use is the fact that there exists no algorithm which, given any set of tiles as input, decides whether or not this set can tile the plane—*tiling is undecidable*[3]. (Here, rotating or reflecting the tiles is not allowed – this would make the problem trivial.[1])

A tiling problem can easily be represented as a ground state energy problem for a *classical* Hamiltonian (i.e. one that is diagonal in a product basis). The mapping is straightforward: with the identification $\mathcal{T} = \{1, \ldots, T\}$ we assign a Hilbert space

[1] If one slightly modifies the rules of the game and requires complementary rather than matching colours for abutting edges, then the problem remains undecidable even if rotations and reflections are allowed.

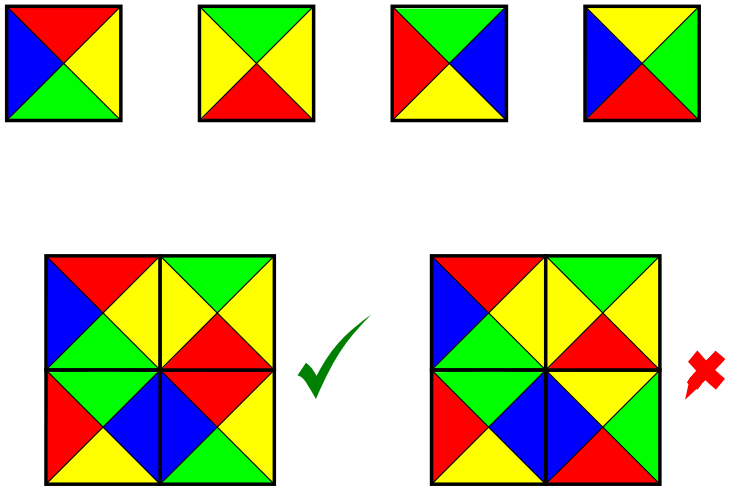

Supplementary Figure 1: Valid and invalid tilings. Examples of valid and invalid tilings (bottom) of a set of four Wang tiles (top).

$\mathcal{H}^{(i)} \simeq \mathbb{C}^T$ to each site i of a square lattice, and define the local interactions via

$$h^{(i,j)} := \sum_{(m,n)\in C^{(i,j)}} |m\rangle\langle m|_{(i)} \otimes |n\rangle\langle n|_{(j)}\,, \tag{8}$$

where the set of constraints $C^{(i,j)} \subseteq \mathcal{T} \times \mathcal{T}$ includes all pairs of tiles (m, n) which are incompatible when placed on adjacent sites i and j. The overall Hamiltonian on the lattice $\Lambda(L)$ is then

$$H_c^{\Lambda(L)} := \sum_{(i,j)\in\mathcal{E}} h_c^{(i,j)}. \tag{9}$$

Undecidability of the ground state energy of H_c with a promise gap of 1 now follows immediately from undecidability of tiling, and this gives undecidability of the ground state energy *density* in the case of open boundary conditions. (A full proof of this is described in [1].)

However, there is a crucial and fundamental limitation to this approach: there is no upper-bound on the local dimension of the Hamiltonian. Rather, the local Hilbert space dimension grows with the number of tile types. And we cannot impose any bound on the latter, or else the tiling problem is restricted to a finite number of cases and is trivially decidable by case enumeration.

On the other hand, this does already allow us to prove a weaker form of our main result: undecidability of the spectral gap for families of Hamiltonians with no constraint on the local dimension.

Nonetheless, from a physical perspective e.g. of characterising the phase diagram of a system, it is unreasonable to allow the local Hilbert space dimension

to grow arbitrarily large – or indeed to change at all – as the parameters of the Hamiltonian are varied. So we are still a long way from proving our main result.

Fundamentally, the problem is that the corresponding Hamiltonians are too simple. For a problem to be algorithmically undecidable, it must admit a countably infinite number of problem instances. If the local Hilbert space dimension is fixed, a translationally-invariant Hamiltonian is completely specified by a finite number of matrix elements defining its local interactions. The only way to encode a countably infinite number of problem instances is to exploit the fact that the matrix elements themselves can take a countable infinity of values (e.g. arbitrarily precise rational numbers, or even arbitrary computable numbers). Whereas the above tiling approach is only sensitive to the pattern of non-zero matrix elements.

To overcome this, we will need an inherently quantum approach, which is the topic of the next section.

3 QMA constructions

There is by now a standard approach to proving complexity-theoretic hardness results for local Hamiltonian problems. The idea, which dates back to Feynman[4] and was significantly developed by Kitaev[5] and others[6;7;8;2], is to construct a Hamiltonian whose ground state encodes the history of a quantum computation *in superposition*. If we divide the system into two parts, a "clock register" and a "computational register", then the desired ground state has the form:

$$\frac{1}{\sqrt{T+1}} \sum_{t=0}^{T} |t\rangle \, |\psi_t\rangle \,, \tag{10}$$

where $|\psi_t\rangle$ denotes the state of the computation after t steps. This superposition over the history of a computation is often called a *computational history state*.

It is not difficult to construct a Hamiltonian with this as its unique ground state. The difficult part is to implement Feynman's idea using a *local* Hamiltonian. This was first done by Kitaev[5], who showed how to construct such a Hamiltonian out of 5-body terms. Kempe, Kitaev and Regev[6] improved this to 2-body, Oliveira and Terhal[7] to nearest-neighbour two-qubit interactions on a 2D square lattice, and Aharonov et al.[8] to nearest-neighbour two-body interactions on a line.

All of these constructions exploit the fact that the interactions can differ from site to site, in order to encode arbitrary computations. Indeed, for translationally-invariant, nearest-neighbour interactions on a regular lattice, the entire Hamiltonian is specified by a finite number of two-body terms (and possibly one single-body term), and it might appear that there are not enough parameters available to encode arbitrary quantum computations. However, in a remarkable paper, Gottesman and Irani[2] showed how to construct a translationally-invariant Hamiltonian which has

as its ground state a computational history state for an arbitrary computation.[1]

The aim of all these local Hamiltonian constructions was to prove QMA-hardness of the finite-size ground state energy problem for the corresponding class of Hamiltonians, by encoding the quantum computation that verifies the witness for a QMA problem and adding a local term to the Hamiltonian that gives an additional energy penalty to the "no" output.

An obvious approach to constructing a Hamiltonian with undecidable ground state energy is to use one of these local Hamiltonian constructions to encode the evolution of a universal (reversible or quantum) Turing Machine, instead of a QMA witness verifier, and give an additional energy penalty to the halting state. If we use the Gottesman-Irani construction[2], the resulting Hamiltonian will consist of translationally-invariant, nearest-neighbour interactions on a line. Since the Halting Problem is undecidable, and the ground state energy depends on whether or not the computation halts, the ground state energy of this Hamiltonian would seem to be undecidable. As in the tiling approach, this is certainly too weak to prove undecidability of the energy *density*. But one might hope that it is sufficient to prove undecidability of the ground state energy.

However, this Feynman-Kitaev Hamiltonian approach does not even achieve the weaker result of the tiling approach. There are now two crucial problems:

(i). The Halting Problem is undecidable for the universal Turing Machine on arbitrary input (or for an arbitrary Turing Machine running on fixed input). As in the tiling approach of the previous section, it is not at all clear how to encode this countably infinite family of problems into the constant number of matrix elements describing the nearest-neighbour interaction.

(ii). The promise gap δ in all known local Hamiltonian constructions, and in particular the translationally-invariant construction of[2], scales inverse-polynomially in the system size. Thus (assuming the limits exist)

$$\lim_{L\to\infty} \lambda_0(H^{\Lambda(L)}) = \begin{cases} \lim_{L\to\infty} f(L) & \text{non-halting} \\ \lim_{L\to\infty} f(L) + \frac{1}{\text{poly}(L)} & \text{halting} \end{cases} \\ = \lim_{L\to\infty} f(L) \tag{11}$$

for some function f. I.e. the ground state energy in the thermodynamic limit is identical in both the halting and non-halting cases.

