

Undecidability of the Spectral Gap

Toby Cubitt^{*1}, David Perez-Garcia^{†2}, and Michael M. Wolf^{‡3}

¹*DAMTP, University of Cambridge, Centre for Mathematical Sciences,
Wilberforce Road, Cambridge CB3 0WA, United Kingdom*

²*Departamento de Análisis Matemático, Facultad de CC Matemáticas,
Universidad Complutense de Madrid, Plaza de Ciencias 3, 28040 Madrid, Spain*

³*Department of Mathematics, Technische Universität München, 85748 Garching, Germany*

February 20, 2015

Abstract

We show that the spectral gap problem is undecidable. Specifically, we construct families of translationally-invariant, nearest-neighbour Hamiltonians on a 2D square lattice of d -level quantum systems (d constant), for which determining whether the system is gapped or gapless is an undecidable problem. This is true even with the promise that each Hamiltonian is either gapped or gapless in the strongest sense: it is promised to either have continuous spectrum above the ground state in the thermodynamic limit, or its spectral gap is lower-bounded by a constant in the thermodynamic limit. Moreover, this constant can be taken equal to the local interaction strength of the Hamiltonian.

This implies that it is logically impossible to say in general whether a quantum many-body model is gapped or gapless. Our results imply that for any consistent, recursive axiomatisation of mathematics, there exist specific Hamiltonians for which the presence or absence of a spectral gap is independent of the axioms.

These results have a number of important implications for condensed matter and many-body quantum theory.

^{*}tsc25@cam.ac.uk

[†]dperezga@ucm.es

[‡]m.wolf@tum.de

Contents

I	7
1 Introduction	7
2 Preliminaries	9
2.1 Gapped versus gapless Hamiltonians	9
2.2 Undecidability	10
3 Main results	12
3.1 Implications of the results	13
4 Extended Overview	14
4.1 Ground state energy density	14
4.2 Wang Tilings	17
4.3 QMA constructions	19
4.4 Constant local dimension	21
4.5 Translational invariance	23
4.6 The thermodynamic limit	25
4.7 Structure of the paper	27
II	29
5 Unconstrained local Hilbert space dimension	29
5.1 Undecidability of the spectral gap via tiling	29
5.1.1 Ingredient 1: a tiling Hamiltonian	29
5.1.2 Ingredient 2: a gapless frustration-free Hamiltonian	30
5.1.3 Reducing tiling to spectral gap	30
5.2 Undecidability of low energy properties	33
5.2.1 Reducing tile completion to a ground state energy problem	33
5.2.2 Reduction of the halting problem to arbitrary low energy properties	36
III	41
6 Quantum Phase Estimation Turing Machine	41
6.1 Quantum Turing Machinery	42
6.1.1 Turing Machine Programming	47

6.1.2	Reversible Turing Machine Toolbox	50
6.2	Quantum phase estimation overview	57
6.3	Preparation stage	58
6.4	Control- U_φ stage	58
6.4.1	cU^k machine	59
6.4.2	Iterating over the control qubits	62
6.5	Locating the LSB	64
6.6	QFT stage	68
6.7	Reset Stage	70
6.8	Analysis	71
7	Encoding QTMs in local Hamiltonians	73
7.1	Preliminaries	74
7.2	Clock Oscillator	77
7.3	Initialisation sweep	79
7.4	Clock Counter	80
7.4.1	Counter TM construction	80
7.4.2	Counter TM Hamiltonian	89
7.5	Clock Hamiltonian	95
7.6	QTM Hamiltonian	98
7.6.1	QTM transition rules	99
7.6.2	QTM Unitarity	100
7.6.3	QTM initialisation sweep	107
7.7	Analysis	107
8	Quasi-periodic tilings	115
8.1	Robinson's tiling	115
8.2	Rigidity of the Robinson tiling	118
9	Putting it all together	124
9.1	Undecidability of the g.s. energy density	124
9.2	From g.s. energy density to spectral gap	133
9.3	Periodic boundary conditions	139
10	Acknowledgements	143

List of Theorems

Definition 1	Gapped	10
Definition 2	Gapless	10
Theorem 3	Algorithmic undecidability of the spectral gap	12
Theorem 4	Axiomatic independence of the spectral gap	12
Proposition 5	Decidability of g.s. energy density with promise gap	16
Theorem 6	Undecidability of g.s. energy density	16
Theorem 7	Reducing tiling to spectral gap	30
Corollary 8	Undecidability of spectral gap for unconstrained dimension	32
Lemma 9	Relating ground state energy to the halting problem	35
Theorem 10	Relating low energy properties to the halting problem	36
Theorem 11	Phase-estimation QTM	41
Definition 12	Turing Machine	42
Definition 13	Quantum Turing Machine	43
Definition 14	Generalised TM and QTM	43
Definition 15	Well-formed QTM	43
Definition 16	Normal form QTM	44
Definition 17	Unidirectional QTM	44
Theorem 18	Local well-formedness	44
Theorem 19	Quantum local well-formedness	45
Definition 20	Proper QTM	46
Lemma 21	Subroutine Lemma	47
Lemma 22	Dovetailing Lemma	48
Lemma 23	Reversal Lemma	48
Lemma 24	Copying machine	50
Lemma 25	Shift-right machine	50
Lemma 26	Equality machine	51
Lemma 27	Increment and decrement machines	53
Lemma 28	Looping Lemma	54
Lemma 29	Binary adder	56
Lemma 30	Unary to binary converter	56
Lemma 31	Controlled- U QTM	59
Definition 32	Computational history state	73
Theorem 33	Local Hamiltonian QTM encoding	73
Definition 34	Bracketed state	75
Definition 35	Legal and illegal states	76
Lemma 36	Regexps	77
Lemma 37	Evolve-to-illegal	83
Definition 38	Well-formed state	95
Lemma 39	Well-formed transitions	95

Lemma 40	Evolve-to-illegal	97
Definition 41	Standard-form Hamiltonian	98
Lemma 42	Invariant subspaces	108
Lemma 43	Clairvoyance Lemma	108
Lemma 44	Geometrical Lemma	109
Proposition 45	Unique ground state	111
Lemma 46	Cross rigidity	119
Lemma 47	Robinson rigidity	119
Lemma 48	Segment bound	122
Lemma 49	Segment rigidity	123
Definition 50	Gottesman-Irani Hamiltonian	124
Lemma 51	Tiling + quantum layers	124
Lemma 52	Robinson + Gottesman-Irani Hamiltonian	128
Proposition 53	Diverging g.s. energy	129
Corollary 54	Undecidability of g.s. energy with promise	133
Theorem 6	Undecidability of g.s. energy density – restated	133
Theorem 3	Undecidability of the spectral gap (restated)	133
Definition 55	Gapless – generalised	140
Lemma 56	Periodic to open b.c.	141

List of Figures

Figure 1	Gapped and gapless	11
Figure 2	From ground state energy density to spectral gap.	15
Figure 3	Valid and invalid tilings.	17
Figure 4	Undecidability of the tiling completion problem.	34
Figure 5	Counting argument for proof of Theorem 10.	39
Figure 6	Quantum phase estimation, control-phase stage.	59
Figure 7	Quantum phase estimation, inverse QFT stage.	68
Figure 8	Evolution of the Track 1 clock oscillator.	78
Figure 9	The five basic tiles of Robinson’s tiling.	115
Figure 10	Parity markings to force the position of the crosses.	115
Figure 11	A possible Robinson tiling of the plane.	116
Figure 12	Robinson tiling squares.	117
Figure 13	The five basic tiles of the modified Robinson tiling.	118
Figure 14	Gottesman-Irani / Robinson “sandwich”.	125
Figure 15	Boundary tiles for periodic boundary condition construction.	141

Part I

1 Introduction

The spectral gap is one of *the* most important properties of a quantum many-body system. It plays a pivotal role in condensed matter, mathematical, and fundamental physics, and also in quantum computing.

One of the main goals of condensed matter theory is to understand phase transitions and phase diagrams. The behaviour of the spectral gap is intimately related to the phase diagram of a quantum many-body system, with quantum phase transitions occurring at critical points where the gap vanishes. The low-temperature physics of the system are governed by the spectral gap: gapped systems exhibit “non-critical” behaviour, with low-energy excitations that behave as massive particles, preventing long-range correlations [MH06]; gapless systems exhibit “critical” behaviour, with low-energy excitations that behave as massless particles, allowing long-range correlations.

Many seminal results in mathematical physics prove that specific systems are gapped or gapless. Examples include the Lieb-Schultz-Mattis result showing that the Heisenberg chain for half-integer spins is gapless [LSM61] (extended to higher dimensions by Hastings [Has04]), or the proof of a spectral gap for the 1D AKLT model [Aff+88]. Other results only apply to gapped systems. Important examples include exponential decay of correlations [MH06], the 1D entanglement area-law [Has07], or the very recent proof that the ground state energy density problem in 1D is in P [LVV13].

The spectral gap also arises as a key quantity in quantum computing, via adiabatic quantum computation [Far+00]. In this type of computation, proven to be exactly as powerful as the usual circuit model for quantum computation [Aha+07], the system is initialised in the ground state of an easy Hamiltonian, that is a Hamiltonian for which the ground state can be easily constructed. Then the interactions are slowly changed, to end up in a target Hamiltonian whose ground state encodes the output of the computation. The key property controlling the efficiency of this method is the minimal spectral gap in the path of Hamiltonians connecting the initial Hamiltonian to the target. A computational problem has an efficient quantum algorithm iff there exists such a Hamiltonian path for which the minimal spectral gap is lower-bounded by an inverse-polynomial in the system size.

Thus understanding whether a given system is gapped or not is one of the fundamental questions in quantum many-body physics. Indeed, there are many famous open problems concerning the spectral gap. A paradigmatic example is the antiferromagnetic Heisenberg model. Even in 1D, the case of integer spin remains open as the “Haldane conjecture” [Hal83], first formulated in 1983. The same ques-

tion in the case of 2D non-bipartite lattices, such as the Kagome lattice, was posed by Anderson in 1973 [And73]. The spectral gap of such systems remains one of the main unsolved questions about the long-sought topological spin liquid phase. This latter problem has gained a lot of attention in recent years [Bal10] due to materials, such as herbertsmithite [Han+12], whose interactions are well-approximated by the Heisenberg coupling, and which—according to the latest numerical evidence [YHW11]—are potential candidates for topological spin liquids. In the related setting of quantum field theory, determining if Yang-Mills theory is gapped is one of the Millennium Prize Problems [JW00], and is closely related to one of the most important open problems in high-energy physics: explaining the phenomenon of quark confinement.

All of these problems are specific cases of the general spectral gap problem: given a quantum many-body Hamiltonian, is the system it describes gapped or gapless? Our main result is a proof that the spectral gap problem is undecidable. This is much stronger than merely showing that a problem is computationally hard. It implies that it is not merely difficult to determine whether a systems is gapped or gapless; it is *logically impossible to say* in general.

This has two subtly distinct meanings. First, the spectral gap problem is undecidable in precisely the same sense as the Halting Problem is undecidable: there cannot exist an algorithm, no matter how inefficient, to determine if a general Hamiltonian is gapped or gapless. This algorithmic undecidability relates to solving the spectral gap for a family of Hamiltonians, as we might have if we want to map out the phase diagram of a system as we vary some external parameters. Second, the spectral gap problem is in general an undecidable question: there exist simple Hamiltonians for which neither the presence nor the absence of a spectral gap is provable from the axioms of mathematics. Theorems 3 and 4 in Section 3 make both of these meanings precise.

The idea that some of the most difficult open problems in physics could be mathematically proven to be “impossible to solve” is not new. Undecidability of other physical quantities has been shown in many-body systems for the much easier cases where either the many-body lattice structure or the translational-invariance is removed [Gu+09; Llo93; PER89]. Indeed, proving such impossibility theorems is highlighted as one of the main open problems in mathematical physics in the list published by the International Association of Mathematical Physics in the late 90’s, edited by Aizenman [Aiz98]. Our result, then, can be seen as a major contribution to this.

2 Preliminaries

2.1 Gapped versus gapless Hamiltonians

The following sets the framework for our analysis and defines the objects and terms used. 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 . Two cases will be considered: *periodic boundary conditions* where the outer rows and columns are connected along the same direction as well, so that $\Lambda(L)$ becomes a square lattice on a torus, and *open boundary conditions* where these connections are not made.

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 possibly on-site terms in a translational 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

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

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

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

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

Let $\text{spec } H^{\Lambda(L)} := \{\lambda_0, \lambda_1, \dots\}$ 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 *ground state energy*, the corresponding eigenvector *ground state* and a Hamiltonian $H^{\Lambda(L)}$ *frustration-free* if its ground state energy is zero while 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.

The quantity we are interested in is the asymptotic behaviour of the *spectral gap*

$$\Delta(H^{\Lambda(L)}) := \lambda_1(H^{\Lambda(L)}) - \lambda_0(H^{\Lambda(L)}) \quad (3)$$

in the thermodynamic limit, that is, when $L \rightarrow \infty$. The main result of this paper is, loosely speaking, that it is in general undecidable whether or not this gap closes in the thermodynamic limit. To be more precise we need the following definitions:

Definition 1 (Gapped) *We say that a family $\{H^{\wedge(L)}\}$ of Hamiltonians, as described above, characterises a gapped system if there is a constant $\gamma > 0$ and a system size L_0 such that for all $L > L_0$, $\lambda_0(H^{\wedge(L)})$ is non-degenerate and $\Delta(H^{\wedge(L)}) \geq \gamma$. In this case, we say that the spectral gap is at least γ .*

Definition 2 (Gapless) *We say that a family $\{H^{\wedge(L)}\}$ of Hamiltonians, as described above, characterises a gapless system if there is a constant $c > 0$ such that for all $\varepsilon > 0$ there is an $L_0 \in \mathbb{N}$ so that for all $L > L_0$ any point in $[\lambda_0(H^{\wedge(L)}), \lambda_0(H^{\wedge(L)}) + c]$ is within distance ε from $\text{spec } H^{\wedge(L)}$.*

Note that gapped is not defined as the negation of gapless; there are systems that fall into neither class. The reason for choosing such strong definitions is to deliberately avoid ambiguous cases (such as systems with degenerate ground states). A Hamiltonian that is gapped or gapless according to these definitions should be recognised as such throughout the literature. Our constructions will allow us to use these strong definitions, because we are able to guarantee that each instance falls into one of the two classes. Indeed, we could further strengthen the definition of “gapless” without changing our undecidability results or their proofs, below, by demanding that $c = c(L)$ grows with L so that $\lim_{L \rightarrow \infty} c(L) = \infty$.

2.2 Undecidability

There are two different albeit related notions of undecidability, which we will briefly and informally recall in the following: one, which can be called *axiomatic undecidability*, is more related to mathematical logic whereas the other, which might be called *algorithmic undecidability*, feels more at home in theoretical computer science. For a rigorous introduction and a more detailed discussion of the notions used below we refer to [Bar77].

A decision problem (i.e. a YES/NO assignment to a each member of a set of inputs) is *algorithmically undecidable* if there is no algorithm that terminates and provides the correct answer upon every input. The notion of an ‘algorithm’ used here is formally defined in terms of Turing machines, but the Church-Turing thesis suggests that any other reasonable formalisation will do as well. A simple but crucial point to note here is that decision problems with only finitely many possible inputs are trivially decidable by case distinction. In particular, single-instance problems like “is the spin-1 anti-ferromagnetic Heisenberg model gapped?” are trivially decidable in the algorithmic sense since either the algorithm “PRINT:’YES’” or the algorithm “PRINT:’NO’” will provide the correct assignment.

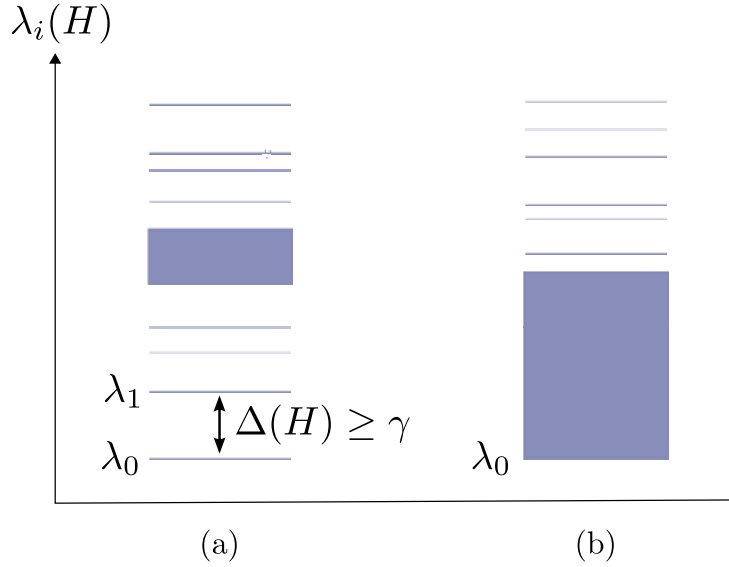


Figure 1: (a) A gapped system has a unique ground state $\lambda_0(H)$ and a constant lower-bound γ on the spectral gap $\Delta(H)$ in the thermodynamic limit. (b) A gapless system has continuous spectrum above the ground state in the thermodynamic limit. See Definitions 1 and 2 for the precise definitions.

This differs from the second notion of axiomatic undecidability, also called *axiomatic independence*. Within a formal system, which comprises a set of axioms together with rules defining formal proofs, a statement is undecidable in the axiomatic sense (or, synonymously, independent) if it can neither be proven nor disproven from the axioms.

There is a simple, but for our purpose important, relation between these two notions of undecidability: if a decision problem is algorithmically undecidable, then for every consistent and recursive formal system in which the problem can be stated, there are infinitely many instances of the problem each of which can neither be proven nor disproven from the axioms. The argument is standard [Poo12], so we only sketch it here informally. Suppose for contradiction that all but finitely many instances are axiomatically decidable. Since the theorems in our formal system are recursively enumerable, this would imply that there exists an algorithm which runs through all theorems, until either a proof or a disproof is found. The exceptions, since there are finitely many, could be dealt with by case distinction.

In the main part of this paper we prove undecidability of problems concerning the spectral gap, the ground state energy density and other low temperature properties of quantum spin systems in the algorithmic sense. Axiomatic independence then follows as a corollary in the way just stated.

3 Main results

For each natural number n , we define $\varphi = \varphi(n)$ as the rational number whose binary decimal expansion contains the digits of n in reverse order after the decimal. We also fix one particular Universal Turing Machine and call it UTM.

Theorem 3 (Algorithmic undecidability of the spectral gap) *We construct explicitly a dimension d , $d^2 \times d^2$ matrices A, B, C, D and a rational number β so that*

- (i). A is Hermitian and with coefficients in $\mathbb{Z} + \beta\mathbb{Z} + \frac{\beta}{\sqrt{2}}\mathbb{Z}$,
- (ii). B, C have integer coefficients,
- (iii). D is Hermitian and with coefficients in $\{0, 1, \beta\}$.

For each natural number n , define:

$$\begin{aligned} h_1(n) &= \alpha(n)\mathbb{1} \quad (\alpha(n) \text{ an algebraic number}) \\ h_{\text{row}}(n) &= D \quad (\text{independent of } n) \\ h_{\text{col}}(n) &= A + \beta \left(e^{i\pi\varphi} B + e^{-i\pi\varphi} B^\dagger + e^{i\pi 2^{-|\varphi|}} C + e^{-i\pi 2^{-|\varphi|}} C^\dagger \right). \end{aligned}$$

Then:

- (i). The local interaction strength is ≤ 1 . I.e. all terms $h_1(n), h_{\text{row}}(n), h_{\text{col}}(n)$ have operator norm bounded by 1.
- (ii). If UTM halts on input n , then the associated family of Hamiltonians $\{H^{\Lambda(L)}(n)\}$ is gapped in the strong sense of Definition 1 and, moreover, the gap $\gamma \geq 1$.
- (iii). If UTM does not halt on input n , then the associated family of Hamiltonians $\{H^{\Lambda(L)}(n)\}$ is gapless in the strong sense of Definition 2.

Using the classic result of Turing [Tur36] that the Halting Problem for UTM on input n is algorithmically undecidable, we can conclude that *the spectral gap problem is also undecidable*. As noted above, this immediately gives the following:

Theorem 4 (Axiomatic independence of the spectral gap) *Let $d \in \mathbb{N}$ be a sufficiently large constant. For any consistent formal system with a recursive set of axioms, there exists a translationally-invariant nearest-neighbour Hamiltonian on a 2D lattice with local dimension d and algebraic entries for which neither the presence nor the absence of a spectral gap is provable from the axioms.*

It is straightforward (if tedious) to extract an explicit value for d in Theorems 3 and 4 from the construction described in this paper.

3.1 Implications of the results

These results have a number of implications for condensed matter physics and fundamental physics. They imply that one can write down simple models whose phase diagrams are uncomputably complicated. The standard approach of trying to gain insight into such models by solving numerically for larger and larger lattice sizes is doomed to failure. The system could display all the features of a gapless model, with the gap of the finite system decreasing monotonically with increasing size. Then, at some threshold size, it may suddenly switch to having a large gap. Not only can this threshold be arbitrarily large; the threshold itself is uncomputable. In general, we can never know whether a large many-body system is approaching the asymptotic behaviour of the thermodynamic limit—one more row of atoms may completely change its properties. Our findings also imply that a result showing robustness of the spectral gap under perturbations, as for the case of frustration-free Hamiltonians in [MP13], cannot hold for general gapped systems.

Phase diagrams with infinitely many phases are known in quantum systems in connection with the quantum Hall effect, where fractal diagrams like the Hofstadter butterfly can be obtained [OA01; Hof76]. Since membership in many fractal sets is not decidable (when formulated in the framework of real computation; see [BS93]), it would be interesting to see whether quantum Hall systems could provide a real-world manifestation of our findings.

Conjectures about the spectral gap, such as the Haldane conjecture, the 2D AKLT conjecture, or the Yang-Mills mass gap conjecture, implicitly assume that these questions can be answered one way or the other. Our results prove that the answer to the general spectral gap question is not determined by the axioms of mathematics. Whilst our results are still a long way from proving that any of these specific conjectures are axiomatically undecidable, they at least open the door to the possibility that these – or similar – questions about physical models may be provably unanswerable.

4 Extended Overview

Since the proof of our spectral gap undecidability result Theorem 3 is quite long, in this section we give an overview and discussion of the main ideas in the proof.

The first and 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.) The intuition behind this is illustrated in Figure 2; the proof can be found in Section 9.2.

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

4.1 Ground state energy density

As in Section 2.1, 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

$$H^{\Lambda(L)} := \sum_{(i,j) \in \mathcal{E}} h^{(i,j)} + \sum_{k \in \Lambda(L)} h_1^{(k)}, \quad (4)$$

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

$$E_\rho := \lim_{L \rightarrow \infty} E_\rho(L), \quad \text{where} \quad E_\rho(L) := L^{-\nu} \lambda_0(H^{\Lambda(L)}). \quad (5)$$

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

$$\lambda_0(H^{\Lambda(L)}) \geq n^\nu \lambda_0(H^{\Lambda(L')}). \quad (6)$$

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

$$\lambda_0(H^{\Lambda(L)}) \leq n^\nu \lambda_0(H^{\Lambda(L')}) + 2\nu n^\nu L'^{\nu-1} c. \quad (7)$$

Dividing both inequalities by L^ν we are left with

$$E_\rho(L') \leq E_\rho(L) \leq E_\rho(L') + \frac{2\nu c}{L'}. \quad (8)$$

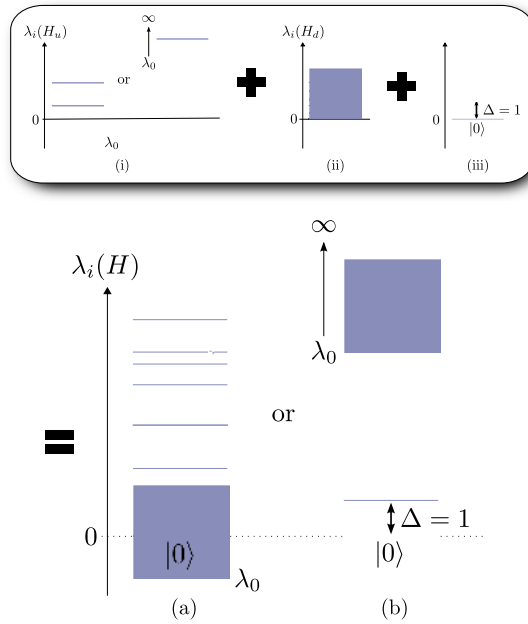


Figure 2: To relate ground state energy density and spectral gap, we need: (i) A Hamiltonian $H_u(\varphi)$ whose ground state energy density is either strictly positive or tends to 0 from below in the thermodynamic limit, but determining which is undecidable, and (ii) a gapless Hamiltonian H_d with ground state energy 0. We combine $H_u(\varphi)$ and H_d to form a new local interaction, $h(\varphi)$, in such a way that $h(\varphi)$ has (iii) an additional non-degenerate 0-energy eigenstate, and the continuous spectrum of H_d is shifted immediately above the ground state energy of H_u . (a) If the ground state energy density of $H_u(\varphi)$ is strictly positive, its ground state energy in the thermodynamic limit must diverge to $+\infty$, and $h(\varphi)$ is gapped. (b) Whereas if the ground state energy density of $H_u(\varphi)$ tends to 0 from below, then its ground state energy in the thermodynamic limit must be ≤ 0 , and $h(\varphi)$ is gapless. A proof can be found in Section 9.2.

Hence, an interval of order $O(1/L')$ contains both \liminf and \limsup of $E_\rho(L)$ so that both must coincide, which proves that $\lim_{L \rightarrow \infty} E_\rho(L)$ 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 precision $\delta > 0$ by exact diagonalisation of $H^{\wedge(L')}$ for any $L' > 2vc/\delta$. This immediately implies that the ground state energy density problem is *decidable* if we provide a finite promise gap δ .

Proposition 5 (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 v -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 as natural number.

This is in sharp contrast to the following, which will form the crucial step in our proof:

Theorem 6 (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 strengths bounded by 1. Then 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 6 that do *not* work. Section 4.2 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 4.3, we briefly review the Kitaev-Feynman-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 4.4, extending ideas from Gottesman and Irani [GI09] in Section 4.5, this approach can give a Hamiltonian with constant local dimension. But it necessarily fails to have the required spectral properties. Finally, in Section 4.6, we discuss how combining ideas from *both* these approaches allows us to achieve the required spectral properties whilst simultaneously keeping the local dimension constant.

4.2 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 [Ber66] classic result that the tiling problem is undecidable.

We will soon see that this approach is too weak to prove Theorem 6. 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 $\mathbb{Z}^2 \rightarrow \mathcal{T}$ 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* [Ber66]. Here, rotations or reflections of the tiles are not allowed – they would make the problem trivial.¹

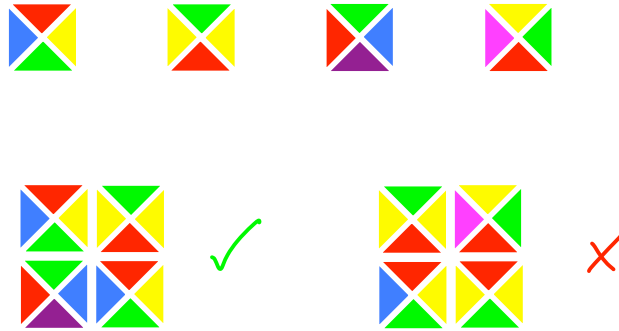


Figure 3: Examples of valid and invalid tilings (bottom) of a set of four Wang tiles (top).

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 $\mathcal{T} = \{1, \dots, T\}$ we assign a Hilbert space $\mathcal{H}^{(i)} \simeq \mathbb{C}^T$ to each site i of a square lattice, and define the local interactions via

$$h^{(i,j)} := \sum_{(m,n) \in C^{(i,j)}} |m\rangle\langle m|_{(i)} \otimes |n\rangle\langle n|_{(j)}, \quad (9)$$

¹If 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.

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 . As is standard in this context, we have used (and will be using throughout) Dirac notation. The overall Hamiltonian on the lattice $\Lambda(L)$ is then

$$H_c^{\Lambda(L)} := \sum_{(i,j) \in \mathcal{E}} h_c^{(i,j)}. \quad (10)$$

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. The full proof is described in Section 5.

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 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. From the discussion in Section 2.2, this immediately implies existence of Hamiltonians for which the spectral gap is axiomatically independent, albeit Hamiltonians whose local dimension is arbitrarily large and unknown (in fact uncomputable). Furthermore, this approach can be extended to prove that *any* low-energy property that distinguishes a Hamiltonian from a gapped system with unique product ground state is undecidable. Section 5 of the paper is devoted to proving these results.

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, the 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 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.

4.3 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 [Fey85] and was significantly developed by Kitaev [KSV02] and others [KKR06; OT08; Aha+09; GI09], is to construct a Hamiltonian whose ground state encodes the history of a quantum computation *in superposition*. If we divide the system into two registers, a “clock” register and a “computational” register, then the desired ground state is

$$\frac{1}{\sqrt{T+1}} \sum_{t=0}^T |t\rangle |\psi_t\rangle, \quad (11)$$

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 [KSV02], who showed how to construct such a Hamiltonian out of 5-body terms. Kempe, Kitaev, and Regev [KKR06] improved this to 2-body, Oliveira and Terhal [OT08] to nearest-neighbour two-qubit interactions on a 2D square lattice, and Aharonov et al. [Aha+09] 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 [GI09] showed how to construct a translationally-invariant Hamiltonian which has as its ground state a computational history state for an arbitrary computation. (In fact, the two-body interaction in the Gottesman and Irani [GI09] construction is the same *fixed* interaction for *all* problem instances. The input is specified by the only remaining free parameter: the length of the chain!)

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 and Irani [GI09] construction, 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 [GI09], scales inverse-polynomially in the system size. Thus (assuming the limits exist)

$$\begin{aligned} \lim_{L \rightarrow \infty} \lambda_0(H^{\Lambda(L)}) &= \begin{cases} \lim_{L \rightarrow \infty} \alpha(L) & \text{non-halting} \\ \lim_{L \rightarrow \infty} \alpha(L) + \frac{1}{\text{poly}(L)} & \text{halting} \end{cases} \\ &= \lim_{L \rightarrow \infty} \alpha(L). \end{aligned} \quad (12)$$

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 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 4.6.

4.4 Constant local dimension

To overcome the unbounded local dimension obstacle we faced in Sections 4.2 and 4.3, 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 4.2, 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.¹ The transition rules of a QTM can have arbitrarily computable numbers as coefficients. (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 [Rad62], or the Holevo bound on the amount of information that can be extracted from a finite-dimensional quantum state [Hol73], 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

¹Indeed, this will essentially be the only point in our construction where we exploit the fact that the Hamiltonian is quantum.

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 [NC00] (running on a QTM) to extract a phase which we encode in a single-qubit unitary, thereby writing out its binary 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 3 becomes vacuous.

It is well-known that the quantum phase estimation circuit can output the unknown phase φ exactly if the phase has a finite binary decimal expansion (i.e. $\varphi = a/2^b$ for some $a, b \in \mathbb{N}$), given the appropriate set of two-qubit gates [NC00]. The universal QTM construction of Bernstein and Vazirani [BV97] shows that any quantum circuit can be implemented on a QTM. But this only implements a quantum circuit up to some error, not exactly. Indeed, it is not at all clear whether one can carry out exact phase estimation for arbitrary φ , even in principle, given that the precision of φ can be arbitrarily high and is unknown to the quantum phase estimation circuit. (Indeed, a no-go result by Nishimura and Ozawa [NO02] shows that no quantum Turing Machine can exactly simulate quantum phase estimation on an arbitrary number of qubits.)

We therefore cannot appeal to previous results to show that such a QTM exists. Instead, in Section 6 we give a detailed and explicit construction of a phase-estimation QTM, that satisfies the very strict requirements of our spectral gap undecidability proof. It may seem as if we have achieved nothing, since now we have the problem of generating the correct input to the phase-estimation QTM!

But in fact this is not an issue because, firstly, the input is encoded in unary, and generating a unary string is much simpler than generating an arbitrary string; and secondly, it does not matter if the upper-bound in the input is larger than the actual number of digits. Our phase-estimation QTM will still write out the phase exactly, and halt deterministically in this case. (The run-time of the QTM will depend on the size of the upper-bound, but run-time is not relevant here.) So we only need to generate a long string consisting entirely of the same symbol! This is very easy, especially when we come to encode this QTM into a local Hamiltonian.

However, it is important to note that, if the input upper-bound is wrong, so that the number of binary digits in the phase exceeds the number specified in the input, then the only guarantee is that the phase-estimation QTM eventually halts; it could leave an arbitrary string (or even quantum state) written to its tape.

In this way, if we supply the phase-estimation QTM with a suitable input and feed its output into a universal TM, the input on which the universal TM runs is entirely determined by the transition amplitudes of the phase-estimation QTM. Assuming the properties of the Gottesman and Irani [GI09] construction carry over, the local Hilbert space dimension will be determined by the alphabet size and number of internal states of the QTMs, which is a constant *independent of the input*. Since the alphabet size and number of internal states of any particular universal TM are clearly constant, the local Hilbert space dimension of the Hamiltonian encoding this sequence of Turing Machines will also be constant.

4.5 Translational invariance

Gottesman and Irani [GI09] 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 and Irani [GI09] construction 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 [GI09] 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 and Irani [GI09] construction 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 and Irani [GI09] construction 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 will require 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 and Irani [GI09] clock construction to count in binary instead of unary, so that the clock can encode at least $\Omega(2^L)$ time-steps. However, this substantially complicates the analysis. We also have to address a subtle but important unitarity issue in order to carry out the required spectral analysis of our Hamiltonian.

We give a complete, self-contained construction in Section 7, heavily inspired by Gottesman and Irani [GI09], together with a rigorous proof of the required spectral properties of the resulting Hamiltonian.²

This construction, together with the constant size phase-estimation quantum Turing Machine from Section 4.4, allows us to construct a family of translationally-invariant 1D local Hamiltonians with *constant* local Hilbert space dimension, whose ground states encode the evolution of a universal Turing Machine running on an input encoded solely into matrix elements of the Hamiltonian. Supplying an appropriate unary input to the phase-estimation QTM (see Section 4.4) can be done simply by adding local terms to the Hamiltonian that force the all the qubits in the initial state of the computational history superposition to be in the $|1\rangle$ state.

Thus we have succeeded in overcoming the constant local dimension obstacle of the QMA constructions discussed above in Section 4.3. 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 in Section 4.3: the energy difference between the halting and non-halting cases decreases polynomially with

¹In their paper, Gottesman and Irani [GI09] 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.

²For this, we prove a general version of the “Clairvoyance Lemma”, which has been used in a number of QMA-hardness proofs in the literature but never formally stated or rigorously proven. This may be of independent interest to Hamiltonian complexity theorists.

the system size. So all dependence of the ground state energy on the outcome of the computation still vanishes in the thermodynamic limit.

4.6 The thermodynamic limit

The more challenging obstacle of the thermodynamic limit still remains.

We first return to tiling problems. However, instead of using these blindly, as in the Wang tiling approach described in Section 4.2 where only undecidability of tiling was used, we prove and then exploit very particular properties of an aperiodic tiling due to Robinson [Rob71]. These will allow us, using the ideas discussed in Section 4.4 and Section 4.5, to encode in a quantum local Hamiltonian the *parallel* execution 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 parallel evolution 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 [Ber66] original proof of undecidability of tiling. Robinson [Rob71] 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.

The Robinson tile set *can* tile the infinite plane, but only aperiodically. The aperiodic pattern generated by a Robinson tiling is shown in Figure 11. We have already seen in Section 4.2 that tiling problems can easily be turned into translationally-invariant classical Hamiltonians. So we can readily turn this into a classical tiling 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 Hamiltonian from the previous section along the borders of all the squares in Figure 11. (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* edge of the squares in the Robinson tiling. The effective tape length of this Turing Machine is limited only by the size of the border it “runs” on. But the Robinson tiling contains borders of all finite sizes.¹ So this Hamiltonian encodes Turing Machines with all possible tape lengths. All of these

¹Actually, all powers of 4, but this makes no difference to the argument.

Turing Machines are running on the same input, encoded in matrix elements of the translationally-invariant local interaction.

If the universal Turing Machine eventually halts on this input, then for all borders 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 border 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 this 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 (border size) is too small to contain a valid upper-bound on the number of digits in the phase, then the quantum state left on the tape by the phase-estimation QTM is arbitrary (see Section 4.4). 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 borders that are too small. But, crucially, we know the maximum size of these borders: it is determined by the number of binary digits in the phase encoded in the matrix elements of the Hamiltonian, *which we chose*. The energy contribution from the small borders is therefore a computable number, e.g. brute-force diagonalisation of the Hamiltonian on the small borders. Indeed, since this is an eigenvalue computation, this quantity is even an algebraic number. The number of borders of any given size grows quadratically in the lattice size, so the total energy contribution from the small borders can be removed by subtracting an appropriately weighted $\mathbb{1}$ term from the Hamiltonian.

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 6 and hence prove Theorem 3. The difference in 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 borders has energy that diverges with the lattice size in the halting case;

once the lattice is large enough, the number of borders 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 cannot be the ground state in the halting case; we can lower the energy by “breaking” some of the Turing Machines, by introducing defects in the tiling layer. In Section 8, we prove strong rigidity properties of the Robinson tiling, which show that the tiling pattern is robust against defects: any defect (non-matching adjacent tiles) in the tiling only affects the pattern of borders in a finite ball around the defect. Thus destroying n borders requires $O(n)$ defects, each defect contributes $O(1)$ energy, and Turing Machines on intact borders contribute $O(1)$ energy. Thus, 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 6.

Making this argument rigorous is somewhat involved, and is the content of Section 9.1. Section 9.2 then combines this result with any frustration-free gapless quantum Hamiltonian (of which there are many well-known examples) to complete the proof of our main result, Theorem 3: undecidability of the spectral gap for 2D translationally-invariant Hamiltonians with fixed finite local Hilbert space dimension and local interactions specified by algebraic matrix elements, even when the Hamiltonians are guaranteed to be either gapped with a spectral gap uniformly lower-bounded by the local interaction strength (Definition 1), or gapless with continuous spectrum above the ground state (Definition 2).

4.7 Structure of the paper

We warm up in Section 5 by giving a complete proof of the weaker but substantially easier undecidability result for the spectral gap of families of Hamiltonians with unbounded local Hilbert space dimension. We also extend this to prove undecidability of essentially all low-energy properties of quantum many-body Hamiltonians.

The bulk of the paper is concerned with proving our main result: undecidability of the spectral gap for Hamiltonians with fixed local dimension. The proof is broken down into several stages, each described in a separate section. Each of these sections is self-contained, in the sense that only the main theorem of the section is needed in subsequent sections.¹ Proofs in any individual section can therefore be skipped on a first reading, as long as the main result of the section is provisionally

¹Except in the case of Section 8, which collects a number of new results on rigidity of the Robinson tiling, needed in subsequent sections.

accepted.

In Section 6, we give an explicit construction of an exact quantum phase estimation Turing machine, which will be crucial to achieving constant local dimension. In Section 7 we describe a modification of the Gottesman and Irani [GI09] encoding of a quantum Turing machine in a 1D, local, translationally-invariant Hamiltonian. Section 8 contains the analysis of the rigidity of the Robinson quasi-periodic tiling, which will allow us to achieve the required spectral properties.

With this machinery in place, Section 9 proves our main result. This is divided into two steps. As discussed in Section 4, the first step is to prove undecidability of the ground state energy density problem, which we do in Section 9.1, making use of all the results from the previous sections. The second step, in Section 9.2, reduces the ground state energy density problem to the spectral gap problem, completing the proof of our main result. Section 9.3 extends these results to periodic boundary conditions.

Part II

5 Unconstrained local Hilbert space dimension

In this part we will describe two approaches that exploit known undecidability results for tiling and completion problems. Based on these, we can derive undecidability results for the spectral gap and for other low energy properties of translational invariant nearest-neighbour Hamiltonians on a square lattice. In contrast to the approach in the main part of this paper, however, the families of Hamiltonians that we construct here, defined by their local interactions $\{(h_{\text{row}}(n), h_{\text{col}}(n))\}_{n \in \mathbb{N}}$, have local Hilbert space dimension d_n that differs for different elements n of the family, in such a way that $\{d_n\}_{n \in \mathbb{N}}$ is unbounded.

5.1 Undecidability of the spectral gap via tiling

As in the more sophisticated constructions that will appear later, the idea is to reduce an undecidable ground state energy problem to the spectral gap problem. If we do not constrain the local Hilbert space dimension, then this reduction can be chosen such that it directly exploits the undecidability of a tiling problem. To this end, we need two ingredients:

5.1.1 Ingredient 1: a tiling Hamiltonian

We saw in Section 4.2 that a tiling problem can easily be represented as a ground state energy problem for the local Hamiltonian

$$h_c^{(i,j)} := \sum_{(m,n) \in C^{(i,j)}} |m\rangle\langle m|^{(i)} \otimes |n\rangle\langle n|^{(j)}, \quad (13)$$

$$H_c^{\wedge(L)} := \sum_{(i,j) \in \mathcal{E}} h_c^{(i,j)}, \quad (14)$$

where the local basis $|m\rangle$ label tile types, and the set of constraints $C^{(i,j)} \subseteq \mathcal{K} \times \mathcal{K}$ includes all pairs of tiles (m, n) which are incompatible when placed on adjacent sites i and j .