These issues are inherent to the spectral gap problem for many-body quantum systems, where the question is only meaningful or interesting in the thermodynamic

[1] In fact, the two-body interaction in the Gottesman-Irani construction[2] is the same *fixed* interaction for *all* problem instances. The input is specified by the only remaining free parameter: the length of the chain!

limit of Hamiltonians with regular structure (of which translational-invariance is the simplest case). Thus they cannot be side-stepped, and overcoming them is the main task in proving the result.

In the following section, we will see that overcoming (i), whilst challenging, can be achieved by exploiting the ability to encode *quantum* computation. Indeed, this will essentially be the only quantum part of our construction.

However, (ii) presents a more serious obstacle to the history state approach. There is an inherent trade-off between run-time and promise gap in Kitaev-style local Hamiltonian constructions, and the run-time is directly related to the system size when the Hamiltonian is constrained to a lattice. But we are necessarily working in the thermodynamic limit of arbitrarily large lattice size. We therefore need a *constant* promise gap, independent of the length of the computation. This cannot be achieved by any known local Hamiltonian construction, and may well be impossible. Without a constant gap between the halting and non-halting cases, the ground state energy problem becomes trivially decidable in the thermodynamic limit. We discuss how we overcome this obstacle in Section 6.

4 Constant local dimension

To overcome the unbounded local dimension obstacle we faced in Sections 2 and 3, we must find a way of encoding the countably infinite family of Halting Problem instances into the finite number of matrix elements describing the local interactions of a system with fixed local Hilbert space dimension.

If we encode the evolution of a quantum Turing Machine into the ground state of a local Hamiltonian using a Feynman-Kitaev-style construction, as described in the previous section, the local dimension will depend on the number of internal states and alphabet size of the QTM. Whichever universal Turing Machine we choose to encode, that particular TM will have a fixed state space and alphabet size. But to encode the Halting Problem, we need a way to feed any desired input to this encoded universal TM. It is not difficult to construct a special-purpose classical TM which outputs any given string, starting from a fixed input. But, exactly analogous to the Wang tiling constructions of Section 2, if there is no upper-bound on the number of different strings that we must be able to produce, then either the number of internal states or the alphabet size of the Turing Machine is necessarily unbounded. This is no use to us, as it would again lead to a family of Hamiltonians with unbounded local dimension.

The only way we can hope to generate arbitrarily long strings using constant alphabet size and a constant number of internal states is to use a genuinely quantum construction. The transition rules of a QTM can have arbitrarily computable numbers as coefficient[9]. (In fact, algebraic numbers will suffice for our purposes.) So, whereas for given alphabet size and number of internal states there is only a

finite number of different classical deterministic TMs, there are a countably infinite number of different QTMs. We will show how the string we want to produce can be encoded in the transition rule coefficients of a QTM, in such a way that the QTM writes out this string and then halts *deterministically*.

At first sight, this might appear to violate the Busy Beaver bound on the runtime of a TM[10], or the Holevo bound on the amount of information that can be extracted from a finite-dimensional quantum state[11], or other results that limit the amount of information that can be extracted from a finite-size system. However, a little more thought reveals there is no contradiction here.

Indeed, something similar is already possible for classical probabilistic Turing Machines. It is a straightforward exercise to construct a classical probabilistic TM with fixed alphabet and number of internal states which, given access to a coin with bias p, outputs the binary expansion of p with high probability, in expected runtime that is a function of the length of the binary expansion. What is perhaps more surprising is that Quantum Turing Machines allow this to be done *deterministically*.

The reason this does not violate the Busy Beaver theorem is that, to simulate a probabilistic or quantum TM on a deterministic TM, the alphabet and/or internal state size must grow with the precision of the entries in the probabilistic or quantum transition function. Nor is there any contradiction with the Holevo bound. We are not encoding the string in a finite-dimensional quantum state, or even in multiple copies of a quantum state. We are encoding the string in the unitary transition rules of a QTM, which we get to apply as many times as we like on any quantum state we like. Applying the transition rules to a fixed quantum state and performing quantum state tomography would already allow us to extract the information encoded in the transition rules to arbitrary precision. Again, perhaps the only somewhat surprising aspect is that, by exploiting the full power of quantum computation, we can recover the encoded string *exactly*, regardless of how long the string is.

The idea behind our construction is to use the quantum phase estimation algorithm[12] (running on a QTM) to extract a phase φ which we encode in a single-qubit unitary

$$U_\varphi = \begin{pmatrix} 1 & 0 \\ 0 & e^{2\pi i\varphi} \end{pmatrix}, \tag{12}$$

thereby writing out its binary fraction expansion to the tape. However, for technical reasons that appear to be insurmountable, it is crucial to our proof that the phase estimation be carried out *exactly*, not merely with high probability. Furthermore, the QTM should halt deterministically after a time that depends only on the input. Without these properties, the matrix elements of the Hamiltonians we construct will not be computable, and Theorem 2 becomes vacuous.

Recall from[12] that the phase estimation algorithm acts on N output qubits initialised on $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) = H|0\rangle$ (H the Hadamard matrix $\left(\begin{smallmatrix} 1 & 1 \\ 1 & -1 \end{smallmatrix}\right)$) and one

auxiliary qubit initialised to $|1\rangle$, and has two stages. In the first stage, it loops throughout all output qubits and applies on the auxiliary qubit the unitary $U_\varphi^{2^{n-1}}$ controlled by the n^{th} output qubit (see Figure 2.) The final state of the output

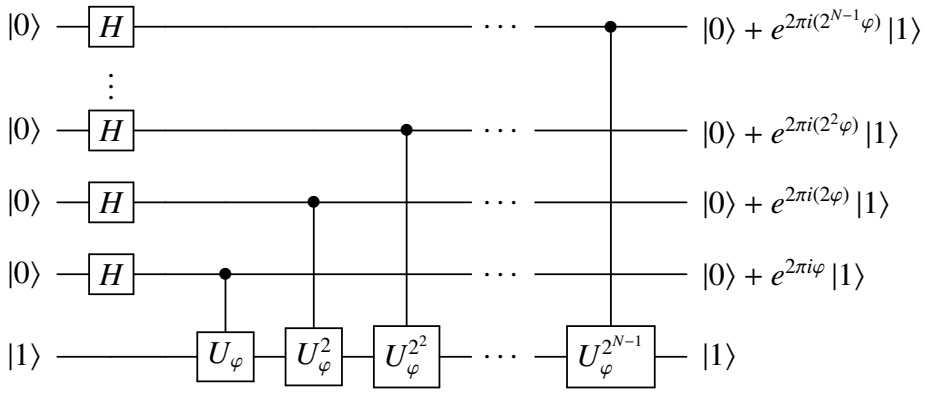

Supplementary Figure 2: Quantum phase estimation, control-phase stage. The first stage of the quantum phase estimation circuit for φ (cf. Fig. 5.2 in Ref. 12).

register after this stage is

$$\frac{1}{2^{\frac{N}{2}}}\left(|0\rangle + e^{2\pi i 0.\varphi_N}|1\rangle\right)\left(|0\rangle + e^{2\pi i 0.\varphi_{N-1}\varphi_N}|1\rangle\right)\cdots \\ \cdots\left(|0\rangle + e^{2\pi i 0.\varphi_2\ldots\varphi_{N-1}\varphi_N}|1\rangle\right)\left(|0\rangle + e^{2\pi i 0.\varphi_1\varphi_2\ldots\varphi_N}|1\rangle\right) \quad (13)$$

where φ_k denotes the k'th digit in the binary fraction expansion of φ.

If φ has an exact binary fraction expansion with $n = |\varphi|$ binary digits, i.e. $\varphi = 0.\varphi_1\varphi_2\cdots\varphi_n$ written in binary with $\varphi_k \in \{0,1\}$ and $\varphi_n = 1$, then the first $N - |\varphi|$ qubits are in the $|+\rangle$ state (which can be converted to $|0\rangle$'s by Hadamard gates). The next register (the one in position $N - |\varphi| + 1$) is in the $|-\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$ state (which is converted to $|1\rangle$ by the Hadamard gate). This qubit is therefore singled out.