By construction we have $h_c^{(i,j)} \geq 0$, $H_c^{\wedge(L)}$ is translational invariant and its spectrum is contained in \mathbb{N}_0 . If there exists a tiling of the plane, then for open boundary conditions $\forall L : 0 \in \text{spec } H_c^{\wedge(L)}$. Similarly, in the case of periodic boundary conditions, the existence of a periodic tiling implies that $0 \in \text{spec } H_c^{\wedge(L)}$ holds for an unbounded sequence of L 's. On the other hand, if there is no tiling, then there is an L_0 such that

$$\forall L > L_0 : \text{spec } H_c^{\wedge(L)} \geq 1. \quad (15)$$

Also note that this implies that in the case of open boundary conditions the ground state energy density is $E_\rho \geq 1/(L_0 + 1)^2$ if no tiling exists, compared to $E_\rho = 0$ if there is a tiling of the plane.

5.1.2 Ingredient 2: a gapless frustration-free Hamiltonian

As a second ingredient we will use that there are frustration-free two-body Hamiltonians

$$H_q^{\Lambda(L)} := \sum_{(i,j) \in \mathcal{E}} h_q^{(i,j)}, \quad (16)$$

such that for $L \rightarrow \infty$ we get $\text{spec } H_q^{\Lambda(L)} \rightarrow \mathbb{R}_+$ in the sense that the spectrum of the finite system approaches a dense subset of \mathbb{R}_+ [FG+12]. We will denote the corresponding single site Hilbert space by $\mathcal{H}_q^{(i)} \simeq \mathbb{C}^D$ and we can in fact choose $D = 2$ for instance by assigning to each row of the lattice the XY-model with transversal field taken at a gapless point with product ground state. Note that in this case the entries of the matrices $h_q^{(i,j)}$ are rational.

5.1.3 Reducing tiling to spectral gap

In order to fruitfully merge the ingredients, we assign a Hilbert space $\mathcal{H}^{(i)} := \mathcal{H}_0^{(i)} \oplus \mathcal{H}_c^{(i)} \otimes \mathcal{H}_q^{(i)} \simeq \mathbb{C}^1 \oplus \mathbb{C}^K \otimes \mathbb{C}^D$ to each site $i \in \Lambda(L)$. A corresponding orthonormal set of basis vectors will be denoted by $|0\rangle \in \mathbb{C}^1$ and $|k, \alpha\rangle := |k\rangle \otimes |\alpha\rangle \in \mathbb{C}^K \otimes \mathbb{C}^D$, respectively. The Hamiltonian is then defined in terms of the two-body interactions

$$H^{(i,j)} := |0\rangle\langle 0|^{(i)} \otimes \mathbb{1}_{c_q}^{(j)} + \mathbb{1}_{c_q}^{(i)} \otimes |0\rangle\langle 0|^{(j)} \quad (17)$$

$$+ \sum_{(m,n) \in \mathcal{C}^{(i,j)}} |m\rangle\langle m|^{(i)} \otimes \mathbb{1}_q^{(i)} \otimes |n\rangle\langle n|^{(j)} \otimes \mathbb{1}_q^{(j)} \quad (18)$$

$$+ \sum_{\alpha, \beta, \gamma, \delta=1}^D \mathbb{1}_c^{(i)} \otimes |\alpha\rangle\langle \beta|^{(i)} \otimes \mathbb{1}_c^{(j)} \otimes |\delta\rangle\langle \gamma|^{(j)} \langle \alpha, \delta | h_q^{(i,j)} | \beta, \gamma \rangle. \quad (19)$$

Here $\mathbb{1}_c, \mathbb{1}_q$ and $\mathbb{1}_{c_q}$ denote the identity operators on $\mathcal{H}_c, \mathcal{H}_q$ and $\mathcal{H}_c \otimes \mathcal{H}_q$, respectively. As before, we define the Hamiltonian assigned to a square $\Lambda(L)$ as

$$H^{\Lambda(L)} := \sum_{(i,j) \in \mathcal{E}} H^{(i,j)}. \quad (20)$$

Theorem 7 (Reducing tiling to spectral gap) *Consider any set \mathcal{K} of Wang tiles, and the corresponding family of two-body Hamiltonians $\{H^{\Lambda(L)}\}_L$ on square lattices $\Lambda(L)$ either with open or with periodic boundary conditions.*

(i). *The family of Hamiltonians is frustration-free.*

(ii). If \mathcal{K} tiles the plane \mathbb{Z}^2 , then in the case of open boundary conditions, we have that $\text{spec } H^{\Lambda(L)} \rightarrow \mathbb{R}_+$ for $L \rightarrow \infty$, i.e., there is no gap and an excitation spectrum that becomes dense in \mathbb{R}_+ . Similarly, if \mathcal{K} tiles any torus, then in the case of periodic boundary conditions $\text{spec } H^{\Lambda(L_i)} \rightarrow \mathbb{R}_+$ for a subsequence $L_i \rightarrow \infty$.

(iii). If \mathcal{K} does not tile the plane, then there is an L_0 such that for all $L > L_0$ $H^{\Lambda(L)}$ has a unique ground state and a spectral gap of size at least one, i.e.

$$\text{spec } H^{\Lambda(L)} \setminus \{0\} \geq 1. \quad (21)$$

Proof Frustration-freeness is evident since $H^{(i,j)} \geq 0$ and $H^{\Lambda(L)} |0, \dots, 0\rangle = 0$. In order to arrive at the other assertions, we decompose the Hamiltonian as $H^{\Lambda(L)} =: H_0 + H_c + H_q$, where the three terms on the right are defined by taking the sum over edges separately for the expressions in (17), (18) and (19), respectively. Let us now assign a *signature* $\sigma \in \{0, \dots, K\}^{L^2}$ to every state of our computational product basis, so that $|0\rangle^{(i)}$ translates to $\sigma_i = 0$ and $|k, \alpha\rangle^{(i)}$ to $\sigma_i = k$, irrespective of α . By collecting computational basis states with the same signature we can then decompose the Hilbert space as

$$\bigotimes_{i \in \Lambda} \mathcal{H}^{(i)} \simeq \bigoplus_{\sigma} \mathcal{H}_{\sigma}. \quad (22)$$

The reason behind this is that the Hamiltonian is block diagonal w.r.t. this decomposition, i.e. it can be written as

$$H^{\Lambda(L)} = \bigoplus_{\sigma} H_{\sigma}, \quad \text{and thus} \quad \text{spec } H^{\Lambda(L)} = \bigcup_{\sigma} \text{spec } H_{\sigma}. \quad (23)$$

In order to identify the spectra coming from different signatures we will distinguish three cases:

- (i). $\sigma = (0, \dots, 0)$: this yields an eigenvalue 0 as already mentioned.
- (ii). $\sigma \neq (0, \dots, 0)$ but $\sigma_i = 0$ for some i : for any unit vector $|\psi_{\sigma}\rangle$ with such a signature we have $\langle \psi_{\sigma} | H_0 | \psi_{\sigma} \rangle \geq 1$. We will denote the part of the spectrum which corresponds to these signatures by $S \subset \mathbb{R}$. The only property we will use is that $S \geq 1$.
- (iii). $\sigma \in \{1, \dots, K\}^{L^2}$: In the subspace spanned by all states whose signature does not contain 0 we have that $H_0 = 0$ and $H_c + H_q = H_c^{\Lambda(L)} \otimes \mathbb{1} + \mathbb{1} \otimes H_q^{\Lambda(L)}$. Consequently, the spectrum stemming from this subspace equals $\text{spec } H_c^{\Lambda(L)} + \text{spec } H_q^{\Lambda(L)}$.

Putting things together, we obtain

$$\text{spec } H^{\Lambda(L)} = \{0\} \cup \left\{ \text{spec } H_c^{\Lambda(L)} + \text{spec } H_q^{\Lambda(L)} \right\} \cup S, \quad (24)$$

so that the claimed equivalence between the impossibility of tiling with \mathcal{K} and the existence of a spectral gap of size at least 1 follows from the properties of the chosen ingredients. Uniqueness of the ground state in cases where no tiling exists follows from the above argument when considering the spectra as multisets rather than sets. \square

From the undecidability of the tiling problem we now obtain:

Corollary 8 (Undecidability of spectral gap for unconstrained dimension)

Let $\epsilon > 0$. Consider all families $\{H^{\Lambda(L)}\}$ of frustration-free translational invariant nearest neighbour Hamiltonians on 2D square lattices with open boundary conditions that are described by a pair of Hermitian matrices $(h_{\text{row}}, h_{\text{col}})$ ¹ with rational entries and operator norm smaller than $1 + \epsilon$. There is no algorithm that upon input of $(h_{\text{row}}, h_{\text{col}})$ decides whether these describe a gapped or a gapless system, even under the promise that one of them is true and that in the gapped case the gap is at least of size 1.

Here, the norm bound follows from the observation that $\|H^{(i,j)}\| \leq 1 + \|h_q^{(i,j)}\|$ where $h_q^{(i,j)}$ can be rescaled to arbitrary small norm.

For the case of periodic boundary conditions we obtain the same statement if we replace our strong definition of ‘gapless’ by the weaker requirement that $\exists c > 0 \forall \epsilon > 0 \forall L_0 \exists L > L_0$ so that every point in $[\lambda_0(H^{\Lambda(L)}), \lambda_0(H^{\Lambda(L)}) + c]$ is within distance ϵ from $\text{spec } H^{\Lambda(L)}$. The reason for this is, that if the period of the torus does not match the required one, then a gap can be generated that, however, disappears again if we enlarge the size of the system.

We emphasise that, so far, there is no constraint on the local Hilbert space dimension. Since every instance in the above construction has a finite local Hilbert space we can, however, at least conclude axiomatic undecidability for finite local Hilbert spaces. That is, for every consistent formal system with a recursive set of axioms, there exists an instance of a Hamiltonian with properties as in the corollary and finite dimensional local Hilbert space such that neither the presence nor the absence of a spectral gap can be proven from the axioms. However, the local dimension, whilst finite, is uncomputable. In the main part of the paper, we will prove the much stronger Theorem 3, which shows that the spectral gap problem remains undecidable even for fixed local dimension. This implies axiomatic independence for Hamiltonians where the finite local dimension is known.

¹Here we make no use of an additional on-site term, so $h_1 = 0$.

In the above construction, gapped cases are related to the impossibility of tiling. Since the latter always admits a proof, the axiomatically undecidable cases coming from this construction have to be gapless. The following section will provide a construction where this asymmetry is reversed.

5.2 Undecidability of low energy properties

5.2.1 Reducing tile completion to a ground state energy problem

We now modify the above construction in order to show that not only the spectral gap but many other low energy properties are undecidable as well. The idea, illustrated in the case of the spectral gap, is to invert the relation between gap and existence of tiling. In the previous construction, the existence of a gap was associated with the impossibility of a tiling. Now we want to associate it with the *existence* of a tiling. The drawback is that we lose the frustration free property in that case. One of the advantages is that we do not have to rely on the undecidability of tiling (which is a rather non-trivial result [Ber66; Rob71]) but can exploit the simpler undecidability of the *completion problem*, already proven by Wang [Wan61]. In the completion problem, one fixes one tile at the left-bottom corner and asks whether the first quadrant can be tiled with a given set of tiles.

Undecidability of the completion problem is relatively simple, and explained in full detail in Robinson [Rob71, Section 4], by reduction from the halting problem for Turing Machines. Any Turing machine is specified by a finite number of states $Q = \{q_0, q_1, \dots\}$, a finite alphabet for the tape $S = \{s_0, s_1, \dots\}$, where q_0 is the initial state and s_0 the blank, and quintuplets of the forms

$$(q, s, s', D, q') \in Q \times S \times S \times \{\text{left, right}\} \times Q$$

meaning that if the head of the machine reads s and the machine is in state q , then it will write s' into the current cell on the tape, transit to state q' and move in the direction D . Starting in q, s , there must be at most one valid quintuplet. If there is none, then the machine halts.¹

Following Robinson [Rob71], we define the tiles in Figure 4, where colours are identified with labelled arrows, in such a way that matching colours are indicated by an arrow head abutting a tail with the same label. We denote the set of tiles by \mathcal{K} , which contains at most $K := |\mathcal{K}| \leq 2 + (3|Q| + 1)|S|$ elements.

We will assume that the Turing machine is constructed such that, when running on the initially blank tape, its head never moves to the left of the starting cell. Note that this can be assumed without loss of generality, since every Turing machine can be simulated by one with this property by standard techniques, which amount to ‘folding’ the tape and thereby doubling the tape alphabet.

¹Note that this definition is slightly different from the one used later in Section 6.




Figure 4: Tiles for encoding a Turing machine with half-infinite tape into a completion problem. The tile in the bottom-left corner is fixed and, together with the second tile in the third row, starts the Turing machines running on the blank half-infinite tape. The tiles in the second row are the 'action tiles' that are able to change the state of the Turing machine and correspond to a left or right moving head. For each pair (q, s) there is at most one action tile.

If the tile that appears in the bottom-left corner in Figure 4, which we denote by \square , is fixed in the bottom-left corner of an $L \times L$ lattice, then any valid tiling corresponds to the evolution of the Turing Machine on the blank tape up to time L , where time goes up in the vertical direction and each row of tiles corresponds to the total configuration (including the tape) of the Turing machine in two consecutive instants of time (represented by the bottom and upper edges of that row of tiles). Hence,

- (1). There exists a valid tile completion for all $L \times L$ squares if and only if the Turing Machine, when running on an initially blank tape, does not halt.
- (2). If such a tiling exists, then it is unique.

Now, the completion problem can be represented as a ground state energy problem in the following straightforward way. Consider a square lattice $\Lambda(L)$ with open boundary conditions and corresponding edge set \mathcal{E} .

We will assign an energy penalty or energy bonus $w_{(i,j)}(m, n) \in \mathbb{Z}$ to an edge $(i, j) \in \mathcal{E}$ on which a pair of tiles $(m, n) \subseteq \mathcal{K} \times \mathcal{K}$ is placed. The choice of the weights is summarised in the following table. Here \square stands for any tile different from \square and the numbers in brackets define the weights in cases where the colours

(i.e., labels and arrows) match.

$$\begin{array}{c|cc}
 & \text{Right tile} & \\
 \hline
 \text{Left tile} & \begin{array}{cc} \boxed{\leftarrow} & \boxed{\bullet} \\ 4 & 2 \ (-1) \\ \hline \boxed{\bullet} & 4 & 2 \ (0) \end{array} & \begin{array}{c|cc}
 & \text{Tile above} & \\
 \hline
 \text{Tile below} & \begin{array}{cc} \boxed{\leftarrow} & \boxed{\bullet} \\ 4 & 2 \ (-1) \\ \hline \boxed{\bullet} & 4 & 2 \ (0) \end{array} & \end{array} \quad (25)
 \end{array}$$

Assigning a Hilbert space $\mathcal{H}_u^{(i)} \simeq \mathbb{C}^K$ to each site $i \in \Lambda(L)$ whose computational basis states are labelled by tiles we define a Hamiltonian $H_u^{\Lambda(L)} := \sum_{(i,j) \in \mathcal{E}} h_u^{(i,j)}$, with

$$h_u^{(i,j)} := \sum_{m,n \in \mathcal{K}} \omega_{(i,j)}(m,n) |m\rangle\langle m|_{(i)} \otimes |n\rangle\langle n|_{(j)}. \quad (26)$$

By construction, the Hamiltonian is diagonal in computational basis, its eigenvectors correspond to tile configurations $\Lambda(L) \rightarrow \mathcal{K}$ and, since all weights are integers, its spectrum lies in \mathbb{Z} . In order to determine the ground state energy note that the tile $\boxed{\leftarrow}$ is the only one that can lead to negative weights. Assume first that a tile $\boxed{\leftarrow}$ is placed on a site different from the bottom-left corner. Then this tile will have at least one neighbour below with an arrow pointing upwards or one neighbour to the left with an arrow pointing to the right. In either case there will be an energy penalty +4 so that the weights involving this tile will sum up to at least 2. If, however, $\boxed{\leftarrow}$ is placed in the bottom-left corner, then no such energy penalty is picked up and the tile can contribute -2 to the energy. If the tiling is then completed following the evolution of the encoded Turing machine on a blank tape, then no additional penalty will occur on $\Lambda(L)$ if until the L 'th time step either the head of the Turing machines moves left of the origin or the Turing machine does not halt. Hence, in both cases $\lambda_0(H_u^{\Lambda(L)}) = -2$ and the ground state is unique and given by a product state.

If the Turing machine halts, then the completion problem has no solution beyond some L_0 and every configuration with $\boxed{\leftarrow}$ at its bottom-left corner will get an additional penalty of at least 2 leading to non-negative energy. Since there is always a configuration that achieves zero energy, e.g., one using only the type of tiles on the top-left in Figure 4, we obtain:

Lemma 9 (Relating ground state energy to the halting problem) *Consider any Turing machine with state set Q and tape alphabet S running on an initially blank tape such that its head never moves left of the starting cell. There is a family of translational invariant nearest neighbour Hamiltonians $\{H_u^{\Lambda(L)}\}_L$ (specified above) on the 2D square lattice with open boundary conditions and local Hilbert space dimension at most $|S|(3|Q| + 1) + 2$ such that*

- (i). If until time L the Turing machines did not halt, then $\lambda_0(H_u^{\wedge(L)}) = -2$. In this case the corresponding ground state is a unique product state and the only eigenstate with negative energy.
- (ii). If the Turing machine halts, then $\exists L_0$ (given by the halting time) such that $\forall L > L_0 : \lambda_0(H_u^{\wedge(L)}) = 0$.

Since every Turing machine can be simulated by one with a half-infinite tape, which automatically satisfies the constraint that its head never moves to the left of the origin, the above construction leads to an undecidable ground state energy problem, which we will exploit in the following.

5.2.2 Reduction of the halting problem to arbitrary low energy properties

Consider translational invariant Hamiltonians on square lattices with open boundary conditions. Let $H_q^{\wedge(L)}$ describe such a family of Hamiltonians on $L \times L$ lattices. Our aim is to prove the undecidability of essentially any low energy property displayed by this Hamiltonian in the thermodynamic limit that distinguishes it from a gapped system with unique product ground state, e.g. the existence of topological order, or non-vanishing correlation functions. This is implied by the following theorem:

Theorem 10 (Relating low energy properties to the halting problem) *Let TM be a Turing machine with state set Q and alphabet S running on an initially blank tape. Consider the class of translational invariant Hamiltonians on square lattices with open boundary conditions. Let $H_q^{\wedge(L)}$ with nearest-neighbour interaction $h_q^{(i,j)} \in \mathcal{B}(\mathcal{H}_q^{(i)} \otimes \mathcal{H}_q^{(j)})$ with $\mathcal{H}_q^{(i)} \simeq \mathbb{C}^q$ and on-site term $h_{q1}^{(k)} \in \mathcal{B}(\mathcal{H}_q^{(k)})$ describe such a Hamiltonian and assume that there is a $c \in \mathbb{R}$ such that $\lambda_0(H_q^{\wedge(L)}) - cL^2 \in [-\frac{1}{2}, \frac{1}{2}]$ for all lattice sizes L . Then there exists a family of Hamiltonians $H^{\wedge(L)}$ (specified in the proof) with the following properties:*

- (i). $H^{\wedge(L)}$ is translational invariant on the square lattice with open boundary conditions and local Hilbert spaces $\mathcal{H}^{(i)} \simeq \mathbb{C}^d$ of dimension $d \leq 2(q+1) + |S|(3|Q|+1)$.
- (ii). The interactions of $H^{\wedge(L)}$ are nearest-neighbour (possibly with on-site terms), computable and independent of L .
- (iii). If up to time L the TM did not halt, then $H^{\wedge(L)}$ has a non-degenerate product ground state and a spectral gap $\Delta(H^{\wedge(L)}) \geq \frac{1}{2}$.
- (iv). If TM halts, then $\exists L_0$ (given by the halting time) such that $\forall L > L_0$ the following identity between multisets holds in the interval $[0, \frac{1}{2})$:

$$\text{spec } H_q^{\wedge(L)} - \lambda_0(H_q^{\wedge(L)}) = \text{spec } H^{\wedge(L)} - \lambda_0(H^{\wedge(L)}). \quad (27)$$

In this case there exist isometries $V^{(i)} : \mathcal{H}_q^{(i)} \rightarrow \mathcal{H}^{(i)}$, $V := \bigotimes_i V^{(i)}$ such that for this part of the spectrum any eigenstate $|\phi_q\rangle$ of $H_q^{\Lambda(L)}$ is mapped to the corresponding eigenstate $|\phi\rangle$ of $H^{\Lambda(L)}$ via $V|\phi_q\rangle = |\phi\rangle$.

Some remarks before the proof. Since, on the Turing machine level, (iii) and (iv) cannot be told apart by any effective algorithm, the same has to hold for any property that distinguishes a gapped system with product ground state from the low energy sector of some frustration free Hamiltonian. With small modifications in the proof, the nearest neighbour assumption for h_q could be replaced by any fixed local interaction geometry. In this case, $H^{\Lambda(L)}$ will of course inherit this interaction geometry.

Proof W.l.o.g. $c = 0$ since we can always compensate for it by adding a multiple of the identity to the on-site part of the Hamiltonian.

For later use we will embed the given frustration free Hamiltonian in a larger systems such that its ground state energy gets an offset -1 . To this end, define a translational invariant nearest neighbour Hamiltonian on an auxiliary system with local Hilbert space $\mathcal{H}_a^{(i)} \simeq \mathbb{C}^2$ via

$$\eta := \sum_{(i,j) \in \mathcal{E}} \eta^{(i,j)}, \quad \eta^{(i,j)} := -\frac{1}{2}|0\rangle\langle 0|^{(i)} \otimes |1\rangle\langle 1|^{(j)} + 2\mathbb{1}^{(i)} \otimes |0\rangle\langle 0|^{(j)}.$$

Note that $\lambda_0(\eta) = -1$ and $\Delta(\eta) = 1$. Now introduce $\mathcal{H}_Q^{(i)} := \mathcal{H}_a^{(i)} \otimes \mathcal{H}_q^{(i)}$ and set $h_Q^{(i,j)} := \eta^{(i,j)} \otimes \mathbb{1}_q^{(i,j)} + \mathbb{1}_a^{(i,j)} \otimes h_q^{(i,j)}$ acting on $\mathcal{H}_Q^{(i)} \mathcal{H}_Q^{(j)}$ (up to reordering of Hilbert spaces). Then $H_Q^{\Lambda(L)} := \sum_{(i,j) \in \mathcal{E}} h_Q^{(i,j)} + \sum_{k \in \Lambda(L)} \mathbb{1}_a^{(k)} \otimes h_{q_1}^{(k)}$ has the property that in the interval $(-\infty, \lambda_0(H_Q^{\Lambda(L)}))$ we have that

$$\text{spec } H_Q^{\Lambda(L)} - 1 = \text{spec } H_q^{\Lambda(L)}$$

holds as an identity between multisets (i.e., including multiplicities) and the corresponding eigenstates are related via a local isometric embedding.

We will now combine this with the Hamiltonian from the previous subsection by enlarging the local Hilbert space to $\mathcal{H}^{(i)} := \mathcal{H}_u^{(i)} \oplus \mathcal{H}_Q^{(i)}$ and defining

$$h^{(i,j)} := h_Q^{(i,j)} + h_u^{(i,j)} + 7\mathbb{1}_Q^{(i)} \otimes \mathbb{1}_u^{(j)}, \quad (28)$$

where $h_Q^{(i,j)}$ and $h_u^{(i,j)}$ act nontrivially on $\mathcal{H}_Q^{(i)} \otimes \mathcal{H}_Q^{(j)}$ and $\mathcal{H}_u^{(i)} \otimes \mathcal{H}_u^{(j)}$ respectively and $\mathbb{1}_u$ and $\mathbb{1}_Q$ denote the projectors onto the respective subspaces. The Hamiltonian on the square $\Lambda(L)$ is now

$$H^{\Lambda(L)} = \sum_{(i,j) \in \mathcal{E}} h^{(i,j)} + \sum_{k \in \Lambda(L)} \mathbb{1}_a^{(k)} \otimes h_{q_1}^{(k)}. \quad (29)$$

We decompose it as $H^{\Lambda(L)} =: \tilde{H}_q + \tilde{H}_u + \tilde{H}_p$, where the three terms on the right are defined by taking the sum over edges separately for the three interactions in (28) and \tilde{H}_q contains the on-site terms in addition. Let us now assign a *signature* $\sigma \in \{0, 1\}^{L^2}$ to every state of our computational product basis, so that $|i\rangle^{(j)}$ translates to $\sigma_j = 0$ for all $|i\rangle \in \mathcal{H}_u^{(j)}$ and $\sigma_j = 1$ for all $|i\rangle \in \mathcal{H}_Q^{(j)}$. By collecting computational basis states with the same signature, we can then decompose the Hilbert space as

$$\bigotimes_{i \in \Lambda} \mathcal{H}^{(i)} \simeq \bigoplus_{\sigma} \mathcal{H}_{\sigma}. \quad (30)$$

The reason behind this is that the Hamiltonian is block diagonal w.r.t. this decomposition, i.e. it can be written as

$$H^{\Lambda(L)} = \bigoplus_{\sigma} H_{\sigma}, \quad \text{and thus} \quad \text{spec } H^{\Lambda(L)} = \bigcup_{\sigma} \text{spec } H_{\sigma}. \quad (31)$$

In order to identify the spectra and eigenstates coming from different signatures we will distinguish three cases:

Case 1: $\sigma = (0, \dots, 0)$. In this sector $\tilde{H}_p = \tilde{H}_q = 0$ so that we are left with $H_u^{\Lambda(L)}$.

Case 2: $\sigma = (1, \dots, 1)$. In this sector $\tilde{H}_p = \tilde{H}_u = 0$ so that we are left with $H_Q^{\Lambda(L)}$.

Case 3: $\sigma \neq (0, \dots, 0)$ but $\sigma_i = 0$ for some i . Let us define $A_{\sigma} := \{i \in \Lambda(L) | \sigma_i = 0\}$. Decompose A_{σ} into rectangles row by row, as in Figure 5 following the following conventions for the rectangles:

- (i) They live on the dual lattice.
- (ii) If one is on top of another of the same size (such as those marked 1 and 2 in the figure), we glue them together into a single larger rectangle.
- (iii) For each rectangle, we choose a neighbouring pair (i, j) such that both i and j are in the same row, $i \in A_{\sigma}$, $j \notin A_{\sigma}$. This selection is illustrated in the figure by black dots (joined by a link). If this is not possible because the rectangle is a whole row (as for the rectangle marked 3 in the figure), we take j in the next (or previous) possible row after applying convention (ii). An example is shown in the figure as a pair of linked purple dots.
- (iv) For each overlapping region between pairs of rectangles (such as those marked with dashed red lines in the figure), we also choose a neighbouring pair (i, j) with $i \in A_{\sigma}$, $j \notin A_{\sigma}$, but now i belongs to one of the rectangles (the larger one is a valid option) and j in the row of the other rectangle. This selection is illustrated in the figure by linked red dots.

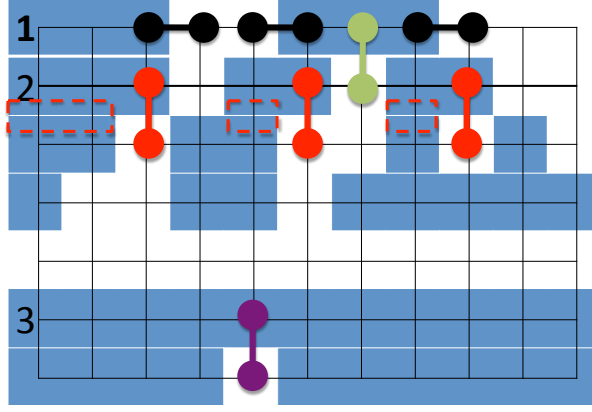


Figure 5: Illustration of the counting explained in the main text.

How can we choose all these pairs (i, j) with the minimal number of repetitions? To repeat a pair, two things may happen. Either we have to choose the same vertical pair to select two different overlapping regions (located at its right and left) – illustrated by a pair of green dots in the figure – or we have to choose the same vertical pair to select one or two overlapping regions and also a rectangle made out of a full row. In any case, each pair is chosen not more than three times. For instance, the purple joint dots in the figure are chosen three times. With this analysis we conclude that if we denote by n_{A_σ} the number of rectangles, m_{A_σ} the number of overlaps between them and k_{A_σ} the number of chosen different pairs (i, j) with $i \in A_\sigma, j \notin A_\sigma$, we have that

$$3k_{A_\sigma} \geq n_{A_\sigma} + m_{A_\sigma}. \quad (32)$$

With these conventions we divide the interactions of \tilde{H}_u into those within the rectangles $H_{\text{rectangles}}$, those in the overlaps of rectangles H_{overlaps} and the rest H_{rest} . We have that for any unit vector $|\psi_\sigma\rangle$ with signature σ :

$$\begin{aligned} \langle \psi_\sigma | H_{\text{rectangles}} | \psi_\sigma \rangle &\geq -2n_{A_\sigma}, \\ \langle \psi_\sigma | H_{\text{overlap}} | \psi_\sigma \rangle &\geq -2m_{A_\sigma}, \\ \langle \psi_\sigma | H_{\text{rest}} | \psi_\sigma \rangle &= 0, \end{aligned}$$

which gives

$$\langle \psi_\sigma | \tilde{H}_u | \psi_\sigma \rangle \geq -2(n_{A_\sigma} + m_{A_\sigma}). \quad (33)$$

The reason behind is that, by the same argument as in the previous section, $\lambda_0(H_u^R) \geq -2$ for all rectangles R . Finally, we have trivially

$$\langle \psi_\sigma | \tilde{H}_q | \psi_\sigma \rangle \geq \lambda_0(H_q^{\Lambda(L)}) - 1, \quad (34)$$

$$\langle \psi_\sigma | \tilde{H}_p | \psi_\sigma \rangle \geq 7k_{A_\sigma}. \quad (35)$$

Therefore, using (32)–(35), we get $\langle \psi_\sigma | H^{\Lambda(L)} | \psi_\sigma \rangle - \lambda_0(H_q^{\Lambda(L)}) \geq k_{A_\sigma} - 1 \geq 0$. Hence, if we denote the part of the spectrum which corresponds to these mixed signatures by $S \subset \mathbb{R}$ we have that $S \geq \lambda_0(H_q^{\Lambda(L)})$.

Putting things together, we get that the spectrum is decomposed according to the signatures as

$$\text{spec } H^{\Lambda(L)} = \text{spec } H_u^{\Lambda(L)} \cup \text{spec } H_Q^{\Lambda(L)} \cup S \quad (36)$$

where

$$\lambda_0(H_Q^{\Lambda(L)}) = \lambda_0(H_q^{\Lambda(L)}) - 1, \quad S \geq -\frac{1}{2}, \quad (37)$$

and according to Lemma 9 either $\text{spec } H_u^{\Lambda(L)} \geq 0$ or $H_u^{\Lambda(L)}$ has one negative eigenvalue -2 , corresponding to a product ground state. Which is the case depends on the evolution of the encoded Turing machine (and thus becomes undecidable). \square

Part III

6 Quantum Phase Estimation Turing Machine

We will see in Section 7 how to encode an arbitrary Quantum Turing Machine (QTM) into a translationally-invariant nearest-neighbour Hamiltonian, in such a way that the local Hilbert space dimension is a function only of the alphabet size and number of internal states. This will allow us to encode any specific universal (reversible) Turing Machine in a Hamiltonian with fixed local dimension. However, to prove undecidability by reduction from the Halting Problem, we will need a way to feed any desired input to this universal TM.

to this end, the main result of this section is to construct a family of quantum Turing Machines, all of which have the same alphabet size and number of internal states, but whose transition rules vary. For any given string, we can choose the transition rules such that the QTM, when started from the empty input, writes out the string to its tape and then halts deterministically.

More precisely, we will prove the following result (see below for the necessarily basic definitions, notations, and facts about QTMs).

Theorem 11 (Phase-estimation QTM) *There exists a family of well-formed, normal form, unidirectional 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 behaves properly, halts deterministically after $O(\text{poly}(N)2^N)$ steps, uses $N + 3$ space, and outputs the binary expansion of n (padded to N digits with leading 0's). (Here, $|n|$ 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*

$$\left\{ 0, 1, \pm \frac{1}{\sqrt{2}}, e^{i\pi\varphi}, e^{i\pi 2^{-|n|}} \right\} \quad (38)$$

where $\varphi \in \mathbb{Q}$.

Part (iii) implies that the only dependence of the transition rules of P_n on n is that implicit in φ , which is defined as the rational number whose binary decimal expansion contains the digits of n in reverse order after the decimal.

Note that the input N does *not* determine the output string that gets written to the tape; it only determines the number of binary digits in the output. The number

represented by that output is determined (up to padding with leading zeros) by the choice of the parameter n for the QTM P_n .

Theorem 11 is essentially the only inherently quantum ingredient in our result, and the precise properties asserted there will be *absolutely crucial* to the proof. For example, if instead of writing out the string and halting deterministically, the QTM did this only with arbitrarily high probability, our proof would not go through.

It is therefore not at all clear a priori whether QTMs fulfilling these strict requirements exist. Since our proof depends so delicately on the precise properties of these QTMs, we cannot appeal to previous results. Instead, in this section we give a detailed and explicit construction of a QTM based on the quantum phase estimation algorithm, that fulfils all the requirements of Theorem 11.

6.1 Quantum Turing Machinery

For completeness we quote the basic definitions of Turing Machines and Quantum Turing Machines verbatim from [BV97].

Definition 12 (Turing Machine – Definition 3.2 in [BV97])

A (deterministic) Turing Machine (TM) is defined by a triplet (Σ, Q, δ) where Σ is a finite alphabet with an identified blank symbol $\#$, Q is a finite set of states with an identified initial state q_0 and final state $q_f \neq q_0$, and δ – the transition function – is a function

$$\delta : Q \times \Sigma \rightarrow \Sigma \times Q \times \{L, R\}. \quad (39)$$

The TM has a two-way infinite tape of cells indexed by \mathbb{Z} and a single read/write tape head that moves along the tape. A configuration of the TM is a complete description of the contents of the tape, the location of the tape head and the state $q \in Q$ of the finite control. At any time, only a finite number of the tape cells may contain non-blank symbols.

For any configuration c of the TM, the successor configuration c' is defined by applying the transition function to the current state and the symbol scanned by the head, replacing them by those specified in the transition function and moving the head left (L) or right (R) according to δ .

By convention, the initial configuration satisfies the following conditions: the head is in cell 0, called the starting cell, and the machine is in state q_0 . We say that an initial configuration has input $x \in (\Sigma \setminus \{\#\})^*$ if x is written on the tape in positions $0, 1, 2, \dots$ and all other tape cells are blank. The TM halts on input x if it eventually enters the final state q_f . The number of steps a TM takes to halt on input x is its running time on input x . If a TM halts, then its output is the string in Σ^* consisting of those tape contents from the leftmost non-blank symbol to the rightmost non-blank symbol, or the empty string if the entire tape is blank. A TM is called reversible if each configuration has at most one predecessor.

Let $\tilde{\mathbb{C}}$ be the set consisting of α such that there is a deterministic algorithm that computes the real and imaginary parts of α to within 2^{-n} in time polynomial in n .

Definition 13 (Quantum Turing Machine – Definition 3.2 in [BV97])

A quantum Turing Machine (QTM) is defined by a triplet (Σ, Q, δ) where Σ is a finite alphabet with an identified blank symbol $\#$, Q is a finite set of states with an identified initial state q_0 and final state $q_f \neq q_0$, and δ – the quantum transition function – is a function

$$\delta : Q \times \Sigma \rightarrow \tilde{\mathbb{C}}^{\Sigma \times Q \times \{L,R\}}. \quad (40)$$

The QTM has a two-way infinite tape of cells indexed by \mathbb{Z} and a single read/write tape head that moves along the tape. We define configurations, initial configurations and final configurations exactly as for deterministic TMs.

Let \mathcal{S} be the inner-product space of finite complex linear combinations of configurations of M with the Euclidean norm. We call each element $\phi \in \mathcal{S}$ a superposition of M . The QTM M defines a linear operator $U_M : \mathcal{S} \rightarrow \mathcal{S}$, called the time evolution operator of M , as follows: if M starts in configuration c with current state p and scanned symbol σ , then after one step M will be in superposition of configurations $\psi = \sum_i \alpha_i c_i$, where each nonzero α_i corresponds to the amplitude $\delta(p, \sigma, \tau, q, d)$ of $|\tau\rangle |q\rangle |d\rangle$ in the transition $\delta(p, \sigma)$ and c_i is the new configuration obtained by writing τ , changing the internal state to q and moving the head in the direction of d . Extending this map to the entire \mathcal{S} through linearity gives the linear time evolution operator U_M .

Here, for convenience of programming, we will consider generalised TMs and QTMs in which the head can also stay still (No-movement), as well as move Left or Right.

Definition 14 (Generalised TM and QTM) A generalised TM or generalised QTM is defined exactly as a standard TM (Definition 12) or standard QTM (Definition 13) except that the set of head movement directions is $\{L, R, N\}$ instead of just $\{L, R\}$.

Following Bernstein and Vazirani [BV97], we define various classes of deterministic and quantum Turing Machines:

Definition 15 (Well-formed – Definition 3.3 in [BV97])

We say that a QTM is well-formed if its time evolution operator is an isometry, that is, it preserves the Euclidean norm.

It is easy to see (Theorem 4.2 in [BV97]) that any reversible TM is also a well-formed QTM where the quantum transition function $\delta(p, \sigma, q, \tau, d) = 1$ if $\delta(p, \sigma) = (q, \tau, d)$ for the reversible TM and 0 otherwise.

Definition 16 (Normal form – Definition 3.13 in [BV97])

A well-formed QTM or reversible TM $M = (\Sigma, Q, \delta)$ is in normal form if

$$\forall \sigma \in \Sigma \quad \delta(q_f, \sigma) = |\sigma\rangle |q_0\rangle |N\rangle. \quad (41)$$

In [BV97] the normal-form condition is $\delta(q_f, \sigma) = |\sigma\rangle |q_0\rangle |R\rangle$. However, the choice of final direction is just an arbitrary convention, since the machine already halted and never carries out those transitions. We replace R with N in the definition for technical reasons (see the Reversal Lemma 23, below).

There is a specific class of QTMs, called *unidirectional*, that play an important role in the general theory developed by Bernstein and Vazirani [BV97].

Definition 17 (Unidirectional – Definition 3.14 in [BV97])

A QTM $M = (\Sigma, Q, \delta)$ is unidirectional if each state can be entered from only one direction. In other words, if $\delta(p_1, \sigma_1, \tau_1, q, d_1)$ and $\delta(p_2, \sigma_2, \tau_2, q, d_2)$ are both non-zero, then $d_1 = d_2$.

Note that in a unidirectional QTM, the direction component in any transition rule triple $|\tau\rangle |q\rangle |D\rangle$ is fully determined by the state $|q\rangle$. Thus we can recover the complete transition rules from the $|\sigma\rangle |q\rangle$ components alone, without the direction component. We call these the *reduced transition rules*, $\delta_r : Q \times \Sigma \mapsto \mathbb{C}^{Q \times \Sigma}$. (Equivalently, there exists an isometry $V : \mathbb{C}^Q \mapsto \mathbb{C}^{Q \times \{L, R, N\}}$ that maps the reduced transition rules to the original transition rules: $V\delta_r(p, \tau) = \delta(p, \tau)$.)

The following theorem gives necessary and sufficient conditions for a (partial) transition function to define a reversible Turing Machine.

Theorem 18 (Local well-formedness – Thm. B.1 and Cor. B.3 in [BV97])

A TM M is reversible iff its transition function satisfies the following conditions:

Unidirection

Each state of M can be entered while moving in only one direction. In other words, if $\delta(p_1, \sigma_1) = (\tau_1, q, d_1)$ and $\delta(p_2, \sigma_2) = (\tau_2, q, d_2)$ then $d_1 = d_2$.

One-to-one

The transition function is one-to-one when direction is ignored.

Furthermore, a partial transition function satisfying these conditions can always be completed to the transition function of a reversible TM.

Bernstein and Vazirani [BV97] proved that one can w.l.o.g. restrict to *unidirectional* QTMs. We therefore restrict the following quantum analogue of Theorem 18 to the unidirectional case:

Theorem 19 (Quantum local well-formedness)

A unidirectional QTM $M = (\Sigma, Q, \delta)$ is well-formed iff its quantum transition function δ satisfies the following conditions:

Normalisation

$$\forall p, \sigma \in Q \times \Sigma \quad \|\delta(p, \sigma)\| = 1. \quad (42a)$$

Orthogonality

$$\forall (p_1, \sigma_1) \neq (p_2, \sigma_2) \in Q \times \Sigma \quad \langle \delta_r(p_1, \sigma_1), \delta_r(p_2, \sigma_2) \rangle = 0. \quad (42b)$$

Furthermore, a partial quantum transition function satisfying these conditions can always be completed to a well-formed transition function.

Bernstein and Vazirani [BV97] proved this result for standard QTMs, that is, where the head must move left or right in each time step. (In fact, they prove it without the restriction to unidirectional QTMs; see Theorem 5.2.2 and Lemma 5.3.4 in [BV97].) We therefore extend the proof of Theorem 19 to generalised QTM, which is not difficult – indeed, it is made particularly straightforward by our restriction to unidirectional machines.

Proof (of Theorem 19) Let U be the time evolution operator of the QTM M . By definition M is well-formed iff U is an isometry, or equivalently iff the columns of U have unit length and are mutually orthogonal. Clearly, the normalisation condition specifies exactly that each column has unit length.

Pairs of configurations whose tapes differ in a cell not under either of the heads, or whose tape heads are more than two cells apart, cannot yield the same configuration in a single time step. Therefore, all such pairs of columns are guaranteed to be orthogonal. The orthogonality condition imposes orthogonality of pairs of columns for configurations that differ only in that one is in state p_1 reading σ_1 while the other is in state p_2 reading σ_2 .