The second stage is to perform the inverse Quantum Fourier Transform on the last n output qubits (see Figure 3), which leaves the n digits of the exact binary fraction expansion of φ written on the last n output qubits. For that, the only gates that are needed are the Hadamard and the controlled-$U_{2^{-|\varphi|}}$ gate.

This shows that, just using the Hadamard, the controlled-U_φ and the controlled-$U_{2^{-|\varphi|}}$ gates (both depending only on the external parameter φ), one is able to design a quantum circuit that outputs the $|\varphi|$ bits of φ *exactly* for the case in which the number of digits in the binary fraction expansion of φ is finite and $\leq N$. By constructing explicitly a QTM which implements the above circuit, it is not difficult – though tedious – to prove the following theorem. The details can be found in[1].

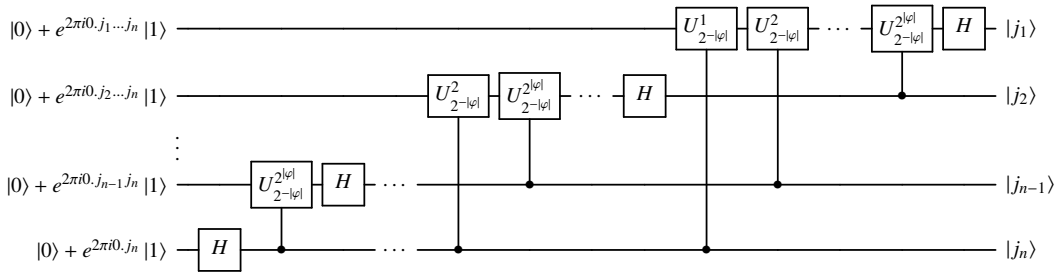

Supplementary Figure 3: Quantum phase estimation, inverse QFT stage. The inverse QFT stage of the quantum phase estimation circuit (cf. Fig. 5.1 in Ref. 12).

Theorem 3 (Phase-estimation QTM) *There exists a family of properly behaved QTMs P_n indexed by $n \in \mathbb{N}$ with the following properties:*

(i). Both the alphabet and the set of internal states are identical for all P_n; only the transition rules differ.

(ii). On input $N \geq |n|$ written in unary, P_n has deterministic head movement, halts deterministically after $O(\mathrm{poly}(N)2^N)$ steps, uses $N + 3$ space, and outputs the binary expansion of n (padded to N digits trailing 0's). (As above, $|n|$ denotes length of the binary expansion of n.)

(iii). For each choice of states p, q, alphabet symbols σ, τ and directions D, the transition amplitude $\delta(p, \sigma, \tau, q, D)$ is, independently of n, one of the elements of the set

$$\left\{0, 1, \pm\tfrac{1}{\sqrt{2}}, e^{i\pi\varphi}, e^{i\pi 2^{-|\varphi|}}\right\}, \tag{14}$$

the only dependence on n being that implicit in φ which is defined as the rational number whose binary fraction expansion contains the digits of n after the decimal point.

Some remarks are in order:

- The input N does *not* determine the output string that gets written to the tape; it only determines the number of qubits in the associated quantum circuit and hence the number of binary digits in the output. As happens in the quantum circuit analysed above, the number represented by that output is determined (up to padding with trailing zeros) by the choice of the parameter n (or equivalently φ) for the QTM P_n.

- By "properly behaved QTM" we mean *well-formed*, *normal form* and *unidirectional* according to the standard definitions in[9], to which we also refer for the formal definition of a QTM.

- A QTM is said to have *deterministic head movement* on an input if, when started with that input, the QTM never enters a configuration in which the head is in a superposition of different locations.

- The form of the elements of the transition amplitudes simply reflects the fact that only Hadamard, controlled-U_φ and controlled-$U_{2-|\varphi|}$ gates are being used in the circuit. The definition of φ for desired output n is also clear by looking at the output of the phase estimation circuit.

- The universal QTM construction of Bernstein and Vazirani[9] shows that any quantum circuit can be implemented on a QTM up to some error, not exactly. Therefore, one cannot rely on previous results and must construct explicitly the desired QTM.

- If $N < |n|$, so that the number of binary digits in the phase exceeds the number of qubits in the circuit, we make no claim about the behaviour of the QTM; it could leave an arbitrary string (or even quantum state) written on its tape, or it could even run forever.

In this way, if we consider the phase-estimation QTM P_n and feed its output n into a universal TM, the local Hilbert space dimension of the Hamiltonian encoding this sequence of Turing Machines will be constant (assuming the properties of the Gottesman-Irani construction[2] carry over), solving the first problem highlighted in Section 3. The next section discusses in more detail how to encode computation into a Hamiltonian.

5 Translational invariance

Gottesman and Irani[2] showed how to construct a *fixed* Hamiltonian on a 1D chain, that can encode in its ground state the evolution of a QTM for a number of time-steps polynomial in the length L of the chain. The input to the QTM in their construction is determined by the chain length. They accomplish this by first constructing a translationally-invariant clock to keep track of time, which runs for a total of L steps. This clock drives a binary counter TM for L steps, leaving the binary representation of L written on the tape. The clock is then reset, and switches over to driving the QTM. The binary counting TM and the QTM share the same tape, so the input to the QTM is the binary representation of L. It is important to note that the local Hilbert space dimension in the Gottesman-Irani construction[2] depends only on the alphabet size and number of internal states of the Turing Machines (plus some constant multiplicative overhead for the clock).

However, in our case we are interested in the thermodynamic limit. The Gottesman and Irani[2] result per se does not achieve what we need. We cannot use

the chain length to encode the input to the QTM, as we are only concerned with the limit as the length tends to infinity. Instead, we want to encode the input to the QTM in the Hamiltonian itself, and carry out the same computation for any chain length.[1]

Given our quantum phase estimation QTM from the previous section, it is clear how we should adapt the Gottesman-Irani construction[2] to achieve what we need. Instead of the binary counter TM, we first run our phase estimation QTM. Provided the chain length is sufficiently large that $L > |n|$, the phase-estimation QTM will write the desired string to the tape and then halt. We then switch to driving a universal reversible TM which shares the same tape.

Whilst this approach does ultimately work, there are a number of technical issues to overcome. In particular, the length of the computation in the Gottesman-Irani construction[2] is limited by the maximum number of time-steps that the clock can encode, which is $O(\text{poly } L)$. Whereas our phase-estimation QTM requires time $O(\text{poly}(L)2^L)$ on input of length L.

There are a number of ways around this. Perhaps the simplest – and the one we adopt – is to modify the Gottesman-Irani clock construction[2] to count in base ζ instead of unary, so that the clock can encode at least $\Omega(\zeta^L)$ time-steps, at the price of substantially complicating the analysis. A full proof can be found in[1], including a detailed analysis of the required spectral properties of the resulting Hamiltonian that was only sketched informally in Gottesman and Irani[2].

Theorem 4 (Local Hamiltonian QTM encoding) *Let* $\mathbb{C}^d = \text{span}\{\circledcirc, \circledcirc\} \oplus \mathbb{C}^C \otimes \mathbb{C}^Q$ *be the local Hilbert space of a 1-dimensional chain of length L (*$\mathbb{C}^C$ *corresponds to the clock register and* $\mathbb{C}^Q$ *to the computational register), so that the Hilbert space of the whole chain is* $\mathcal{H}(L) = (\mathbb{C}^d)^{\otimes} L$. *For any properly behaved Quantum Turing Machine M with alphabet* Σ, *set of states Q and transition amplitude function* δ *and for any constant* $K \in \mathbb{N}$, *we construct a two-body interaction* $h \in \mathcal{B}(\mathbb{C}^d \otimes \mathbb{C}^d)$ *such that the 1-dimensional, translationally-invariant, nearest-neighbour Hamiltonian* $H(L) = \sum_{i=1}^{L+1} h^{(i,i+1)} \in \mathcal{B}(\mathcal{H}(L))$ *on the chain of length* $L \geq K + 3$ *has the following properties:*

(i). d depends only on the alphabet size and number of internal states of M.

(ii). $h \geq 0$, *and the overall Hamiltonian H(L) is frustration-free for all L (i.e. the ground energy of H(L) is* 0 *for all L).*

[1] In their paper[2], Gottesman and Irani also provide a construction for infinitely long chains. However this works by adding terms to the Hamiltonian which effectively break up the chain into segments of length L, and the finite chain-length construction then goes through independently for each segment. This is also not what we want, as it means that, despite the infinitely long chain, there is still a finite bound L on the space available for the computation.

(iii). When restricted to the subspace $S_{br} \subset \mathcal{H}(L)$ with $|\circledless\rangle, |\circledgtr\rangle$ at the left and right ends of the chain, respectively, the unique ground state of $H(L)|_{S_{br}}$ is a computational history state encoding the evolution of M on input consisting of a string of $L-K-3$ '1's, running on a finite tape segment of length $L-3$.

Moreover, if M has deterministic head movement on input consisting of a string of $L-K-3$ '1's, then:

(iv). The computational history state always encodes $\Omega(|\Sigma \times Q|^L)$ time-steps. If M halts in fewer than the number of encoded time steps, exactly one $|\psi_t\rangle$ has support on a state $|\top\rangle$ that encodes a halting state of the QTM. The remaining time steps of the evolution encoded in the history state leave M's tape unaltered, and have zero overlap with $|\top\rangle$.

(v). If M runs out of tape within a time T less than the number of encoded time steps (i.e. in time-step $T+1$ it would move its head before the starting cell or beyond cell $L-3$), the computation history state only encodes the evolution of M up to time T. The remaining steps of the evolution encoded in the computational history state leave M's tape unaltered.

(vi). If P_n is a family of QTMs which satisfies part (iii) of Theorem 3, then h has the following form

$$h = A + (e^{i\pi\varphi}B + e^{i\pi 2^{-|\varphi|}}C + \text{h.c.}) \tag{15}$$

with $B, C \in \mathcal{B}(\mathbb{C}^d \otimes \mathbb{C}^d)$ independent of n with coefficients in $\mathbb{Z}$, and $A \in \mathcal{B}(\mathbb{C}^d \otimes \mathbb{C}^d)$ Hermitian, independent of n, with coefficients in $\mathbb{Z} + \frac{1}{\sqrt{2}}\mathbb{Z}$.

Thus we have succeeded in overcoming the constant local dimension obstacle of the QMA constructions discussed in Section 3. However, simply adding a local term to this Hamiltonian that gives an additional energy penalty to the halting state does not work, for the reasons discussed previously: the energy difference between the halting and non-halting cases decreases polynomially with the system size. So all dependence of the ground state energy on the outcome of the computation still vanishes in the thermodynamic limit.

6 The thermodynamic limit

The more challenging obstacle of the thermodynamic limit still remains. To address this, we first return to tiling problems. However, instead of using these blindly, as in the Wang tiling approach described in Section 2 where only undecidability of tiling was used, we prove and then exploit very particular properties of an aperiodic tiling

due to Robinson[13]. These will allow us, using the ideas discussed in Sections 4 and 5, to encode in a quantum local Hamiltonian the execution of *many copies* of the *same* universal Turing Machine running on the same chosen input, but running on tapes of all possible finite lengths and for every possible finite run-time.

The idea of encoding the evolution of many copies of the same Turing Machine on tapes of all possible finite lengths, instead of encoding the evolution of a single Turing Machine on an infinite tape, dates back to Berger's original proof of undecidability of tiling[3]. [13] also used the same idea in his simplification of Berger's proof. However, our reason for exploiting this idea is somewhat different. In our case, it is the crucial ingredient that allows us to decouple the energy dependence of the computational history ground state – a purely quantum property – from the overall system size.

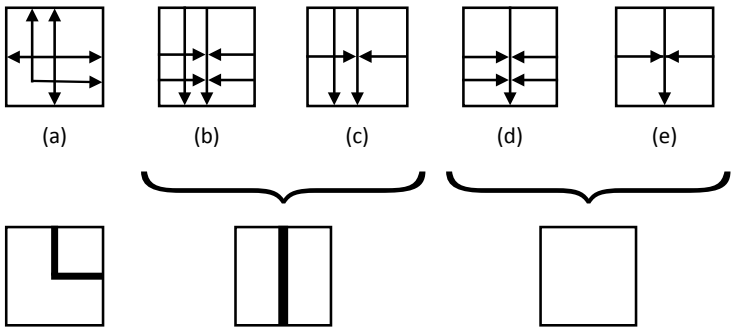

Supplementary Figure 4: The five basic tiles of Robinson's tiling. The tile set includes the basic tiles (top) and all rotations and reflections thereof. These tiles are unambiguously identified by a simplified schematic representation (bottom) showing only the orientation of complete off-centre arrows.

The Robinson tile set *can* tile the infinite plane, but only aperiodically. The aperiodic pattern generated by a Robinson tiling is shown in Figure 5, which has the crucial (for our purposes) quasi-periodic structure consisting of squares of increasing size. The pattern is produced by the five basic tiles shown in Figure 4 and all rotations and reflections thereof, together with extra colour and parity markings not shown in the figure (see[13] or[1] for details). The top edges of the squares in Figure 5 are called *segments*. These have sizes $4^n + 1$, $n \in \mathbb{N}$, and repeat with period 2^{2n+1}.

We have already seen in Section 2 that tiling problems can easily be turned into translationally-invariant classical Hamiltonians. So we can readily turn this into a classical Hamiltonian whose ground state has the same quasi-periodic structure as the Robinson tiling. The idea is to add another "layer" on top of this tiling Hamiltonian, and use this second "quantum" layer to place copies of the history-state Hamiltonian from the previous section along the top edges (segments) of

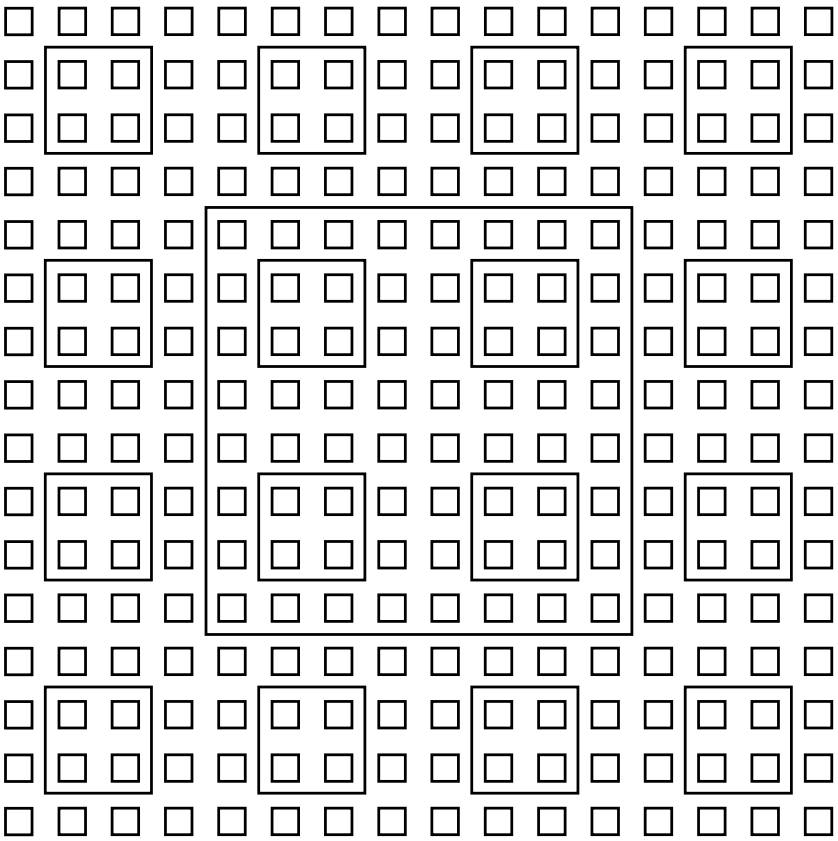

Supplementary Figure 5: Robinson's aperiodic tiling. Schematic depiction of the pattern produced by Robinson's aperiodic tile set. Only the part of the pattern relevant to the proof is shown, using the simplified tile representation of Figure 4. (Tile boundaries are omitted for clarity; cf. Figure 3 in the main text.)

all the squares in Figure 5. (Of course, this has to be done by adding additional translationally-invariant local terms to the Hamiltonian that effectively restrict where the 1D Hamiltonian acts, not by literally restricting the 1D Hamiltonian to the square borders, which would break translational invariance.)