It remains to consider pairs of configurations whose heads are one or two cells apart, differing at most in the symbol written in these cells and in their states. However, unidirectionality implies that any update triples that share the same state must share the same direction. Thus either the state or the head location necessarily differs for all such pairs of columns, hence these too are orthogonal.

The final claim follows straightforwardly from the fact that the normalisation and orthogonality conditions imply that a partial unidirectional reduced transition function δ_r is well-formed iff it defines an isometry on the space of states and tape symbols, and this can always be extended to a unitary. This fills in the undefined entries of δ_r by extending $\delta_r(q, \sigma)$ to an orthonormal basis for the space of states and symbols. \square

We will often only be interested in the behaviour of a QTM (or reversible TM) on a particular subset of inputs, since the machine will only be run on those. A *proper* machine is guaranteed to behave appropriately on some subset of inputs, but not necessarily on all possible inputs.

Definition 20 (Proper QTM) *A QTM behaves properly (or is proper) on a subset \mathcal{X} of initial superpositions if whenever initialised in $\phi \in \mathcal{X}$, the QTM halts in a final superposition where each configuration has the tape head in the start cell¹, the head never moved to the left of the starting cell, and the QTM never enters a configuration in which the head is in a superposition of different locations. We will refer to this latter property as deterministic head movement.*

Similarly, we say that a deterministic TM behaves properly (or is proper) on $\mathcal{X} \subset (\Sigma \setminus \{\#\})^$ if the head never moves to the left of the starting cell, and it halts on every input $x \in \mathcal{X}$ with its tape head back in the starting cell.*

When the set \mathcal{X} is clear from the context, we will omit it.

Note that, for a QTM to have deterministic head movement, it is *not* sufficient that none of its transition rules $\delta(\sigma, \tau)$ produce a superposition of head directions. The head can also end up in a superposition of different locations because the tape state is in a superposition, so that two transition rules with different deterministic head movement apply in superposition.

Behaving properly is not a severe restriction on classical Turing Machines.² In fact, given any TM, there is always an equivalent proper TM that computes the same function. One way to see this is to recall that all computable functions are computable by Turing Machines restricted to one-way infinite tapes (see any standard text book on the theory of computation, e.g. Kozen [Koz97]), and these clearly behave properly if the tape is extended to be two-way infinite. (Returning the head to the starting cell at the end of the computation poses no great difficulty.) In particular, this means that there exist proper universal Turing Machines.

Quantum Turing Machines were originally defined in Bernstein and Vazirani [BV97] to have two-way infinite tapes. Indeed, those authors point out that there are trivial well-formed machines (such as the always-move-right machine) whose evolution would not be unitary on a one-way infinite tape, since the starting configuration would have no predecessor. However, when we come to encode our QTMs into local Hamiltonians, we will only be able to simulate tapes with a boundary.³ To avoid technical issues, we follow Bernstein and Vazirani [BV97] in defining QTMs on two-way infinite tapes, but we will ensure that none of the QTMs

¹This property is called *stationarity* in [BV97]

²Nor is it a severe restriction on QTMs, but we will not need or prove this here.

³Effectively, we can only encode quantum bounded linear automata rather than QTMs, though there is no limit on the finite tape length we can encode.

(or reversible TMs) that we construct ever move their head before the starting cell. Thus, when encoding the QTM in a Hamiltonian, we can ignore all of the tape to the left of the starting cell.

In fact, the local Hamiltonians encoding the QTMs will only be able to simulate the evolution on a finite (but arbitrarily large) section of tape. We will therefore be interested in keeping very tight control on the space requirements of all the reversible and quantum TMs that we construct. By carefully controlling the space overhead, it will then be sufficient for our purposes to simulate the evolution of the QTM on a finite portion of tape that is essentially no longer than the input.

6.1.1 Turing Machine Programming

For a multi-track Turing machine with k tracks and alphabet Σ_i on track i , we will denote the contents of a tape cell by a tuple of symbols $[\sigma_1, \sigma_2, \dot{\sigma}, \sigma_k] \in \Sigma_1 \times \Sigma_2 \times \cdots \times \Sigma_k$ specifying the symbol written on each track. Similarly, the configuration of the tape will be specified by a tuple $[c_1, c_2, \dots, c_n] \in \Sigma_1^* \times \Sigma_2^* \times \cdots \times \Sigma_k^*$, where by convention all c_i are aligned to start in the same cell (which will be the starting cell unless otherwise specified). We will use \cdot to stand for an arbitrary track symbol. We will often describe Turing machines that act only on a subset of tracks, and leave the contents of all other tracks alone. In this case, we will only write out the states of the acted-upon tracks in the transition function; this transition function should be understood to be extended to the remaining tracks in the obvious way.

As a shorthand, a transition rule on a tuple containing a \cdot on one or more tracks should be understood to stand for the set of transition rules that leave the tracks marked \cdot unchanged, and act as indicated on the remaining tracks. We will often only specify some of the elements of a transition function when defining a reversible or quantum TM. We will call such partial transition functions “well-formed” if the elements that *are* defined satisfy the conditions of Theorem 18 or Theorem 19, since by those theorems this is sufficient to guarantee that the partial transition functions can be completed to well-formed transition functions. QTM (or reversible TM) will sometimes use a finite number of auxiliary tracks. These are assumed always to start and finish in the all blank configuration.

We will have frequent recourse to the following basic TM and QTM programming primitives from Bernstein and Vazirani [BV97], which we slightly generalise here to account for additional properties of the resulting QTMs that will be important to us later. These primitives allow different QTMs to be combined in various ways to build up a more complex QTM.

Lemma 21 (Subroutine Lemma)

Let M_1 be a two-track, normal form, reversible TM and M_2 a two-track normal form reversible TM (or well-formed, normal form, unidirectional QTM) with the same alphabet and the following properties:

- (i). M_1 is proper on initial configurations in \mathcal{X}_1 and M_2 is proper on \mathcal{X}_2 .
- (ii). When started on \mathcal{X}_1 , M_1 leaves the second track untouched and when started on \mathcal{X}_2 , M_2 leaves the first track untouched.
- (iii). There is a state q on M_1 that, when started on \mathcal{X}_1 , can only be entered with the head in the starting cell.
- (iv). \mathcal{X}_2 contains all the output superpositions (with q_f replaced by q_0) of k consecutive executions of M_2 started from an initial configuration in \mathcal{X}_1 for all $0 \leq k \leq r$, where $r \in \mathbb{N} \cup \{\infty\}$ is the maximum number of times that q is entered when M_1 runs on input in \mathcal{X}_1 .

Then there is a normal form, reversible TM (or well-formed, normal form, unidirectional QTM) M which behaves properly on \mathcal{X}_1 and acts as M_1 except that each time it would enter state q , it instead runs machine M_2 .

Proof Exactly as Lemma 4.8 in [BV97]. □

Lemma 22 (Dovetailing Lemma) *Let M_1 and M_2 be well-formed normal form unidirectional QTMs (resp. normal form reversible TMs) with the same alphabet, so that M_1 is proper on \mathcal{X}_1 , M_2 is proper on \mathcal{X}_2 and \mathcal{X}_2 contains all final superpositions (resp. configurations) of M_1 started on \mathcal{X}_1 . Then, there is a well-formed normal form unidirectional QTM (resp. normal form reversible TM) M which carries out the computation of M_1 followed by the computation of M_2 and that is also proper on \mathcal{X}_1 .*

Proof Exactly as Lemma 4.9 in [BV97]. □

Lemma 23 (Reversal Lemma) *If M is a well-formed, normal form, unidirectional QTM (resp. normal form reversible TM) then there is a well formed, normal form, unidirectional QTM (resp. normal form reversible TM) M^\dagger that reverses the computation of M while taking two extra time steps and using the same amount of space. Moreover, if M is proper on \mathcal{X} , then M^\dagger is proper on the set of final superpositions (resp. configurations) of M started on \mathcal{X} .*

Proof The proof is very similar to that of Lemma 4.12 in [BV97]. We will prove it for QTMs, since reversible TMs are a special case of these. Consider an initial superposition $|\phi_0\rangle \in \mathcal{X}$, and let $|\phi_1\rangle, \dots, |\phi_n\rangle$ be the evolved sequence obtained by $|\phi_{i+1}\rangle = U_M |\phi_i\rangle$, where $|\phi_n\rangle$ is a final superposition. Since M is normal form, $|\phi_1\rangle, \dots, |\phi_n\rangle$ do not have support on q_0 .

Let $|\phi'_n\rangle$ be the superposition obtained by replacing the state q_f in $|\phi_n\rangle$ with the initial state of the new machine q'_0 . Let $|\phi'_0\rangle$ be the superposition obtained by

replacing the state q_f in $|\phi_0\rangle$ with the new final state q'_f . We want to construct a QTM M^\dagger that, when started from $|\phi'_n\rangle$, halts on $|\phi'_0\rangle$ in $n + 2$ steps. M^\dagger will have the same alphabet and set of states as M , together with the new initial and final states q'_0, q'_f . We define the transition function δ' in the following way:

(i). $\delta'(q'_0, \sigma) = |\sigma\rangle |q_f\rangle |\bar{d}_{q_f}\rangle$.

(ii) For each $q \in Q \setminus \{q_0\}$ and each $\tau \in \Sigma$,

$$\delta'(q, \tau) = \sum_{p,q} \delta(p, \sigma, \tau, q, d_q)^* |\sigma\rangle |p\rangle |\bar{d}_p\rangle. \quad (43)$$

(iii) $\delta'(q_0, \sigma) = |\sigma\rangle |q'_f\rangle |d_{q_0}\rangle$.

(iv) $\delta'(q'_f, \sigma) = |\sigma\rangle |q'_0\rangle |N\rangle$.

Here, for any state q , d_q is the unique direction in which that state can be entered, and \bar{d}_q is the opposite direction. Since M is unidirectional, Theorem 19 implies that M^\dagger is a well-formed, normal form, unidirectional QTM. Given a configuration c in state q , $\pi(c)$ is defined as the configuration derived from c by moving the head one step in the direction \bar{d}_q . π can be extended by linearity to \mathcal{S} .

Let us now analyse the behaviour of M^\dagger started from $|\phi'_n\rangle$. By (i), M^\dagger maps $|\phi'_n\rangle$ in one step to $\pi(|\phi_n\rangle)$. Now consider (i). Since M is normal form, it maps superposition $|\psi_1\rangle$ to superposition $|\psi_2\rangle$ with no support on state q_0 if and only if $|\psi_1\rangle$ has no support on state q_f . Denote $Q_0 = Q \setminus \{q_0\}$, $Q_f = Q \setminus \{q_f\}$. If M takes a configuration c_1 with a state from Q_f with amplitude α to a configuration c_2 (necessarily with a state in Q_0), then (i) ensures that M^\dagger takes configuration $\pi(c_2)$ to configuration $\pi(c_1)$ with amplitude α^* . Let S_0 (resp. S_f) be the space of superpositions using only states in Q_0 (resp. Q_f). Since M is well-formed, the restriction of the evolution operator U_M to S_f is an isometry into S_0 . Hence M^\dagger implements, up to conjugation by π , the inverse of U_M restricted to $U_M(S_f)$.

As an aside, note that this implies that the evolution operator of M is indeed a unitary and not just an isometry. Indeed, by (i) to (i) above, the evolution operator of M^\dagger is also an isometry from S_0 into S_f . Since π is trivially a unitary on \mathcal{S} , this implies that U_M restricted to S_f is a unitary *onto* S_0 . Now, since M is normal form, U_M achieves any possible configuration with state q_0 by starting from the same configuration but with q_0 replaced by q_f . These two facts together show that U_M is indeed surjective and hence a unitary.

Returning to the proof of the lemma, we have seen how (i) implies that M^\dagger executes the sequence of n steps $\pi(|\phi_{n-1}\rangle), \dots, \pi(|\phi_0\rangle)$. Finally, by (i), in the $(n+2)$ 'th step, M^\dagger maps $\pi(|\phi_0\rangle)$ to $|\phi'_0\rangle$ as desired.

Moreover, it is trivial to see that M^\dagger uses exactly the same space as M , and behaves properly. \square

Note that our way of defining normal form (Definition 16 as opposed to that in [BV97]) is crucial to guarantee that the reversal machine is proper. As a corollary of the proof, we have also shown that the evolution of a well-formed, normal form unidirectional QTM is a unitary operator. (The proof of this fact for non-unidirectional QTMs is much more involved, and can be found in [BV97]).

6.1.2 Reversible Turing Machine Toolbox

It will be helpful to have a toolbox of reversible TMs which carry out various elementary computations. In order to satisfy the space constraints of Theorem 11, we will be interested in keeping tight control of the space requirements of all our constructions.

Lemma 24 (Copying machine) *There is a two-track, normal form, reversible TM COPY with alphabet $\Sigma \times \Sigma$ that, on input s written on the first track, behaves properly, copies the input to the second track and runs for time $2|s| + 1$, using $|s| + 1$ space.*

Proof We simply step the head right, copying the symbol from the first track to the second. However, we defer copying the starting cell until the end of the computation, so that we can locate the starting cell again. It is straightforward to verify that the following normal form transition function implements COPY:

$$\begin{array}{c|ccc}
 & [\sigma, \#] & [\#, \#] & [\tau, \tau] \\
 \hline
 q_0 & ([\sigma, \#], q_1, R) & ([\#, \#], q_f, N) & \\
 q_1 & ([\sigma, \sigma], q_1, R) & ([\#, \#], q_2, L) & \\
 q_2 & ([\sigma, \sigma], q_f, N) & & ([\tau, \tau], q_2, L) \\
 q_f & ([\sigma, \#], q_0, N) & ([\#, \#], q_0, N) & ([\tau, \tau], q_0, N)
 \end{array} \tag{44}$$

$\forall \sigma, \tau \in \Sigma / \#$

This partial transition function verifies the two conditions of Theorem 18, so it is well-formed and can be completed to give a reversible TM. \square

Lemma 25 (Shift-right machine) *There exists a normal form, reversible TM SHIFT with alphabet Σ that, on input s behaves properly, shifts s one cell to the right and runs for time $2|s| + 2$, using space $|s| + 2$.*

Proof It is straightforward to verify that the following normal-form transition

function implements `SHIFT`:

$$\begin{array}{c|cc}
 & \# & \sigma \\
 \hline
 q_0 & & (\#, q_1^\sigma, R) \\
 q_1^\nu & (\nu, q_2, R) & (\nu, q_1^\sigma, R) \\
 q_2 & (\#, q_3, L) & \\
 q_3 & (\#, q_f, N) & (\sigma, q_3, L) \\
 q_f & (\#, q_0, N) & (\sigma, q_0, N)
 \end{array} \tag{45}$$

$\forall \sigma, \nu \in \Sigma/\#$

As this partial transition function verifies the two conditions of Theorem 18, it is well-formed and can be completed to give a reversible TM. \square

Lemma 26 (Equality machine) *There is a three-track, normal form, reversible TM `EQL` with alphabet $\Sigma \times \Sigma \times \{\#, 0, 1\}$ that, on input $s; t; b$, where s and t are arbitrary input strings and b is a single bit, behaves properly and outputs $s; t; b'$, where $b' = \neg b$ if $s = t$ and $b' = b$ otherwise. Furthermore, `EQL` runs for time $2|s| + 1$ and uses $|s| + 1$ space.*

Proof Implementing a non-reversible equality-testing machine is trivial: just scan the head right checking if the symbols on the two input tracks match. If we encounter a non-matching pair, return to the starting cell and flip the output bit. If we reach the end of the input without encountering a non-matching pair, return to the starting cell and leave the output bit unchanged.

Doing this reversibly requires more care. The problem is that the computation splits into two possible paths, depending on whether a non-match was encountered or not, and we must merge these two divergent computations back together again *reversibly*. For example, we cannot simply halt after either flipping the output bit or leaving it unchanged, as that would give multiple transitions into the final state that write the same symbol to the starting cell.

The trick is to return to the point at which the computational paths diverged (either the first non-matching pair of symbols, or the end of the input) after setting the output bit, in order to merge the two computational paths back together, before returning to the starting cell again to halt.

Because we want to avoid ever moving the head before the starting cell, there is also the issue of how to identify the starting cell so that we can return to it again. If the symbols in the first cell differ, we can set the output bit immediately and halt. If they are identical, we temporarily change the symbol on the second input track to a $\#$ to mark the starting cell. We can recover the original second-track input at the end of the computation by copying over the symbol from the first track.

The following normal-form transition function carries out this procedure:

	$[\sigma, \sigma, b]$	$[\sigma, \tau, b]$	$[\sigma, \#, b]$	$[\#, \tau, b]$
q_0	$([\sigma, \#, b], q_1, R)$	$([\sigma, \tau, b], q_f, N)$	$([\sigma, \#, b], q_f, N)$	$([\#, \tau, b], q_f, N)$
q_1				
q_2			$([\sigma, \#, -b], q_3, R)$	
q'_2			$([\sigma, \#, b], q'_3, R)$	
q_3				
q'_3				
q_4			$([\sigma, \sigma, b], q_f, N)$	
q_f	$([\sigma, \sigma, b], q_0, N)$	$([\sigma, \tau, b], q_0, N)$	$([\sigma, \#, b], q_0, N)$	$([\#, \tau, b], q_0, N)$
	$[\sigma, \sigma, \#]$	$[\sigma, \tau, \#]$	$[\sigma, \#, \#]$	$[\#, \tau, \#]$
q_1	$([\sigma, \sigma, \#], q_1, R)$	$([\sigma, \tau, \#], q'_2, L)$	$([\sigma, \#, \#], q'_2, L)$	$([\#, \tau, \#], q'_2, L)$
q_2	$([\sigma, \sigma, \#], q_2, L)$			
q'_2	$([\sigma, \sigma, \#], q'_2, L)$			
q_3	$([\sigma, \sigma, \#], q_3, R)$			
q'_3	$([\sigma, \sigma, \#], q'_3, R)$	$([\sigma, \tau, \#], q_4, L)$	$([\sigma, \#, \#], q_4, L)$	$([\#, \tau, \#], q_4, L)$
q_4	$([\sigma, \sigma, \#], q_4, L)$			
q_f	$([\sigma, \sigma, \#], q_0, N)$	$([\sigma, \tau, \#], q_0, N)$	$([\sigma, \#, \#], q_0, N)$	$([\#, \tau, \#], q_0, N)$
	$[\#, \#, \#]$			
q_1	$([\#, \#, \#], q_2, L)$			
q_2				
q'_2				
q_3	$([\#, \#, \#], q_4, L)$			
q'_3				
q_4				
q_f	$([\#, \#, \#], q_0, N)$			

$$\forall \sigma \neq \tau \in \Sigma/\#, \forall b \in \{0, 1\}$$

(46)

One can verify that this partial transition function satisfies the two conditions of Theorem 18, so it is well-formed and can be completed to give a reversible TM, as required. \square

It is somewhat easier to construct reversible implementations of basic arithmetic operations if the numbers are written on the tape in little-endian order (i.e. least-significant bit first), as it avoids any need to shift the entire number to the right to accommodate additional digits. We adopt this convention for all the following basic arithmetic machines. We do not allow numbers to be padded with leading 0's as that would allow multiple binary representations of the same number, which

is inconvenient when constructing reversible machines. Note that this means the number zero is represented by the blank string, not the string “0”.

Lemma 27 (Increment and decrement machines) *There exist normal form, reversible TMs INC and DEC with alphabet $\{\#, 0, 1\}$ that, on little-endian binary input n (with $n > 1$ for DEC), behave properly and output $n + 1$ or $n - 1$ respectively. Both machines run for time $O(\log n)$ and uses at most $|n| + 2$ space.*

Incrementing a binary number on a non-reversible TM is straightforward: simply step the head along the number starting from the least significant bit, flipping 1’s to 0’s to propagate the carry until the first 0 or #, then flip that 0 or # to 1 and halt. (Little-endian order avoids any need to shift the whole input to the right to accommodate an additional digit, should one be required.) However, making this procedure reversible is more fiddly. One option of course is to use Bennett’s [Ben73] general reversible simulation and uncomputation of a non-reversible TM, but this comes at a cost of polynomial space overhead. A more careful construction allows us to implement INC directly, using just two additional tape cells.

Proof (of Lemma 27) We first use the SHIFT machine from Lemma 25 to shift the entire input one cell to the right, as a convenient way of getting a # in the starting cell so that we can return to it later. We then reversibly increment the binary number written on the tape, and finish by running SHIFT^\dagger (the reversal of the SHIFT machine, constructed using Lemma 23) to shift the output back one cell to the left. The following normal-form transition function implements the increment part of this procedure:

	#	0	1	
q_0	$(\#, q_1, R)$			
q_1	$(1, q'_2, R)$	$(1, q_2, R)$	$(0, q_1, R)$	
q_2		$(0, q_3, L)$	$(1, q_3, L)$	
q'_2	$(\#, q_3, L)$			(47)
q_3			$(1, q_4, L)$	
q_4	$(\#, q_f, N)$	$(0, q_4, L)$		
q_f	$(\#, q_0, N)$	$(0, q_0, N)$	$(1, q_0, N)$	

This partial transition function verifies the two conditions of Theorem 18, so it is well-formed and can be completed to give a reversible TM.

This completes the construction of INC. To implement DEC, simply construct the reversal of INC using Lemma 23. □

We can use the INC construction to construct a looping primitive. (This gives a slightly more general version of the Looping Lemma from [BV97, Lemma 4.13], which also has tighter control on the space requirements.)

Lemma 28 (Looping Lemma) *There is a two-track, normal form, reversible TM LOOP2 with alphabet $\{\#, 0, 1\}$, which has the following properties. On input $n; m$, with $n < m$ both little-endian binary numbers, LOOP2 behaves properly, runs for time $O((m - n) \log m)$, uses space $|m| + 2$ and halts with its tape unchanged. Moreover, LOOP2 has a special state q such that on input $n; m$, it visits state q exactly $m - n$ times, each time with its tape head back in the start cell.*

There is also a one-track, normal form, proper, reversible TM LOOP which, on input $m \geq 1$, behaves as LOOP2 on input $0; m$.

Proof Our construction closely follows the proof of Bernstein and Vazirani [BV97, Lemma 4.2.10]. We will use two auxiliary tracks in addition to the two input tracks, both with alphabet $\{\#, 0, 1\}$ and both initially blank.

The core of LOOP2 machine is a reversible TM M' constructed out of two proper, normal form, reversible TMs M_1 and M_2 . M_1 has initial and final states q_0, q_f , and transforms input $n; m; x; b$ into $n; m; x + 1; b'$, where b is a bit and $b' = \neg b$ if $x = n$ or $x = m - 1$ but not both, otherwise $b' = b$. M_2 has new initial and final states q_α, q_ω , with q_0 and q_f as its only other normal (non-initial or final) internal states. M_2 behaves as follows:

- (i). If it is in state q_α with $b = 0$, it flips b to 1 and enters state q_0 .
- (ii). If it is in state q_f with $b = 0$, it enters state q_0 .
- (iii). If it is in state q_f with $b = 1$, it flips b to 0 and halts.

The transition rules for M' are constructed by deleting the q_f to q_0 transition rules from M_1 , and adding all the remaining M_1 rules to those of M_2 . The initial and final states of M' are q_α, q_ω . On input $n; m; n; 0$, M' will therefore run M_2 until it enters the q_0 state, then run M_1 until it halts and re-enters M_2 . It will continue to alternate in this way between M_2 and M_1 until the former halts. Thus M' goes through the following sequence of configurations:

$$\begin{aligned} n; m; n; 0 &\xrightarrow{(i)} n; m; n; 1 \xrightarrow{M_1} n; m; n+1; 0 \xrightarrow{(ii)} n; m; n+1; 0 \xrightarrow{M_1} n; m; n+2; 0 \xrightarrow{(ii)} \dots \\ &\dots \xrightarrow{(ii)} n; m; m-1; 0 \xrightarrow{M_1} n; m; m; 1 \xrightarrow{(iii)} n; m; m; 0 \quad (48) \end{aligned}$$

It therefore runs exactly $m - n$ times and enters the state q_f once in each run, so we can take q_f as the special state q .

To initialise the tracks, we dovetail a proper, reversible TM before M' which transforms the input $n; m; \#; \#$ into the configuration $n; m; n; 0$. This is easily constructed by dovetailing the COPY machine from Lemma 24 (acting on the first and third tracks) with a simple reversible TM which changes the first cell of the fourth track from $\#$ to 0 and halts (implementing the latter is trivial). To return

all tracks to their initial configuration at the end, we dovetail another proper, reversible TM after M' which transforms the configuration $n; m; m; 0$ into the final output $n; m; \#; \#$. This is easily constructed by dovetailing the reversal COPY^\dagger of the copying machine from Lemma 24 (acting on the second and third tracks) with a simple reversible TM which changes the first cell of the fourth track from 0 back to $\#$. From Lemma 29, these initialisation and reset machines run for time $O(\log n)$ and $O((m - n) \log m)$, respectively, and use $|n| + 1$ and $|m| + 1$ space.

It remains to construct M_1 and M_2 , and show that combining them as described above to form M' gives a proper, normal form, reversible TM. M_1 is constructed by dovetailing together the EQL machine from Lemma 26 (with the first and third tracks as its input tracks and the fourth track as its output), then the INC machine from Lemma 27 (acting only on the third track), and finally another EQL machine (this time with the second and third tracks as its input tracks and the fourth track as its output). By Lemma 22 and the fact that all the constituent machines are proper, normal form and reversible, M_1 is proper, normal form and reversible. From Lemmas 26 and 27, M_1 runs for time $O(\log m)$ and takes at most $|m| + 2$ space.

The following normal form, partial transition function is essentially the same as the corresponding construction in Lemma 4.2.6 of Bernstein and Vazirani [BV97], but we can simplify slightly by exploiting the fact that we are allowing generalised TMs. It acts only on the fourth track, and implements a machine M_2 that clearly satisfies the properties (i) to (iii), above:

$$\begin{array}{c|ccc}
 & \# & 0 & 1 \\
 \hline
 q_\alpha & & (1, q_0, N) & \\
 q_f & & (0, q_0, N) & (0, q_\omega, N) \\
 q_\omega & (\#, q_\alpha, N) & (0, q_\alpha, N) & (1, q_\alpha, N)
 \end{array} \tag{49}$$

M_2 is clearly normal form, and satisfies the two conditions of Theorem 18.

Since M_1 is normal form, it has no transitions into q_0 or out of q_f other than the ones we deleted before constructing M' . All remaining internal states of M_1 and M_2 are distinct. Thus, since M_1 is reversible, the transition rules for M' also satisfy the two conditions of Theorem 18. M' can therefore be completed to a normal form, reversible TM. Since M_1 is proper, it is easy to see that M' is also proper.

M' loops on M_1 $m - n$ times (with constant time overhead and no space overhead coming from M_2). Each run of M_1 takes time $O(\log m)$, so the complete Loop2 implementation takes time $O(\log m)$. None of the constituent TMs use more than $|m| + 2$ space, so neither does Loop2 . This completes the construction of Loop2 .

To construct Loop , we simply remove the second input track from the Loop2 machine, and replace the EQL machine that acts on that track with a trivial machine that checks whether the starting cell of the third track contains the $\#$ symbol. \square

The reversible incrementing, decrementing, and looping machines constructed above are sufficient to implement all arithmetic operations. The only ones we will need are addition and subtraction. These are now easy to construct.

Lemma 29 (Binary adder) *There exist two-track, normal form, reversible TMs ADD and SUB with alphabet $\{\#, 0, 1\}^{\times 3}$, which have the following properties. On input $n; m$ with n and m both little-endian binary numbers and $m \geq 1$, ADD behaves properly and outputs $n + m; m$. On input $n; m$ with $n > m$, SUB behaves properly and outputs $n - m; m$. Both TMs run for time $O(m \log m \log n)$ and use $\max(|n|, |m|) + 2$ space.*

Proof To construct ADD , we simply run the LOOP TM of Lemma 28 on the second track, inserting for its special state the INC machine of Lemma 27 acting on the first track. Since both LOOP and INC are proper and normal form, the resulting machine is also proper and normal form. SUB is simply the reversal of ADD , constructed using Lemma 23. \square

We will also need a reversible TM to convert the input from unary representation to binary. Again, it is not difficult to construct this using our reversible incrementing TM.

Lemma 30 (Unary to binary converter) *There exists a two-track, normal form, reversible TM UTOB with alphabet $\{\#, 1\} \times \{\#, 0, 1\}$ with the following properties. On input $1n; \#$ (n written in unary on the first track), UTOB behaves properly and outputs $1n; n$ (n written in little-endian binary on the second track). Furthermore, UTOB runs for $O(n^2 \log n)$ steps and uses $n + 1$ space.*

Proof The basic idea is to step the head right along the unary track until we reach the end of the unary input string, running the INC machine of Lemma 27 once each time we step right. However, INC needs to be started with its head in the starting cell. So each time we run it, we need to temporarily mark the current head location, move the head back to the starting cell in order to run INC , then return the head to its previous location and continue stepping right.

Consider the following normal form transition function, acting only on the unary track:

	#	1	
q_0	$(\#, q_f, N)$	$(1, q, N)$	
q_1	$(\#, q_4, L)$	$(\#, q_2, L)$	
q_2	$(\#, q, N)$	$(1, q_2, L)$	
q	$(\#, q'_2, R)$	$(\#, q_1, R)$	
q'_2	$(1, q_1, R)$	$(1, q'_2, R)$	
q_4	$(1, q_f, N)$	$(1, q_4, L)$	
q_f	$(\#, q_0, N)$	$(1, q_0, N)$	(50)

This verifies the two conditions of Theorem 18, so it is well-formed and defines a reversible TM. It is also easy to see that this TM is proper.

The transition function implements a machine that steps right over the 1's on the unary track until it reaches the end of the input, at which point it moves its head back to the starting cell and halts. However, each time it steps right, it temporarily marks the current head location by overwriting the 1 on the unary track with a #, returns the head to the starting cell, enters an apparently useless internal state q for one time step, then returns the head back to where it was before, restores the 1 on the unary track, and continues stepping right from where it left off. The purpose of this apparently pointless procedure is that, by Lemma 21, we can substitute the INC machine from Lemma 27 for the state q , with INC acting on the (initially blank) binary track.

The above TM enters the q state precisely n times. Thus, when we substitute INC for q , INC will be run precisely n times, leaving the number n written on the binary track as required. Each iteration of INC takes time $O(\log n)$ and at most $|n| + 2$ space. The step-right TM constructed above adds time overhead $O(n^2)$ and uses $n + 1$ space. Neither machine ever moves its head before the starting cell. Thus the overall machine satisfies the time and space claims of the Lemma. \square

6.2 Quantum phase estimation overview

The qubits used in the quantum phase estimation circuit can be divided into two sets: the “output” qubits which will ultimately contain the binary decimal expansion of the phase in the black-box unitary, and the “ancilla” qubits on which the black-box unitary is applied. In our case, the black-box unitary will be the single qubit unitary $U_\varphi = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi\varphi} \end{pmatrix}$, so the ancilla set will be a single qubit. As stated above, φ will refer to the number whose binary decimal expansion contains the digits of n in reverse order after the decimal.

The quantum phase estimation algorithm proceeds in five stages: (1) A preparation stage, in which the qubits are first initialised in the $|0\rangle$ state, then Hadamard gates are applied to all the output qubits and the ancilla qubit is prepared in an eigenstate of U_φ (in our case $|1\rangle$); (2) The controlled-unitary stage, during which control- U_φ operations are applied between the output qubits and the ancilla qubit; (3) A stage in which we locate the least significant bit of the output; (4) The quantum Fourier transform stage, in which the inverse quantum Fourier transform is applied to the output qubits; (5) A reset stage, which resets all the auxiliary systems used during the computation to their initial configuration.

We will construct QTMs for each of these stages separately, and use the Dovetailing Lemma 22 to chain them together. It will be useful to divide the tape into multiple tracks. A “quantum track” with alphabet $\Sigma_q = \{\#, 0, 1\}$ will store the qubits involved in the phase estimation algorithm. The input $1N$ will be supplied

as a string of N 1's on the quantum track. All the other tracks will essentially be classical; after each stage they will be left in a single standard basis state (i.e. neither entangled with other tracks, nor in a superposition of basis states). These classical tracks will be used to implement the classical processing needed to control the quantum operations applied to the quantum track.

6.3 Preparation stage

We will use the first cell of the quantum track as the ancilla qubit, and the following N cells as the output qubits of the quantum phase estimation algorithm.

It will be convenient to first store a copy of the input given in unary on the quantum track on an auxiliary “input track”, but written in binary, i.e. to transform the configuration $1N; \#$ of the quantum and input tracks into $1N; N$. Running the unary-to-binary converter U_{toB} from Lemma 30 on the quantum and input tracks carries out the desired transformation.

We want to initialise the $N + 1$ qubits that will be used in the phase estimation, by preparing the ancilla qubit in the state $|1\rangle$ and the output qubits in the $|+\rangle$ state. The first N qubits are already in the $|1\rangle$ state thanks to the input string. We therefore prepare the desired state by first temporarily flipping the state of the first qubit on the quantum track $|\#\rangle$ to mark the starting cell, then stepping the QTM head right rotating each qubit into the $|+\rangle$ state, until we reach the first $|\#\rangle$ state which we again rotate to $|+\rangle$ to initialise the $N + 1$ 'th qubit. We then move the head back to the start location, flip the state of the first cell of the quantum track from $|\#\rangle$ to $|1\rangle$, and halt. (The contents of all other tracks are ignored.) It is straightforward to verify that the following partial transition function satisfies the conditions of Theorem 19, so is well-formed, and implements a proper, well-formed, normal form, unidirectional QTM that does exactly what we want.

	#	0	1	
q_0			$ \#\rangle q_1\rangle R\rangle$	
q_1	$\frac{1}{\sqrt{2}}(0\rangle + 1\rangle)$	$ q_2\rangle R\rangle$	$\frac{1}{\sqrt{2}}(0\rangle + 1\rangle)$	$ q_1\rangle R\rangle$
q_2	$ \#\rangle$	$ q_3\rangle L\rangle$		
q_3	$ 1\rangle$	$ q_f\rangle N\rangle$	$ 0\rangle$	$ q_3\rangle L\rangle$
q_f	$ \#\rangle$	$ q_0\rangle N\rangle$	$ 0\rangle$	$ q_0\rangle N\rangle$

(51)

6.4 Control- U_φ stage

The second stage of the phase estimation procedure is to apply control- $U_\varphi^{2^{n-1}}$ operations between the n 'th output qubit and the ancilla qubit (see Figure 6). Constructing a QTM that implements this is more complex. The basic idea is to supply the QTM with an internal state q which causes the U_φ rotation to be applied to the quantum track at the current head location. We then construct classical

control machinery which iterates over the output qubits, and loops on q to apply U_φ for a total of 2^{n-1} times.




Figure 6: The first stage of the quantum phase estimation circuit for φ (cf. Fig. 5.2 in [NC00]).

We will use three auxiliary tracks. One, with alphabet $\Sigma_m = \{t, c, m_0, m_1, \#\}$, will be used to mark the position of the current control and target qubits. (The $m_{0,1}$ states will be used to temporarily store an auxiliary qubit on the mark track of the starting cell.) The other two auxiliary tracks, with alphabet $\Sigma_l = \{0, 1, \#\}$, will constitute the work tapes of two different reversible looping TMs from Lemma 28.

6.4.1 cU^k machine

We will need a QTM which applies the control- U_φ operation k times. We give a construction for an arbitrary controlled single-qubit unitary U , as this will be useful later.

Lemma 31 (Controlled- U QTM) *For any single-qubit unitary U , there exists a well-formed normal form unidirectional QTM cU^k with the following properties. The QTM has three tracks: a looping track with alphabet $\Sigma_l = \{0, 1, \#\}$, a mark track with alphabet $\Sigma_m = \{t, c, m_0, m_1, \#\}$, and a quantum track with alphabet $\Sigma_q = \{0, 1, \#\}$. The input consists of a number $k \geq 1$ written in binary on the looping track in little-endian order, a configuration containing a single t and a single c within the first n tape cells on the mark track (and all other cells blank), and an arbitrary n -qubit state on the quantum track.*

On such input, the QTM applies the control- U operation k times between the control and target qubits on the quantum track marked by c and t , and then halts, having run for time $O(kn + k \log k)$, used at most $\max(n, |k|) + 2$ space, behaving properly and leaving the configurations of the looping and mark tracks unchanged.

Proof It will be convenient to first run the `SHIFT` machine from Lemma 25 acting only on the mark and quantum tracks, to shift this part of the input one cell to the right. This produces a `#` on the mark and quantum tracks of the start cell, which we can use later to return to this cell. The remainder of the construction returns the quantum and mark tracks of the start cell to the `#` state, so we can run the corresponding `SHIFT`[†] machine at the very end (where `SHIFT`[†] is the reversal of M_1 constructed using Lemma 23) to shift these tracks back one cell to the left, so that the final output is correctly aligned. These shift operations take $O(n)$ time and use $n + 2$ space.

We construct the core of the cU^k QTM out of two simpler machines, M_1 and M_2 , dovetailed together in the sequence M_1, M_2, M_1^\dagger (where M_1^\dagger is the reversal of M_1). Machine M_1 scans right until it encounters a c on the mark track, at which point it applies a CNOT between the quantum track and an internal qubit, returns the head to the starting cell, applies a CNOT between the internal qubit and an auxiliary qubit on the mark track, and halts. Clearly, M_1 is proper, runs for time at most $O(n)$ and uses at most $n + 1$ space.

The following well-formed, normal form, partial quantum transition function implements M_1 (which acts only on the mark and quantum tracks). Note that one of the internal states of M_1 comes in two varieties, indicated by a superscript $j = \{0, 1\}$, which are used to temporarily store a qubit in the internal state of the machine:

	[#, #]	[#, i]	[t , \cdot]	[c , j]	[m_k , \cdot]
q_0	$ m_0, \# \rangle q_1 \rangle R \rangle$				
q_1		$ \#, i \rangle q_1 \rangle R \rangle$	$ t, \cdot \rangle q_1 \rangle R \rangle$	$ c, j \rangle q_2^j \rangle L \rangle$	
q_2^j		$ \#, i \rangle q_2^j \rangle L \rangle$	$ t, \cdot \rangle q_2^j \rangle L \rangle$		$ m_{k \oplus j} \rangle q_f \rangle N \rangle$
q_f	$ \#, \# \rangle q_0 \rangle N \rangle$	$ \#, i \rangle q_0 \rangle N \rangle$	$ t, \cdot \rangle q_0 \rangle N \rangle$	$ c, j \rangle q_0 \rangle N \rangle$	$ m_k, \cdot \rangle q_0 \rangle N \rangle$

$\forall i, j, k \in \{0, 1\}$

(52)

Machine M_2 loops k times, where k is specified by the number written in binary on the looping track, applying a control- U operation between the auxiliary qubit stored on the mark track of the start cell and the target qubit, and then halts. We construct M_2 using the reversible looping TM `LOOP` from Lemma 28, and inserting for its special state a QTM M' that is very similar to M_1 . The M' machine first applies a CNOT between the auxiliary qubit stored in the mark track of the start cell and an internal qubit. It then scans right until it finds the t , and applies a single control- U operation between the internal qubit and the target qubit. Finally, it moves the head back to the starting cell, and applies another CNOT between the internal qubit and the auxiliary qubit, before halting.

If u_{ij} denotes the i, j 'th matrix element of U , then the following well-formed, normal form, partial quantum transition function (which acts only on the mark and

quantum tracks) implements M' :

	$[m_0, \cdot]$	$[m_1, \cdot]$	$[\#, \cdot]$	$[t, j]$	$[c, \cdot]$
q_0	$ m_0, \cdot\rangle q_1^0\rangle R\rangle$	$ m_1, \cdot\rangle q_1^1\rangle R\rangle$			
q_1^0			$[\#, \cdot] q_1^0\rangle R\rangle$	$ t, j\rangle q_2^0\rangle L\rangle$	$[c, \cdot] q_1^0\rangle R\rangle$
q_1^1			$[\#, \cdot] q_1^1\rangle R\rangle$	$\sum_i u_{ij} t, i\rangle q_2^1\rangle L\rangle$	$[c, \cdot] q_1^1\rangle R\rangle$
q_2^0	$ m_0, \cdot\rangle q_f\rangle N\rangle$		$[\#, \cdot] q_2^0\rangle L\rangle$		$[c, \cdot] q_2^0\rangle L\rangle$
q_2^1		$ m_1, \cdot\rangle q_f\rangle N\rangle$	$[\#, \cdot] q_2^1\rangle L\rangle$		$[c, \cdot] q_2^1\rangle L\rangle$
q_f	$ m_0, \cdot\rangle q_0\rangle N\rangle$	$ m_1, \cdot\rangle q_0\rangle N\rangle$	$[\#, \cdot] q_0\rangle N\rangle$	$ t, j\rangle q_0\rangle N\rangle$	$[c, \cdot] q_0\rangle N\rangle$

$\forall i, j, k \in \{0, 1\}$

(53)

This machine is proper, never alters the mark track, and, for given mark track configuration, always runs for the same number of time steps. So substituting it in the looping machine produces a proper QTM M_2 . Furthermore, since M' takes $O(n)$ time and at most $n + 1$ space, M_2 runs for time $O(kn + k \log k)$, uses at most $\max(n, |k|) + 2$ space, and never moves the head before the starting cell.

The only *qubits* on which the overall QTM acts are the two qubits in the quantum tracks of the cells marked by t and c , the auxiliary qubit stored on the mark track (let's label this m), and the two internal qubits stored in the internal states q_2^j of M_1 and $q_{1,2}^j$ of M_2 (let's label these auxiliary M_1 and M_2 internal qubits by $a_{1,2}$). Qubits m , a_1 and a_2 are initially in the $|0\rangle$ state. Apart from classical processing to move the head into the correct location, M_1 just applies a CNOT between the c and a_1 qubits, followed by a CNOT between the a_1 and m qubits.