In this way, we construct a 2D translationally-invariant local Hamiltonian whose ground state contains an encoding of the evolution of the universal Turing Machine along *each* segment in the Robinson tiling. The effective tape length of this Turing Machine is limited only by the size of the segment it "runs" on. But the Robinson tiling contains segments of all sizes of the form $4^n + 1$. So this Hamiltonian encodes Turing Machines with all possible power-of-4 tape lengths. Note that all of these Turing Machines are running on the *same* input, encoded in the phase φ which appears in a matrix element of the translationally-invariant local interaction.

If the universal Turing Machine eventually halts on this input, then for all segments above a certain size the effective tape will be sufficiently long for the machine to halt before it runs out of tape space. If we add a translationally-invariant local term that penalises the halting state, then the ground state will pick up an additional energy from the history states encoding Turing Machines that halt. This energy still decreases with the size of the system it acts on. But, crucially, this size is now the size of the segment it is "running" on, *not* the overall system size. We have decoupled the ground state energy from the overall system size. It now depends only on the space required for the universal TM to halt.

If the universal Turing Machine never halts on the given input, then it will not pick up any additional energy, providing the tape is sufficiently long for the phase-estimation QTM to operate correctly. However, if the effective tape length (segment length) is too small to contain the number of digits in the exact binary fraction expansion of the phase φ, then the quantum state left on the tape by the phase-estimation QTM is some arbitrary quantum state (see Section 4). In this case, we cannot assume that the universal Turing Machine runs forever, since its input may be corrupted. Thus, even in the non-halting case, the ground state will pick up some additional energy from segments that are too small. But, crucially, we know the maximum size of these segments: it is determined by the number of binary digits in φ – an external parameter of the Hamiltonian *which we choose*. The contribution to the energy density from the small segments is therefore a computable number, α, given e.g. by brute-force diagonalisation of the Hamiltonian on the small segments. Indeed, since this is an eigenvalue computation, this quantity is in fact an algebraic number. The number of segments of any given size grows quadratically in the lattice size, so the total energy contribution from the small segments (which is also present in the halting case) can be removed by subtracting an appropriately weighted one-body term $-\alpha\mathbb{1}$ from the Hamiltonian, which simply amounts to an overall energy shift.

We have managed to construct a family of Hamiltonians whose ground state energies depend on the solution of the corresponding Halting Problem. However, this is still not sufficient to prove Theorem 2. The difference in ground state energy between the halting and non-halting cases (the "promise gap") depends inverse-polynomially on the space required for the universal Turing Machine to halt. Thus, not only does this fail to provide a uniform bound on the promise gap as required, this promise gap is uncomputable.

In fact, this is a limitation of the above analysis, not the true situation. Instead of being uncomputably small, as the above argument suggests, the true promise gap for this Hamiltonian is infinite! To show this requires a more careful analysis of the low-lying eigenvalues (low-energy excitations). It is easy to see that the eigenstate consisting of a valid tiling together with computational history states along the segments has energy that diverges with the lattice size in the halting case; once the lattice is large enough, the number of segments that are sufficiently large for the encoded Turing Machine to halt grows quadratically, and each of them contributes a small but non-zero energy. The difficulty is that, since its energy diverges, this eigenstate may not be the ground state in the halting case; we might be able to lower the energy by introducing defects in the tiling layer, thereby "breaking" some of the Turing Machines.

In[1], we prove strong rigidity properties of the Robinson tiling, which show that the tiling pattern is robust against defects: any defect (pair of non-matching adjacent tiles) in the tiling only affects the pattern of segments in a finite ball around the defect. Thus destroying n segments requires $O(n)$ defects, each defect contributes $O(1)$ energy due to the tiling part of the Hamiltonian, and Turing Machines on intact segments contribute $O(1)$ energy. Therefore, no matter how many defects we introduce, the energy will still grow quadratically with the lattice size. The promise gap therefore diverges quadratically in the thermodynamic limit. Thus we in fact obtain a stronger result than the uniform bound on the promise gap required for our main result; since the ground state energy diverges quadratically and our system is on a 2D lattice, this in fact proves undecidability of the ground state energy *density*, as claimed in Theorem 2.

We will devote the rest of this supplementary material to making this argument rigorous, based on Theorem 3, Theorem 4 and the following theorem, which formalises the rigidity property of the Robinson tiling.

Theorem 5 (Segment rigidity) *In any tiling of an $L\times H$ rectangle (width L, height H) with d defects using modified Robinson tiles, the total number of segments of size 4^n+1 is at least $\lfloor H/2^{2n+1}\rfloor(\lfloor L/2^{2n+1}\rfloor-1)-2d$. For the case of no defects $d=0$, this minimum can be attained simultaneously for all n.*

Remark: This is proven in[1] only for a suitable modification of Robinson tiles which still gives the desired pattern of Figure 5.

7 Undecidability of the g.s. energy density – proof

To prove Theorem 2, we will need 1D translationally-invariant Hamiltonians with a particular set of properties. For conciseness, we will call these "Gottesman-Irani Hamiltonians", captured in the following definition:

Definition 6 (Gottesman-Irani Hamiltonian) *Let $\mathbb{C}^Q$ be a finite-dimensional Hilbert space with two distinguished orthogonal states labelled $|\langle\rangle\rangle$, $|\rangle\rangle$. A* Gottesman-Irani Hamiltonian *is a 1D, translationally-invariant, nearest-neighbour Hamiltonian $H_q(r)$ on a chain of length $r+1$ with local interaction $h_q \in \mathcal{B}(\mathbb{C}^Q \otimes \mathbb{C}^Q)$, which satisfies the following properties:*

(i). $h_q \geq 0$.

(ii). $[h_q, |\langle\rangle\rangle\langle\langle\rangle| \otimes |\langle\rangle\rangle\langle\langle\rangle|] = [h_q, |\langle\rangle\rangle\langle\langle\rangle| \otimes |\rangle\rangle\langle\rangle|] = [h_q, |\rangle\rangle\langle\rangle| \otimes |\langle\rangle\rangle\langle\langle\rangle|] = [h_q, |\rangle\rangle\langle\rangle| \otimes |\rangle\rangle\langle\rangle|] = 0$.

(iii). $\lambda_0(r) := \lambda_0(H_q(r)|_{S_{br}}) < 1$, *where S_{br} is the subspace of states with fixed boundary conditions $|\langle\rangle\rangle$, $|\rangle\rangle$ at the left and right ends of the chain, respectively.*

(iv). $\forall n \in \mathbb{N} : \lambda_0(4^n) \geq 0$ *and* $\sum_{n=1}^{\infty} \lambda_0(4^n) < 1/2$.