Similarly, M' applies a CNOT between the m and a_2 qubits, a control- U operation between the a_2 and t qubits, and a final CNOT between the m and a_2 qubits. Letting cU denote the control- U gate, and recalling that the a_2 qubit starts off in the $|0\rangle$ state, the overall effect of this is:

$$CNOT_{ma_2} cU_{a_2t} CNOT_{ma_2} |\psi\rangle_{mt} |0\rangle_{a_2} = cU |\psi\rangle_{mt} |0\rangle_{a_2}, \quad (54)$$

i.e. M' effectively applies a single control- U operation between the m and t qubits. M_2 repeats the M' operations k times, so the overall action of M_2 is to apply a control- U operation k times between the m and t qubits.

The time-reversal of M_1 simply applies its CNOTs in the reverse order. So dovetailing M_1, M_2, M_1^\dagger carries out the operation

$$CNOT_{ca_1} CNOT_{a_1m} cU_{mt}^k CNOT_{a_1m} CNOT_{ca_1} |\psi\rangle_{ct} |0\rangle_{a_1} |0\rangle_m = cU_{ct}^k |\psi\rangle_{ct} |0\rangle_{a_1} |0\rangle_m. \quad (55)$$

Thus the overall action of the entire QTM is to apply a cU^k operation between the control and target qubits, as required, leaving the looping and mark tracks back in the configuration they started in. Furthermore, the QTM runs for time

$O(kn + k \log k)$, uses at most $\max(n, |k|) + 2$ space, and is well-formed, normal form, unidirectional and proper, as claimed. \square

6.4.2 Iterating over the control qubits

We want to apply a $cU_\varphi^{2^{n-1}}$ operation between the n 'th output qubit and the ancilla qubit, for each of the N output qubits. (Note that the qubits are stored in reverse order on the quantum track, so the 1st output qubit is stored in the $N + 1$ 'th cell of the quantum track and the N 'th output qubit is stored in the 2nd cell.) As we will need to use a second looping machine to iterate over the qubits, we refer to the looping track for the cU^k machine as the “ cU looping track”, and the track for the second looping machine used here as the “outer looping track”.

We will need a reversible TM which, on suitable mark and cU looping track configurations (cf. Lemma 31), moves the c marker on the mark track one cell to the left, and doubles the value k written on the cU looping track. It is convenient (though slightly less efficient) to divide the implementation into two parts, dovetailed together: STEP_c which shifts the c marker one cell to the left, and DBL which doubles the value on the cU looping track. The following well-formed, normal form, partial transition function implements the reversible TM STEP_c (which only acts on the mark track):

	#	t_i	c	
q_0		(t_i, q_1, R)		
q_1	$(\#, q_1, R)$		$(\#, q_2, L)$	
q_2	(c, q_3, L)			
q_3	$(\#, q_3, L)$	(t_i, q_f, N)		
q_f	$(\#, q_0, N)$	(t_i, q_0, N)	(c, q_0, N)	(56)

$\forall i \in \{0, 1\}$

Doubling a number in binary simply appends a 0 onto the binary representation of the number. The following well-formed, normal form, partial transition function for DBL (which acts only on the cU looping track) does precisely this, taking advantage of the fact that k is always a power of 2 (so, in little-endian order, always consists of a string of 0's followed by a single 1) to simplify the construction:

	#	0	1	
q_0		$(\#, q_1, R)$	$(\#, q_2, R)$	
q_1		$(0, q_1, R)$	$(0, q_2, R)$	
q_2	$(1, q_3, L)$			
q_3	$(0, q_f, N)$	$(0, q_3, L)$		
q_f	$(\#, q_0, N)$	$(0, q_0, N)$	$(1, q_0, N)$	(57)

We are now in a position to implement the complete controlled- U_φ stage of the quantum phase estimation algorithm. We initialise the states of the auxiliary

tracks, by first running a simple reversible TM that changes the $[\#, \#]$ in the first cell of the mark and cU looping tracks into $[t_0, 1]$, then returns to the starting cell and halts. (Constructing a reversible TM for this is trivial.) We then dovetail this with a proper, normal form reversible TM that scans to the end of the output qubits, changes the $\#$ on the mark track into a c , and halts. The following normal form, partial transition function (which acts only on the mark and quantum tracks) accomplishes this:

	$[\#, \#]$	$[\#, i]$	$[t_i, j]$
q_0			$([t_i, j], q_1, R)$
q_1	$([\#, \#], q_2, L)$	$([\#, i], q_1, R)$	
q_2		$([c, i], q_3, L)$	
q_3		$([\#, i], q_3, L)$	$([t_i, j], q_f, N)$
q_f	$([\#, \#], q_0, N)$	$([\#, i], q_0, N)$	$([t_i, j], q_0, N)$

(58)

$\forall i, j \in \{0, 1\}$

On the mark track, this leaves a t_0 and c over the ancilla qubit and the first output qubit,¹ respectively, and blanks everywhere else. The configuration prepared on the cU looping track corresponds to the number 1 written in binary.

Next, we copy N from the input track to the outer looping track using the COPY machine from Lemma 24. We then run a reversible looping TM from Lemma 28 which uses this track as its input track (so it will loop N times in total). For the special state of this looping machine, we substitute a QTM which dovetails the cU^k machine from Lemma 31 with the STEP_c and DBL machines constructed above.

On appropriate input, the cU^k machine from Lemma 31 runs for a number of steps which depends only on the classical part of the input supplied on the mark and looping tracks, and not on the quantum state in the quantum track. The STEP_c and DBL reversible TMs don't touch the quantum track at all. So as long as the mark and cU looping tracks are always in a suitable configuration (cf. Lemma 31) before the cU^k machine is run in each iteration, the outer looping machine will be proper.

Now, we already initialised the cU looping track to $k = 1$, above, and marked the ancilla qubit as the target and the first output qubit as the control. So the initial configuration of the mark and cU looping tracks is suitable input for a cU^k machine. When cU^k is run in the first iteration of the outer looping machine, it applies a $cU^1 = cU$ between the first output qubit and the target qubit. STEP_c and DBL machines then run, which shifts the control marker one cell to the left onto the second output qubit and doubles k to 2, ready for the next iteration. In the second iteration, the looping machine therefore applies a cU^2 between the second output

¹Recall that the first output qubit is the last one on the quantum track

qubit and the ancilla, and doubles k to 4. The looping machine goes through a total of N iterations, each time feeding suitable input to the cU^k machine to make it apply a $cU^{2^{n-1}}$ between the n 'th output qubit and the ancilla, before halting with N written on the outer looping track, a t_0 and c in the first and second cells of the mark track, and $N + 1$ written on the cU looping track.

Thus the outer looping machine is well-formed, normal form unidirectional and proper. It runs for a total of N iterations with overhead $N \log N$ (cf. Lemma 28). In the n 'th iteration it runs the cU^k machine from Lemma 31 once with $k = 2^{n-1}$, and runs the STEP_c and DBL machines once each. The n 'th iteration of the cU^k machine takes time $O(N + 2^{n-1} \log 2^{n-1})$, and the STEP_c and DBL machines always take time $O(n)$. So the outer looping machine runs for a total time $O(N^2 2^N)$.

Finally, we can uncompute (reset) the auxiliary tracks by running a reversible TM which transforms the configuration $N; N; t_0c; N + 1$ left on the input, outer looping, mark, and cU looping tracks, into $N; \#; \#; \#$. We do this by running the reversal COPY^\dagger of the copying machine from Lemma 24 on the input and outer looping tracks to erase the outer looping track, decrementing the cU looping track and running COPY^\dagger again to erase that track, and running a trivial reversible TM that changes t_0c on the mark track into $\#\#$.

Putting everything together, the control- U stage runs for a total time $O(N^2 2^N)$ and requires space $N + 3$ (2 more cells than the number of qubits, which is $N + 1$). When dovetailed with the preparation stage from Section 6.3, the overall QTM we have constructed implements the control- U circuit of Figure 6. If φ_k denotes the k 'th digit in the binary decimal expansion of φ , then this prepares the state [NC00]

$$\frac{1}{2^N} \left(|0\rangle + e^{2\pi i 0.\varphi_N} |1\rangle \right) \left(|0\rangle + e^{2\pi i 0.\varphi_{N-1}\varphi_N} |1\rangle \right) \dots \dots \left(|0\rangle + e^{2\pi i 0.\varphi_2 \dots \varphi_{N-1}\varphi_N} |1\rangle \right) \left(|0\rangle + e^{2\pi i 0.\varphi_1 \varphi_2 \dots \varphi_N} |1\rangle \right) \quad (59)$$

on the N output qubits (ordered as they are on the quantum track). All other tracks are blank, except for the input track which still has the input N written on it as a little-endian binary number.

6.5 Locating the LSB

The final stage of the quantum phase estimation algorithm is to apply the inverse quantum Fourier transform to the quantum state generated by the control- U stage. The structure of the quantum circuit for the QFT is rather reminiscent of that of the control- U stage. It again involves applying a ‘‘cascade’’ of control- U^{2^n} gates between pairs of qubits, which we already know how to implement! Only now there is a new cascade starting from each output qubit (see Figure 7).

However, the phase 2^{-N} of the control- $U_{2^{-N}}$ rotation needed in the QFT circuit depends on the total number of qubits N that the QFT is being applied to. If we

simply implemented the inverse QFT circuit directly on all N output qubits, the entries of the transition function in our cU^k machine from Section 6.4 would have to depend on the input N , which is not allowed. It is not at all obvious whether the inverse QFT circuit on N qubits can be implemented *exactly* when N is given as input.¹

On the other hand, the entries of the transition function *are* allowed to depend on the phase φ that we are estimating – indeed, they necessarily do so, since the output of the P_n QTM in Theorem 11 depends on n (or equivalently on φ). Rather than implementing the QFT on N qubits, we take a different approach. We supply our QTM with the $U_{2^{-|\varphi|}}$ gate, and implement the inverse QFT on $|\varphi|$ qubits, independent of the input N . (I.e. just enough qubits to hold all the digits of the binary decimal expansion of φ .) However, for this to work, we must first identify which output qubit holds the least significant bit (LSB) of φ , so that we know which qubits to apply the $|\varphi|$ -qubit inverse QFT to. The role of the input N is merely to provide us with an upper-bound on $|\varphi|$, which allows us to identify the LSB in finite time and space.

Recall that the control- U stage prepared the state [NC00]

$$\frac{1}{2^{\frac{N}{2}}} \left(|0\rangle + e^{2\pi i 0.\varphi_N} |1\rangle \right) \left(|0\rangle + e^{2\pi i 0.\varphi_{N-1}\varphi_N} |1\rangle \right) \dots \\ \dots \left(|0\rangle + e^{2\pi i 0.\varphi_2 \dots \varphi_{N-1}\varphi_N} |1\rangle \right) \left(|0\rangle + e^{2\pi i 0.\varphi_1 \varphi_2 \dots \varphi_N} |1\rangle \right) \quad (60)$$

on the N output qubits (where φ_k denotes the k 'th digit in the binary decimal expansion of φ).

The least significant bit of φ is the $|\varphi|$ 'th bit, and by assumption (cf. Theorem 11) $|\varphi| \leq N$. So the $|\varphi|$ 'th bit of φ is 1 and the $|\varphi| + 1 \dots N$ 'th bits are 0. Thus the last $N - |\varphi|$ output qubits are in the $|+\rangle$ state, and the $|\varphi|$ 'th qubit, which corresponds to the least significant bit of φ , is in the $|-\rangle$ state. (Recall that the output qubits are stored on the quantum track in reverse order.)

Therefore, if we construct a QTM that steps right, applying Hadamard rotations to the output qubits in the quantum track until it obtains a $|1\rangle$ and halts, we will be able to locate the least significant bit of φ . More precisely, this will rotate the first $N - |\varphi|$ output qubits into the $|0\rangle$ state, and halt with the $N - |\varphi| + 1$ 'th qubit (which corresponds to the $|\varphi|$ 'th bit of φ) rotated into the $|1\rangle$ state. The least significant bit of φ is therefore identified by the first $|1\rangle$ on the quantum track after this machine has finished running.

The following well-formed, normal form, partial quantum transition function (which acts only on the quantum track) implements a proper QTM that steps along the quantum track, applying Hadamards until it obtains a $|1\rangle$. It can be verified that

¹The universal QTM construction of Bernstein and Vazirani [BV97] can approximate the QFT on N qubits to arbitrary precision, but it does not implement it exactly.

it obeys the three conditions of Theorem 19, so can be completed to a well-formed QTM:

	#	0	1	
q_0				$ \#\rangle q_1\rangle R\rangle$
q_1	$ \#\rangle q_3\rangle L\rangle$	$\frac{1}{\sqrt{2}}(0\rangle + 1\rangle) q_2\rangle N\rangle$	$\frac{1}{\sqrt{2}}(0\rangle - 1\rangle) q_2\rangle N\rangle$	(61)
q_2		$ 0\rangle q_1\rangle R\rangle$	$ 1\rangle q_3\rangle L\rangle$	
q_3	$ \#\rangle q_f\rangle N\rangle$	$ 0\rangle q_3\rangle L\rangle$		
q_f	$ \#\rangle q_0\rangle N\rangle$	$ 0\rangle q_0\rangle N\rangle$	$ 1\rangle q_0\rangle N\rangle$	

(Note that the qubit in the starting cell of the quantum track is still in the state $|1\rangle$ from the preparation and control- U stages. As we will not need this qubit again, this machine sets it to $|\#\rangle$ to mark the starting cell and leaves it in that state when it halts, for later convenience.)

The configuration written on the quantum track after this machine has finished is the $N + 1$ -qubit state:

$$|\#\rangle |0\rangle \overset{N-|\varphi|}{\dots} |0\rangle |1\rangle \frac{1}{2^{\frac{|\varphi|-1}{2}}} \left(|0\rangle + e^{2\pi i 0 \cdot \varphi_{|\varphi|-1} \varphi_{|\varphi|}} |1\rangle \right) \left(|0\rangle + e^{2\pi i 0 \cdot \varphi_{|\varphi|-2} \varphi_{|\varphi|-1} \varphi_{|\varphi|}} |1\rangle \right) \dots$$

$$\dots \left(|0\rangle + e^{2\pi i 0 \cdot \varphi_1 \varphi_2 \dots \varphi_{|\varphi|-1} \varphi_{|\varphi|}} |1\rangle \right). \quad (62)$$

Note that the first $N - |\varphi| + 2$ cells of the quantum track are not entangled with the rest of the track.

It will be helpful for later to record the number of output bits that we are skipping, i.e. to compute $N - |\varphi|$ and store the result on a separate auxiliary “count track”. We can do this using a construction that is very similar to the unary-to-binary converter of Lemma 30. Specifically, we will construct a reversible TM COUNT that steps right over $|0\rangle$'s until it finds a $|1\rangle$ on the quantum track, running the INC machine once each time it steps right. (This means first returning the head to the starting cell, running INC, then returning the head back to where it was previously.)

Consider the following partial transition function, which satisfies the conditions of Theorem 18 so can be completed to a proper, normal form, reversible TM:

	#	0	1	
q_0	$(\#, q_1, R)$			
q_1	$(\#, q_4, L)$	$(\#, q_2, L)$	$(1, q_4, L)$	
q_2	$(\#, q, N)$	$(0, q_2, L)$		(63)
q	$(\#, q_3, R)$			
q_3	$(0, q_1, R)$	$(0, q_3, R)$		
q_4	$(\#, q_f, N)$	$(0, q_4, L)$		
q_f	$(\#, q_0, N)$	$(0, q_0, N)$	$(1, q_0, N)$	

This machine steps right over 0's until it encounters a 1 (or reaches the end of the input), at which point it moves its head back to the starting cell and halts. (It assumes that the first cell is marked with a #.) However, each time it is about to step right, it first temporarily marks the current head location by overwriting the 0 with a #, returns the head to the starting cell, enters an internal state q for one time step, then returns the head back to where it was before, resets the quantum track to 0, and continues stepping right from where it left off. By Lemma 21, we can substitute for the state q the INC machine from Lemma 27 (acting on an initially blank “count track”). If we run the above TM on the quantum track, it will step right precisely $N - |\varphi|$ times before finding the first $|1\rangle$, so it will halt with $N - |\varphi|$ written to the count track.

The configuration left on the tape now consists of the little-endian binary number N written on the input track, the little-endian binary number $N - |\varphi|$ written on the count track, and the $N + 1$ -qubit state from (62) written on the quantum track.

As the qubits storing the $N - |\varphi|$ trailing 0's of the N -digit binary expansion of φ now play no further role, it is convenient to shift the starting cell of subsequent machines to the tape cell containing the LSB. To this end, we copy the values N and $N - |\varphi|$ from the input and count tracks onto new auxiliary tracks. We then use a LOOP machine from Lemma 28, which uses the original count track as its input and work track, to run a SHIFT machine from Lemma 25 acting on the tracks holding the copies of the count and input tracks. Thus the SHIFT machine will run a total of $N - |\varphi|$ times, thereby shifting the binary strings on the copies so that they start in the $N - |\varphi|$ 'th cell – the one containing the LSB of $|\varphi|$ in the quantum track. We then run a trivial TM that steps the head right over the initial string of $|0\rangle$ s on the quantum track, until it encounters the first $|1\rangle$ identifying the LSB, and halts with the head at this LSB cell. When we refer to the input and count tracks in the following section, we mean the copies shifted to start at the LSB cell. (At the very end of the computation, we will uncompute the shifted copies of the count and input tracks and step the head left back to the original starting cell, to ensure the overall QTM remains proper.)

Now, we have been careful to ensure that the head is never moved to a location before the starting cell in any of the reversible and quantum TMs that we have constructed. Furthermore, the section of the complete tape configuration located before the LSB cell is unentangled with the rest (see (62), and note that all tracks other than the quantum track are in classical configurations). Thus, if we start any of our TMs in the LSB cell, it acts as if its input were the section of the tape configuration located after (and including) the LSB cell. We can therefore ignore the tape configuration of the first $N - |\varphi|$ cells in the following section, and restrict our attention to the following $|\varphi|$ cells.

6.6 QFT stage

We are now ready to apply the $|\varphi|$ -qubit inverse QFT to the $|\varphi|$ -qubit state stored on the quantum track. The inverse QFT circuit on $|\varphi|$ qubits consists of cascades of $cU_{2^{-|\varphi|}}^k$ gates, very reminiscent of the control- U stage we have already implemented (see Figure 7). The construction will therefore be similar. In the following, as we have shifted the starting cell, we relabel the output qubits and refer to the LSB qubit as the 1st qubit, the last one as the $|\varphi|$ 'th qubit.

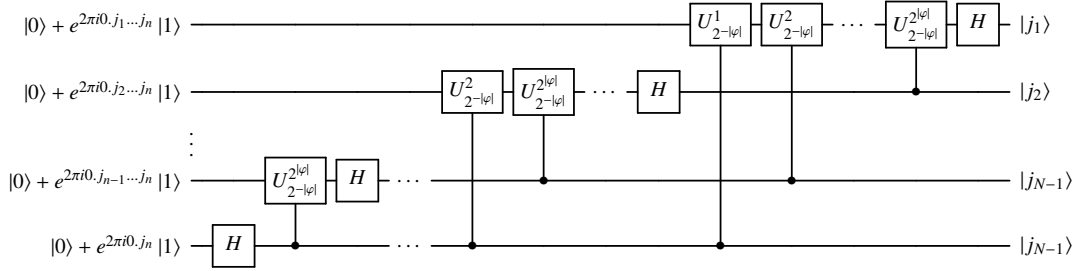


Figure 7: The inverse QFT stage of the quantum phase estimation circuit (cf. Fig. 5.1 in [NC00]).

In each cascade, we want to apply a $cU_{2^{m-n}} = (cU_{2^{-|\varphi|}})^{2^{|\varphi|+m-n}}$ gate between the m 'th and n 'th qubit ($m < n$). Once again, we will use mark and cU looping tracks to hold the input to a $cU_{2^{-|\varphi|}}^k$ machine from Lemma 31. However, the main loop will now consist of two nested loops: an outer loop to iterate the control qubit m of the $cU_{2^{-|\varphi|}}^{2^{|\varphi|+m-n}}$ gate over each output qubit, and an inner loop to iterate the target qubit n over qubits $m + 1$ through $|\varphi|$.

We first run a TM that initialises the mark and cU looping tracks, so that the mark track contains the configuration ct in the first two cells, and the cU looping track contains a string of $|\varphi| - 1$ 0's followed by a 1. Note that this initial configuration of the cU looping track is the little-endian binary representation of the number $2^{|\varphi|-1}$. (Constructing a proper normal form, reversible TM that implements all of this is an easy exercise.)

We also use the Inc machine from Lemma 27 to increment the number $N - |\varphi|$ stored on the count track to $N - |\varphi| + 1$. We then run the SUB machine from Lemma 29, with the input and count tracks as the input tracks and the inner looping track as the output track, to write the number $N - (N - |\varphi| + 1) = |\varphi| - 1$ to the inner looping track.