(v). $\lambda_0(H_q(r)|_{S_<}) = \lambda_0(H_q(r)|_{S_>}) = 0$, *where $S_<$ and $S_>$ are the subspaces of states with, respectively, a $|\langle\rangle\rangle$ at the left end of the chain or a $|\rangle\rangle$ at the right end of the chain.*

Lemma 7 (Robinson + Gottesman-Irani Hamiltonian)
Let $h_c^{\text{row}}, h_c^{\text{col}} \in \mathcal{B}(\mathbb{C}^C \otimes \mathbb{C}^C)$ be the local interactions of the tiling Hamiltonian associated with the modified Robinson tiles. For a given ground state configuration (tiling) of H_c, let $\mathcal{L}$ denote the set of all segments of the lattice, that is, horizontal lines that lie between  *and*  *tiles (inclusive). Let $h_q \in \mathcal{B}(\mathbb{C}^Q \otimes \mathbb{C}^Q)$ be the local interaction of a Gottesman-Irani Hamiltonian $H_q(r)$, as in Definition 6.*

Then there is a Hamiltonian on a 2D square lattice of width L and height H with nearest-neighbour interactions $h^{\text{row}}, h^{\text{col}} \in \mathcal{B}(\mathbb{C}^{C+Q+1} \otimes \mathbb{C}^{C+Q+1})$ such that, for any L, H, the ground state energy

$$\lambda_0(H^{\Lambda(L)}) = \min_{\mathcal{L} \subset \Lambda(L)} \sum_{\ell \in \mathcal{L}} \lambda_0(|\ell|), \tag{16}$$

where the minimisation is over all valid tilings of the $L \times H$ rectangle.

Proof The idea is to sandwich the two Hamiltonians H_c and H_q together in two "layers", so that the overall Hamiltonian acts as H_c on the c-layer, with constraints between the layers that force low-energy configurations of the q-layer to be in the

auxiliary $|0\rangle$ "blank" state, *except* between pairs of $\left|\square\right\rangle$ and $\left|\square\right\rangle$ states appearing in the same row of the c-layer, where the q-layer acts like H_q on that line segment. (See Figure 3 in the main text.)

To this end, define the local Hilbert space to be $\mathcal{H} := \mathcal{H}_c \otimes (\mathcal{H}_e \oplus \mathcal{H}_q) \simeq \mathbb{C}^C \otimes (|0\rangle \oplus \mathbb{C}^Q)$. The Hamiltonian H is defined in terms of the two-body interactions as follows:

$$
\begin{aligned}
h^{\text{col}}_{j,j+1} =& h^{\text{col}}_c \otimes \mathbb{1}^{(j)}_{eq} \otimes \mathbb{1}^{(j+1)}_{eq} && \text{(17a)}\\
h^{\text{row}}_{i,i+1} =& h^{\text{row}}_c \otimes \mathbb{1}^{(i)}_{eq} \otimes \mathbb{1}^{(i+1)}_{eq} && \text{(17b)}\\
&+ \mathbb{1}^{(i)}_{c} \otimes \mathbb{1}^{(i+1)}_{c} \otimes h_q && \text{(17c)}\\
&+ |\square\rangle\langle\square|^{(i)}_c \otimes (\mathbb{1}_{eq} - |\textcircled{<}\rangle\langle\textcircled{<}|)^{(i)} \otimes \mathbb{1}^{(i+1)}_{ceq} && \text{(17d)}\\
&+ (\mathbb{1}_c - |\square\rangle\langle\square|_c)^{(i)} \otimes |\textcircled{<}\rangle\langle\textcircled{<}|^{(i)} \otimes \mathbb{1}^{(i+1)}_{ceq} && \text{(17e)}\\
&+ \mathbb{1}^{(i)}_{ceq} \otimes |\square\rangle\langle\square|^{(i+1)}_c \otimes (\mathbb{1}_{eq} - |\textcircled{>}\rangle\langle\textcircled{>}|)^{(i+1)} && \text{(17f)}\\
&+ \mathbb{1}^{(i)}_{ceq} \otimes (\mathbb{1}_c - |\square\rangle\langle\square|_c)^{(i+1)} \otimes |\textcircled{>}\rangle\langle\textcircled{>}|^{(i+1)} && \text{(17g)}\\
&+ \mathbb{1}^{(i)}_{c} \otimes |0\rangle\langle 0|^{(i)}_e \otimes |\square\rangle\langle\square|^{(i+1)}_c \otimes \mathbb{1}^{(i+1)}_{eq} && \text{(17h)}\\
&+ |\square\rangle\langle\square|^{(i)}_c \otimes \mathbb{1}^{(i)}_{eq} \otimes \mathbb{1}^{(i+1)}_{c} \otimes |0\rangle\langle 0|^{(i+1)}_e && \text{(17i)}\\
&+ \mathbb{1}^{(i)}_{c} \otimes |0\rangle\langle 0|^{(i)}_e \otimes (\mathbb{1}_c - |\square\rangle\langle\square|_c)^{(i+1)} \otimes (\mathbb{1}_{eq} - |0\rangle\langle 0|_e)^{(i+1)} && \text{(17j)}\\
&+ (\mathbb{1}_c - |\square\rangle\langle\square|_c)^{(i)} \otimes (\mathbb{1}_{eq} - |0\rangle\langle 0|_e)^{(i)} \otimes \mathbb{1}^{(i+1)}_{c} \otimes |0\rangle\langle 0|^{(i+1)}_e\,, && \text{(17k)}
\end{aligned}
$$

where $\mathbb{1}_c$, $\mathbb{1}_{eq}$ and $\mathbb{1}_{ceq}$ are the identity operators on the corresponding Hilbert spaces. The Hamiltonian can be understood as follows. (17d) and (17e) force a $|\textcircled{<}\rangle$ in the q-layer whenever there is an $\left|\square\right\rangle$ in the c-layer. (17f) and (17g) do the same with $|\textcircled{>}\rangle$ and $\left|\square\right\rangle$. (17h) and (17i) force non-blank to the left and right of an $\left|\square\right\rangle$ or $\left|\square\right\rangle$, respectively. Finally, (17j) and (17k) force a non-blank to the left and right of any other non-blank in the q-layer, except when a non-blank coincides with an $\left|\square\right\rangle$ or $\left|\square\right\rangle$ in the c-layer.

One can easily see[1] that there is a basis of eigenstates of H of the form $|T\rangle_c\,|\psi\rangle_q$, where $|T\rangle_c$ is a product state in the canonical basis of the c-layer.

For a given classical tile configuration $|T\rangle_c$ on the c-layer, let $\mathcal{L}$ denote the set of all horizontal line segments ℓ that lie between an $\left|\square\right\rangle$ and an $\left|\square\right\rangle$ (inclusive) in the classical configuration $|T\rangle_c$ (without any other $\square$ or $\square$ in between them). Let $\mathcal{L}_L$ denote the set of all horizontal line segments between an $\square$ and the right boundary of the region, and similarly $\mathcal{L}_R$ the horizontal line segments between the left boundary and an $\square$ (in both cases, also without any other $\square$ or $\square$ in between).

The associated energy $\langle T|_c\,\langle\psi|_q\,H\,|T\rangle_c\,|\psi\rangle_q$ can be see to be $\geq \langle T|H_c|T\rangle + \sum_{\ell\in\mathcal{L}} \lambda_0(|\ell|)$[1]. Moreover, for a configuration T given by a valid tiling, the associated

energy is indeed

$$\sum_{\ell\in\mathcal{L}} \lambda_0(|\ell|). \tag{18}$$

This is attained by choosing the state $|\psi_0\rangle_q$ consisting of the ground state of $H_q(\ell)$ in the q-layer for each $\ell \in \mathcal{L}$, a 0-energy eigenstate of $H_q(\ell)$ in the q-layer for each $\ell \in \mathcal{L}_L \cup \mathcal{L}_R$, and $|0\rangle$ everywhere else in the q-layer. In this case the set $\mathcal{L}$ is given by exactly the *segments*, that is, top borders of the squares appearing in the pattern of Figure 5. By Theorem 5, (18) is minimised among all valid tilings by the quantity

$$E(0 \text{ defects}) = \sum_{n=1}^{\lfloor \log_4(L/2)\rfloor} \left(\left\lfloor \frac{H}{2^{2n+1}}\right\rfloor \left(\left\lfloor \frac{L}{2^{2n+1}}\right\rfloor - 1\right)\right) \lambda_0(4^n). \tag{19}$$

On the other hand, since each defect in the classical tile configuration contributes energy at least 1 from the h_c term, Theorem 5 implies that the energy of an eigenstate with d defects on the $L \times H$ rectangle is at least

$$E(d \text{ defects}) \geq d + \sum_{\ell\in\mathcal{L}} \lambda_0(|\ell|) \geq d + \sum_{n=1}^{\lfloor \log_4(L/2)\rfloor} \left(\left\lfloor \frac{H}{2^{2n+1}}\right\rfloor \left(\left\lfloor \frac{L}{2^{2n+1}}\right\rfloor - 1\right) - 2d\right) \lambda_0(4^n). \tag{20}$$

Since $\sum_r \lambda_0(r) < 1/2$ by assumption for a Gottesman-Irani Hamiltonian (see Definition 6), we have for all $d > 0$

$$E(d \text{ defects}) - E(0 \text{ defects}) \geq d\left(1 - 2\sum_{n=1}^{\lfloor \log_4(L/2)\rfloor} \lambda_0(4^n)\right) > 0. \tag{21}$$

The Lemma follows. □

We can now apply this Lemma to construct a Hamiltonian h_u with ground state energy that is undecidable even with a constant promise on the energy gap.

Proposition 8 (Diverging g.s. energy)
We can construct a family of interactions $h_u^{row}(n), h_u^{col}(n) \in \mathcal{B}(\mathbb{C}^U \otimes \mathbb{C}^U)$ and $h_u^{(1)}(n) \in \mathcal{B}(\mathbb{C}^U)$ with operator norm $\leq \beta$ and algebraic matrix entries, and strictly positive functions $\alpha_1^l(n), \alpha_0(n), \delta_2(n), \alpha_1^u(n), \delta_1(n)$ (where the α functions are computable and the δ functions are uncomputable), such that either $\lambda_0(H_u^{\Lambda(L)}(n)) = -L\alpha_1^l(n) + \alpha_0(n)$, or $\lambda_0(H_u^{\Lambda}(n)) = L^2\delta_2(n) - L[\alpha_1^u(n) + \delta_1(n)]$, but determining which is undecidable.

Moreover, the interactions can be taken to have the following form: $h_u^{(1)}(n) = \alpha_2(n)\mathbb{1}$ with $\alpha_2(n)$ an algebraic number $\leq \beta$ computable from n, $h_u^{row}(n)$ $\{0,\beta\}$-valued and independent of n and

$$h_u^{col}(n) = \beta\left(A + e^{i\pi\varphi}B + e^{i\pi 2^{-|\varphi|}}C\right) + \text{h.c.} \tag{22}$$

where $A \in \mathcal{B}(\mathbb{C}^U \otimes \mathbb{C}^U)$ is independent of n and has coefficients in $\mathbb{Z} + \frac{1}{\sqrt{2}}\mathbb{Z}$, $B, C \in \mathcal{B}(\mathbb{C}^U \otimes \mathbb{C}^U)$ are independent of n and have coefficients in $\mathbb{Z}$, and $\beta \in \mathbb{Q}$ is independent of n and can be taken as small as desired. Recall that φ is defined as the rational number whose binary fraction expansion contains the digits of n after the decimal point.

Proof Let h_{q0} be the Hamiltonian obtained by applying Theorem 4 with $K = 3$ to the QTM from Theorem 3 with a properly behaved reversible universal TM dovetailed after it. The Hamiltonian $h_q(n)$ in Lemma 7 will then be $h_q(n) = h_{q0}(n) + |\top\rangle\langle\top| \otimes \mathbb{1} + \mathbb{1} \otimes |\top\rangle\langle\top|$, where $|\top\rangle$ is the halting state of the universal TM. h_q clearly has the form given in part (vi) of Theorem 4. Moreover, this Hamiltonian is a Gottesman-Irani Hamiltonian according to Definition 6. The key observation for this (see Ref. 1) is the following estimate.

Let $|\psi\rangle = \frac{1}{\sqrt{T}} \sum_{t=1}^{T} |\phi_t\rangle |\psi_t\rangle$ be the normalised computational history state for the QTM, where $T = \Omega(|\Sigma \times Q|^r)$ and $|\psi_t\rangle$ is the state encoding the t^{th} step of the computation. Note that $|\psi\rangle$ is a zero-energy eigenstate of H_{q0}, and at most one $|\psi_t\rangle$ can have support on the state $|\top\rangle$ that represents the halting state of the universal TM, by Theorem 4. For $r > 2$, we have

$$\begin{aligned}\lambda_0(r) \leq \langle\psi|H_q(r)|\psi\rangle &= \langle\psi|\left(\sum_i h_{q0}^{(i,i+1)}(n) + |\top\rangle\langle\top|_i \otimes \mathbb{1}_{i+1} + \mathbb{1}_i \otimes |\top\rangle\langle\top|_{i+1}\right)|\psi\rangle \\ &= \sum_{t=1}^{T} \frac{1}{T}\langle\psi_t|\left(\sum_i |\top\rangle\langle\top|_i \otimes \mathbb{1}_{i+1} + \mathbb{1}_i \otimes |\top\rangle\langle\top|_{i+1}\right)|\psi_t\rangle \leq O\left(\frac{1}{|\Sigma \times Q|^r}\right).\end{aligned} \tag{23}$$

Let $\tilde{h}_u^{\text{row}}(n), \tilde{h}_u^{\text{column}}(n)$ be the local interactions obtained by applying Lemma 7 to $h_q(n)$. Let $N(n) := \max\{\|h^{row}(n)\|, \|h^{col}(n)\|\}$, and take a rational number $\beta \leq \frac{1}{N(n)}$ for all n. Such β exists by the form of h_q guaranteed by part (vi) in Theorem 4 and the definition of $\tilde{h}_u^{\text{row}}(n), \tilde{h}_u^{\text{column}}(n)$ based on h_q. Define the normalised local interactions $h_u^{\text{row}}(n) := \beta\tilde{h}_u^{\text{row}}(n)$, $h_u^{\text{column}}(n) := \beta\tilde{h}_u^{\text{column}}(n)$.

For any $r \geq |n| + 6$, the QTM from Theorem 3 has sufficient tape and time to finish, and we can be sure that the reversible universal TM starts. (Here, $|n|$ once again denotes the number of digits in the binary expansion of n.) If the universal TM does *not* halt on input n, then for all $r \geq |n| + 6$ we have that $\lambda_0(r) = 0$. By Theorem 5, the minimum number of r-segments in any tiling of an $L \times L$ square (for $r = 4^m$, $m \in \mathbb{N}$) is $\lfloor L/2r \rfloor (\lfloor L/2r \rfloor - 1)$, and this minimum can be attained for all r simultaneously. Hence, as long as we take $L \geq L_0(n)$ where $L_0(n)$ is the minimal L such that the modified Robinson tiling of $\Lambda(L)$ necessarily contains a 4^m-segment

of size $4^m \geq |n| + 6$, then Lemma 7 gives a ground state energy for $H^{\Lambda(L)}(n)$ of

$$\lambda_0(H^{\Lambda(L)}) = \beta \min_{\mathcal{L}\subset\Lambda(L)} \sum_{\ell\in\mathcal{L}} \lambda_0(|\ell|) \quad = \beta \sum_{\substack{1\leq r\leq |n|+6 \\ r=4^m,\, m\in\mathbb{N}}} \left\lfloor \frac{L}{2r} \right\rfloor \left(\left\lfloor \frac{L}{2r} \right\rfloor - 1 \right) \lambda_0(r) \tag{24a}$$

$$= L^2 \left[\beta \sum_{\substack{1\leq r\leq |n|+6 \\ r=4^m,\, m\in\mathbb{N}}} \frac{\lambda_0(r)}{4r^2} \right] \; - \; L \left[\beta \sum_{\substack{1\leq r\leq |n|+6 \\ r=4^m,\, m\in\mathbb{N}}} \frac{\lambda_0(r)}{2r} \left(2\,\mathrm{frac}\left(\frac{L}{2r}\right) + 1 \right) \right]$$
$$+ \left[\beta \sum_{\substack{1\leq r\leq |n|+6 \\ r=4^m,\, m\in\mathbb{N}}} \lambda_0(r)\,\mathrm{frac}\left(\frac{L}{2r}\right) \left(\mathrm{frac}\left(\frac{L}{2r}\right) + 1 \right) \right] \tag{24b}$$

$$=: L^2\,\alpha_2(n) - L\,\alpha_1(n, L) + \alpha_0(n, L)\,, \tag{24c}$$

where $\mathrm{frac}(x) := x - \lfloor x \rfloor$ denotes the fractional part of x. Note that the number of terms in the sums are finite, and for all finite r the quantity $\lambda(r)$ is an eigenvalue of a finite-dimensional matrix. Therefore, $\alpha_2(n)$ is always a algebraic computable number. We also have

$$\alpha_1^l(n) := \beta \sum_{\substack{1\leq r\leq |n|+6 \\ r=4^m,\, m\in\mathbb{N}}} \frac{\lambda_0(r)}{2r} \leq \alpha_1(n, L) < \beta \sum_{\substack{1\leq r\leq |n|+6 \\ r=4^m,\, m\in\mathbb{N}}} \frac{3\lambda_0(r)}{2r} =: \alpha_1^u(n), \tag{25}$$

$$0 \leq \alpha_0(n, L) \leq \beta \sum_{\substack{1\leq r\leq |n|+6 \\ r=4^m,\, m\in\mathbb{N}}} \frac{\lambda_0(r)}{r} =: \alpha_0(n). \tag{26}$$

If the universal TM *does* halt on input n, then for any r larger than the size of tape needed to halt it is clear that $\lambda_0(r) > 0$. This follows immediately from the fact that the computational history state encoding the evolution (which necessarily has support on $|\top\rangle$) is the unique ground state of h_{q0}, and $h_{q0} \geq 0$. Let $r_1(n)$ be the minimal such r of the form 4^m. Then by Lemma 7, the ground state energy of $H^{\Lambda(L)}$

is

$$\lambda_0(H^{\Lambda(L)}) = \beta \min_{\mathcal{L}\subset\Lambda(L)} \sum_{\ell\in\mathcal{L}} \lambda_0(|\ell|) \tag{27a}$$

$$= \beta \sum_{\substack{1\le r\le |n|+6 \\ r=4^m,\, m\in\mathbb{N}}} \left\lfloor \frac{L}{2r} \right\rfloor \left(\left\lfloor \frac{L}{2r} \right\rfloor - 1 \right) \lambda_0(r) \quad + \beta \sum_{\substack{r\ge r_1(n) \\ r=4^m,\, m\in\mathbb{N}}} \left\lfloor \frac{L}{2r} \right\rfloor \left(\left\lfloor \frac{L}{2r} \right\rfloor - 1 \right) \lambda_0(r) \tag{27b}$$

$$\begin{aligned} &= L^2 \left(\alpha_2(n) + \beta \sum_{\substack{r\ge r_1(n) \\ r=4^m,\, m\in\mathbb{N}}} \frac{\lambda_0(r)}{4r^2} \right) \\ &\quad - L \left(\alpha_1(n) + \beta \sum_{\substack{r\ge r_1(n) \\ r=4^m,\, m\in\mathbb{N}}} \frac{\lambda_0(r)}{2r} \left(2 \operatorname{frac}\left(\frac{L}{2r} \right) + 1 \right) \right) \\ &\quad + \left(\alpha_0(n) + \beta \sum_{\substack{r\ge r_1(n) \\ r=4^m,\, m\in\mathbb{N}}} \lambda_0(r) \operatorname{frac}\left(\frac{L}{2r} \right) \left(\operatorname{frac}\left(\frac{L}{2r} \right) + 1 \right) \right) \end{aligned} \tag{27c}$$

$$=: L^2[\alpha_2(n) + \delta_2(n)] - L[\alpha_1(n, L) + \delta_1(n, L)] + \alpha_0(n, L) + \delta_0(n, L). \tag{27d}$$

Note that $\delta_2(n) > 0$, since in the halting case $\lambda_0(r) > 0$ for all $r \ge r_1(n)$, and

$$\beta \sum_{\substack{r\ge r_1(n) \\ r=4^m,\, m\in\mathbb{N}}} \frac{\lambda_0(r)}{2r} \le \delta_1(n, L) < \beta \sum_{\substack{r\ge r_1(n) \\ r=4^m,\, m\in\mathbb{N}}} \frac{3\lambda_0(r)}{2r} =: \delta_1(n), \tag{28}$$

$$0 \le \delta_0(n, L) \le \beta \sum_{\substack{r\ge r_1(n) \\ r=4^m,\, m\in\mathbb{N}}} \frac{\lambda_0(r)}{r} := \delta_0(n). \tag{29}$$

We now modify $h(n)$ by adding the 1-body term $h_u^{(1)} = -\alpha_2(n)\mathbb{1}$ acting at each site. The ground state energy of $H^{\Lambda(L)}$ is simply shifted by exactly $-L^2\alpha_2(n)$, so in the non-halting case we have

$$\lambda_0(H^{\Lambda(L)}) = -L\alpha_1(n, L) + \alpha_0(n, L) \le -L\alpha_1^l(n) + \alpha_0(n). \tag{30}$$

In the halting case, we have

$$\begin{aligned} \lambda_0(H^{\Lambda(L)}) &= L^2\delta_2(n) - L[\alpha_1(n, L) + \delta_1(n, L)] + \alpha_0(n, L) + \delta_0(n, L) \\ &\ge L^2\delta_2(n) - L[\alpha_1^u(n) + \delta_1(n)]. \end{aligned} \tag{31}$$

The Proposition follows from undecidability of the Halting Problem. □

Undecidability of the ground state energy density (Theorem 2) is now immediate from Proposition 8 by the definition $E_\rho := \lim_{L\to\infty} \lambda_0(H^{\Lambda(L)})/L^2$ of the ground state energy density. Undecidability of the spectral gap then follows from the construction discussed in the main text.

References

1. Toby S. Cubitt, David Perez-Garcia, and Michale M. Wolf. Undecidability of the spectral gap. 2015.

2. Daniel Gottesman and Sandy Irani. 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*, pages 95–104. IEEE, 2009. arXiv:0905.2419[quant-ph].

3. R. Berger. The undecidability of the domino problem. *Mem. Amer. Math. Soc.*, 66:1–72, 1966.

4. Richard Feynman. Quantum mechanical computers. *Optics News*, 11:11–20, 1985.

5. A.Y. Kitaev, A.H. Shen, and M.N. Vyalyi. *Classical and Quantum Computation*, volume 47. AMS, 2002.

6. J. Kempe, A. Kitaev, and O. Regev. The complexity of the local hamiltonian problem. *SIAM Journal of Computing*, 35:1070–1097, 2006.

7. R. Oliveira and B. Terhal. The complexity of quantum spin systems on a two-dimensional square lattice. *QIC*, 8:0900–0924, 2008. arXiv:0504050[quant-ph].

8. Dorit Aharonov, Daniel Gottesman, Sandy Irani, and Julia Kempe. The power of quantum systems on a line. *Commun. Math. Phys.*, 287:41–65, 2009. arXiv:0705.4077v2 [quant-ph].

9. Ethan Bernstein and Umesh Vazirani. Quantum complexity theory. *SIAM Journal on Computing*, 26(5):1411–1473, 1997. www.cs.berkeley.edu/~vazirani/pubs/bv.ps.

10. Tibor Rado. On non-computable functions. *Bell System Technical Journal*, 41: 877–884, 1962.

11. Alexander S. Holevo. Bound for the quantity of information transmitted by a quantum communication channel. *Problems of Information Transmission*, 9: 177–183, 1973.

12. Michael A. Nielsen and Isaac L. Chuang. *Quantum computation and quantum information*. Cambridge University Press, 2000.

13. R. Robinson. Undecidability and nonperiodicity for tilings of the plane. *Invent. Math.*, 12:177–209, 1971.