Inner loop The inner looping machine is very similar to the main looping machine from the control- U stage. We use a Loop machine from Lemma 28, with the inner looping track as its input and work track. In each iteration, we first run a TM to divide the value on the cU looping track by 2. (This is simply the reversal

DBL[†] of the machine constructed in Section 6.4.2.) We dovetail this with the $cU_{2^{|\varphi|}}^k$ machine from Lemma 31. Finally, we dovetail this with a TM that moves the target marker t on the mark track one cell to the right. (This can be implemented by taking the reversal STEP _{c} [†] of the machine from Section 6.4.2 and replacing c with t in its transition rules to give a STEP _{t} [†] machine.)

Assume that the inner looping track is initialised to $|\varphi| - m$, the control and target markers on the mark track are initially over the m 'th and $m + 1$ 'th qubits, and the cU looping track is initialised to $2^{|\varphi|}$. Then the effect of this inner looping machine is to apply $cU_{2^{m-n}}$ operations between the m 'th and n 'th qubit for all $m < n \leq |\varphi|$ – the cascade of cU 's starting from the m 'th qubit in Figure 7.

This inner looping TM leaves the mark track in the configuration with a c in the same cell that it started out in, a t in the $|\varphi| + 1$ 'th cell (i.e. the cell after the last output qubit), and 2^{m-1} written on the cU looping track. (The latter is the configuration consisting of an initial string of 0's, followed by a 1 in the cell containing the c on the mark track.)

Outer loop For the outer looping machine, we use a LOOP2 machine from Lemma 28 running on the input and count tracks (which hold the numbers N and $N - |\varphi| + 1$, respectively). By Lemma 28, this machine will therefore loop $|\varphi| - 1$ times.

In each iteration, we first run the inner looping machine. Note that for the first iteration, the auxiliary tracks are already initialised as assumed above for the value $m = 1$. We dovetail this with a simple QTM that applies a Hadamard operation to the current control qubit (the one marked by a c on the mark track).

We then run a TM that changes the 1 on the cU looping track to a 0, shifts the c on the mark track one cell to the right, steps right, and changes $[\#, \#]$ on the mark and cU looping tracks to $[t, 0]$. The machine then steps right, changing $\#$ to 0 on the cU looping track as it goes, until it reaches the end of the output qubits. Whereupon it changes $\#$ on the cU looping track to 1, steps right, and changes the t on the mark track to $\#$. Finally, it returns to its starting cell and halts. (Again, by now, constructing a proper, normal form, reversible TM for this is straightforward.) We dovetail this with the DEC TM, acting on the inner looping track.

The effect of all this is to reset the configuration of the auxiliary tracks, ready for the next iteration of the inner loop. The control marker c is shifted to the next qubit along, qubit m say, and the target marker t is placed over the adjacent $m + 1$ 'th qubit. The cU looping track is reset to $2^{|\varphi|-1}$, and the inner looping track is decremented to $|\varphi| - m$, as required. In other words, the auxiliary tracks are initialised to the desired configuration for the new value of m .

Thus the outer looping machine runs the inner looping machine $|\varphi| - 1$ times for each of the output qubits $m = 1 \dots |\varphi| - 1$. Each time it runs, the inner looping

machine applies the desired cascade of cU operations between the m 'th and n 'th qubits for all $m < n \leq |\varphi|$. The outer looping machine then applies a final Hadamard operation to the m 'th qubit, before moving onto the next qubit. One can therefore see that the overall effect is to apply the inverse quantum Fourier Transform circuit to the $|\varphi|$ output qubits.

6.7 Reset Stage

To reset all the auxiliary tracks, we dovetail the outer looping TM with a sequence of TMs that uncompute all the configurations we previously prepared on these tracks.

The first of these TMs changes the ct left on the mark track over the final two qubits to $\#\#$, and changes the final configuration $1; 2^{|\varphi|-1}$ of the inner looping and cU looping tracks to the blank configuration. (Note that $2^{|\varphi|-1}$ is the configuration consisting of a leading string of 0's followed by a single 1 on the final output qubit, so is straightforward to reset without any arithmetic computations.)

To reset the final configuration $N; N - |\varphi| + 1$ of the input and count tracks, we start by decrementing the count track to $N; N - |\varphi|$ (using DEC). Recall that we redefined the location of the starting cell before implementing the inverse QFT machine, so that the starting cell became the cell containing the LSB of φ . We also shifted the copies of the input and count tracks accordingly (Section 6.5). We now step the head back to the original starting cell (which is easily located as we left a $\#$ written there on the quantum track, cf. Section 6.5). We can then run the reversal SHIFT^\dagger (constructed using Lemma 23) of the machine used in Section 6.5, to shift the copies of the input and count tracks used in the QFT stage back left again to the original starting cell, and run the reversal COPY^\dagger of the copy machines to erase the copies of the input and count tracks.

To erase the original count track, which contains $N - |\varphi|$, we must run the reversal COUNT^\dagger that we constructed in Section 6.5. (The first $N - |\varphi|$ qubits remain in the $|0\rangle$ state and the LSB qubit in the $|1\rangle$ state, so COUNT^\dagger will indeed uncompute $N - |\varphi|$.) Finally, to erase the original input track, we first use the SHIFT^\dagger machine from Lemma 25 on the *quantum* track, to shift the entire final state of the output qubits left. Note that this will work because the first qubit (the ancilla qubit of the control- U stage) was set to $\#$ in Section 6.5. We can then run the reversal UTOB^\dagger of the unary-to-binary machine from Lemma 30, acting on the quantum track and input track, to uncompute the binary conversion of the unary input encoded in the length of the qubit state. Note that the UTOB machine only checks whether the symbols on the unary track are $\#$ or non- $\#$. So the fact that the configuration of the quantum track is no longer necessarily a string of N $|1\rangle$'s does not matter; all that matters is that it is a string of N non- $\#$'s.

The end result of all this is to reset all the auxiliary tracks to the blank state,

leaving the output of the inverse QFT circuit stored in the first N cells of the quantum track.

6.8 Analysis

By dovetailing together the preparation stage (Section 6.3), control- U_φ stage (Section 6.4), inverse-QFT stage (Section 6.6) and reset stage (Section 6.7), we have succeeded in constructing a family of well-formed, normal form, unidirectional quantum Turing Machines P_n that behave properly on input $N \geq |\varphi|$ written in unary, and implement the quantum phase estimation algorithm on N qubits for phase φ . By construction, P_n satisfies part (i) of Theorem 11.

Now, we know [NC00] that the quantum phase estimation algorithm outputs the exact binary decimal expansion of the phase, as long as we run it on enough qubits to store the entire binary decimal expansion of the phase. Thus for $N \geq |\varphi| = |n|$, the P_n writes out the binary decimal expansion of φ to N bits (in little-endian order, padding with 0's as necessary). We choose φ to be the rational number whose binary decimal expansion contains the digits of n in reverse order after the decimal, where n indexes P_n . So our QTM P_n implements the computation claimed in part (ii) of Theorem 11.

We were careful throughout to keep tight control on the space requirements of the reversible and quantum TMs that we used to construct P_n . In fact, none of them used more than $N + 3$ space. Finally, all steps of the computation take time $O(\text{poly } N)$, except the cU_φ^k and cU_φ^k computations, which take time $O(2^N)$. Thus the overall run-time is $O(\text{poly}(N)2^N)$. Thus P_n fulfils the space and time requirements in part (ii) of Theorem 11.

The last thing we must check in order to finish the proof of Theorem 11 is part (iii). But the partial transition function defined so far for P_n satisfies part (iii). It is trivial to check that one can complete this to a full transition function, whilst still satisfying part (iii). Indeed, by Theorem 19 the problem is equivalent to completing an orthonormal set of vectors with coefficients in

$$\mathcal{S} = \left\{ 0, 1, \pm \frac{1}{\sqrt{2}}, e^{i\pi\varphi}, e^{i\pi 2^{-|\varphi|}} \right\} \quad (64)$$

to a full orthonormal basis with coefficients in \mathcal{S} . Since any normalised vector with coefficients in \mathcal{S} must be proportional to a vector in the canonical basis $\{e_j\}_j$ or of the form $\frac{1}{\sqrt{2}}(e_j \pm e_k)$, $j \neq k$, the result is immediate. This completes the proof of Theorem 11.

Note that the QTM we have constructed only implements the computation correctly when supplied a valid upper-bound on the number of digits in the binary decimal expansion of the phase. If the input is not actually a correct upper bound, we make no claim about the behaviour of the QTM. (This will be significant later,

when we come in Section 9 to bound ground state energies of a Hamiltonian constructed out of this QTM.)

7 Encoding QTMs in local Hamiltonians

In this section, we give a self-contained analysis of a local Hamiltonian construction that generalises that of Gottesman and Irani [GI09], in order to prove a general result about local Hamiltonian encodings of arbitrary quantum computations.

A full, rigorous analysis of the resulting Hamiltonian will then prove the main result of this section:

Definition 32 (Computational history state) A computational history state $|\psi\rangle_{CQ} \in \mathcal{H}_C \otimes \mathcal{H}_Q$ is a state of the form

$$|\psi\rangle_{CQ} = \frac{1}{\sqrt{\dim \mathcal{H}_C}} \sum_{t=1}^{d_C} |t\rangle |\psi_t\rangle, \quad (65)$$

where $\{|t\rangle\}$ is an orthonormal basis for \mathcal{H}_C , and $|\psi_t\rangle = \prod_{i=1}^t U_i |\psi_0\rangle$ for some initial state $|\psi_0\rangle \in \mathcal{H}_Q$ and set of unitaries $U_i \in \mathcal{B}(\mathcal{H}_Q)$.

\mathcal{H}_C is called the clock register and \mathcal{H}_Q is called the computational register. If U_t is the unitary transformation corresponding the t 'th step of a quantum computation, then $|\psi_t\rangle$ is the state of the computation after t steps. We say that the history state $|\psi\rangle$ encodes the evolution of the quantum computation.

Theorem 33 (Local Hamiltonian QTM encoding)

Let $\mathbb{C}^d = (\text{span}\{|\ominus\rangle, |\otimes\rangle\} \oplus \mathbb{C}^C) \otimes (\text{span}\{|\ominus\rangle, |\otimes\rangle\} \oplus \mathbb{C}^Q)$ be the local Hilbert space of a 1-dimensional chain of length L .

For any well-formed unidirectional Quantum Turing Machine $M = (\Sigma, Q, \delta)$ and any constant $K > 0$, we can 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 .
- (iii). If we call $\mathcal{H}(L-2) := (\mathbb{C}^C)^{\otimes L-2} \otimes (\mathbb{C}^Q)^{\otimes L-2} =: \mathcal{H}_C \otimes \mathcal{H}_Q$ and define $S_{br} \subset \mathcal{H} = \text{span}(|\ominus\rangle \otimes |\ominus\rangle)_1 \otimes \mathcal{H}(L-2) \otimes \text{span}(|\otimes\rangle \otimes |\otimes\rangle)_L$, the unique ground state of $H(L)|_{S_{br}}$ is a computational history state encoding the evolution of M on input corresponding to the unary representation of the number $L - K - 3$, running on a finite tape segment of length $L - 3$.

Moreover, if M is proper on input given by the unary representation of $L - K - 3$, 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 computational 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). *Finally, if M satisfies part (iii) of Theorem 11, then h has the following form*

$$h = A + (e^{i\pi\varphi} B + e^{i\pi 2^{-|\varphi|}} C + \text{h.c.}), \quad (66)$$

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

Though we will not need it for our purposes, it is not difficult to prove that the next-highest eigenstate of $H(L)|_{S_{\text{br}}}$ has energy $O(1/T^3)$, where T is the total number of time-steps encoded in the computational history state.

7.1 Preliminaries

As our construction draws heavily on [GI09], we will follow their notation and terminology, which we summarise here. We divide the chain into multiple tracks:

\langle	...	Track 1: Clock oscillator	...	\rangle
	...	Track 2: Counter TM head and state	...	
	...	Track 3: Counter TM tape	...	
	...	Track 4: QTM head and state	...	
	...	Track 5: QTM tape	...	
	...	Track 6: Time-wasting tape	...	

The local Hilbert space at each site is:

$$\begin{aligned}
\mathcal{H} := & \left(\text{span}\{|s, p, \sigma\rangle\} \oplus \text{span}\{|\langle\rangle, |\rangle\rangle\} \right) \otimes \left(\text{span}\{|q, \tau, \gamma\rangle\} \oplus \text{span}\{|\langle\rangle, |\rangle\rangle\} \right) \\
s \in & \left\{ \begin{array}{c} \begin{array}{c} \rightarrow \\ \circ \\ \leftarrow \end{array}, \begin{array}{c} \leftarrow \\ \circ \\ \rightarrow \end{array}, \begin{array}{c} \leftarrow \\ \circ \\ \leftarrow \end{array}, \begin{array}{c} \rightarrow \\ \circ \\ \rightarrow \end{array}, \begin{array}{c} \rightarrow \\ \circ \\ \leftarrow \end{array}, \begin{array}{c} \leftarrow \\ \circ \\ \leftarrow \end{array}, \circ, \odot \end{array} \right\} \text{ for } i \in \{1, \dots, K\}, \\
p \in & P' \cup \{\circ, \odot\} \text{ where } P := P_L \cup P_N \cup P_R \text{ and } P' := P \cup P'_R, \\
q \in & Q' \cup P \cup P'_L \cup \{r_x\} \cup \{\circ, \odot\} \text{ where } Q := Q_L \cup Q_N \cup Q_R, Q' := Q \cup Q'_L \\
& \text{and } x \in Q \\
\sigma \in & \Sigma, \\
\tau \in & \Xi \text{ where } \Xi := \{\vdash, \#, 0, \dots, \zeta - 1\} \text{ with } \zeta = |\sigma \times Q|, \\
\gamma \in & \Xi \cup \{\vdash_q\} \text{ where } q \in Q'.
\end{aligned} \tag{67}$$

Σ is the tape alphabet of our given QTM M . Ξ is the alphabet of the counter TM. P_L, P_N, P_R are the sets of internal states of the counter TM that can be entered by the TM head moving left, not moving, or moving right (respectively). The states $p' \in P'_R$ duplicate the states $p \in P_R$, and $p' \in P'_L$ duplicate those in P_L . Similarly for the internal states Q_L, Q_N, Q_R of the QTM, with $q' \in Q'_L$ duplicating the states $q \in Q_L$. Recall that by the unidirection property of reversible and quantum Turing Machines (Theorems 18 and 19), these sets are disjoint. The \circ and \odot Track 2 and 4 symbols are used for cells that do not currently hold the head.¹ The bracket states \langle, \rangle are the special boundary states $|\langle\rangle, |\rangle\rangle$ from Theorem 33 that will only appear at the ends of the chain. The role of the Track 1 states is described in Section 7.2. That of the Track 6 “time-wasting” tape, as well as the role of the additional Track 4 $P \cup P'_L$ states, will become apparent in Section 7.6.2.

This set of states defines a standard basis for the single-site Hilbert space \mathcal{H} . The product states over this single-site basis then give a basis for the Hilbert space $\mathcal{H}^{\otimes L}$ of the chain. We call these the *standard basis states*.

Definition 34 (Bracketed state) *We call standard basis states that match the regular expression $\langle \circ^* \rangle$ bracketed states (where \langle, \rangle stand for the states $|\langle\rangle_{\text{Tracks } 1-3}, |\langle\rangle_{\text{Tracks } 4-6}$ and $|\rangle\rangle_{\text{Tracks } 1-3}, |\rangle\rangle_{\text{Tracks } 4-6}$ respectively, and \circ stands for any single-site state other than \langle, \rangle). We denote by S_{br} the subspace spanned by the bracketed states.*

¹There are two blank symbols because we will need different symbols to the left and right of the head, in order to enforce the constraint that there is only one head on the track.

We will often denote configurations of multiple tracks by writing them vertically, e.g.

$$\begin{array}{|c|} \hline \leftarrow \\ \hline p_1 \\ \hline \end{array}, \quad \begin{array}{|c|c|} \hline \circ & \leftarrow \\ \hline \end{array}, \quad \begin{array}{|c|c|} \hline \leftarrow & \odot \\ \hline \circ & p_1 \\ \hline 0 & 1 \\ \hline \end{array}$$

denote, respectively, a single-site configuration of two tracks, a configuration on two neighbouring sites of one track, and a configuration on two neighbouring sites of three tracks. Since we will restrict throughout to the subspace S_{br} of bracketed states, the \leftarrow and \rightarrow states will only ever appear at the left and right ends of the chain, across all the tracks. As a shorthand, we will therefore denote configurations involving \leftarrow and \rightarrow with symbols stretching vertically across all tracks, e.g.

$$\begin{array}{|c|c|} \hline \leftarrow & \leftarrow \\ \hline \leftarrow & p_1 \\ \hline \end{array}, \quad \begin{array}{|c|c|} \hline \odot & \\ \hline p_1 & \rightarrow \\ \hline 1 & \\ \hline \end{array}.$$

We will also sometimes need to specify configurations in which the tracks are in any configuration *other than* \leftarrow , \rightarrow . We denote this with a $\neg\leftarrow$ or $\neg\rightarrow$ symbol stretching vertically across all tracks, e.g.

$$\begin{array}{|c|c|} \hline \neg\leftarrow & \leftarrow \\ \hline \neg\leftarrow & p_1 \\ \hline \end{array}, \quad \begin{array}{|c|c|} \hline \odot & \\ \hline p_1 & \neg\rightarrow \\ \hline 1 & \\ \hline \end{array}.$$

As usual in such constructions, the two-body Hamiltonian will contain two types of terms: *penalty terms* and *transition rule terms*. Penalty terms have the form $|ab\rangle\langle ab|$ where a, b are standard basis states. This adds a positive energy contribution to any configuration containing a to the left of b . We call ab an *illegal pair*, and denote a penalty term $|ab\rangle\langle ab|$ in the Hamiltonian by its corresponding illegal pair. We will sometimes also make use of single-site illegal states a . Note that single-site illegal states are easily implemented in terms of illegal pairs, by adding penalty terms $|ax\rangle\langle ax|$ and $|xa\rangle\langle xa|$ for all pairs ax and xa in which the single-site state appears.

Definition 35 (Legal and illegal states) *We call a standard basis state legal if it does not contain any illegal pairs, and illegal otherwise.*

We will make repeated use of Lemma 5.2 from [GI09], which lets us use penalty terms to restrict to configurations matching a regular expression:

Lemma 36 (Regexps) *For any regular expression over the standard basis of single-site states in which each state appears at most once, we can use penalty terms to ensure that any legal standard basis state for the system is a substring of a string in the regular set.*

All the regular expressions we will use start with a left-bracket state $\textcircled{<}$ and end with a right-bracket state $\textcircled{>}$. So whenever we apply this Lemma, the only substrings of the regular set that are *bracketed states* are in fact complete strings in the regular set, not substrings.

Transition rule terms have the form $\frac{1}{2}(|\psi\rangle - |\varphi\rangle)(\langle\psi| - \langle\varphi|)$, where $|\psi\rangle, |\varphi\rangle$ are states on the same pair of adjacent sites. We will often take $|\psi\rangle = |ab\rangle$ and $|\varphi\rangle = |cd\rangle$ to be standard basis states. This forces any zero-energy eigenstate with amplitude on a configuration containing ab to also have equal amplitude on the configuration in which that ab is replaced by cd . Conversely, if a zero-energy eigenstate has amplitude on a configuration containing cd , it must also have equal amplitude on the configuration in which that cd is replaced by ab . Thus we can think of these terms as implementing the transition rules $ab \rightarrow cd$, which we arbitrarily call the “forwards” direction, and the corresponding “backwards” transition $cd \rightarrow ab$. Following the notation in [GI09], we will denote transition rule Hamiltonian terms by their associated forwards transitions $ab \rightarrow cd$ or, more generally, $|\psi\rangle \rightarrow |\varphi\rangle$.

When a transition rule acts diagonally on a subset of the tracks, we will sometimes consider the *restriction* of the transition rule to those tracks. That is, for a transition rule with the general form $|ab\rangle_T |\psi\rangle_{T^c} \rightarrow |cd\rangle_T |\varphi\rangle_{T^c}$, where T is some subset of the tracks, the restriction of the rule to T is given by $|ab\rangle_T \rightarrow |cd\rangle_T$.

When we specify a Hamiltonian term only on a subset of the tracks, we implicitly mean that it acts as identity on the remaining (unspecified) tracks. We will assume throughout this section that the ground state subspace of the Hamiltonian is restricted to the subspace S_{br} of bracketed states, with $\textcircled{<}$ at the left end of the chain and $\textcircled{>}$ at the right. (We show how to enforce this later, in Section 9.)

7.2 Clock Oscillator

Every clock – even one unusual enough to be encoded in superposition into the ground state of a quantum many-body Hamiltonian! – needs some form of *oscillator* that oscillates with a fixed period (such as a pendulum), and a *counter* that is incremented after each complete oscillation of the oscillator (such as the hands on a clock). For the counter, we will use a reversible counter Turing Machine, described in Section 7.4. This section describes the Track 1 clock oscillator.

The legal configurations of Track 1 are states matching the regular expression $\textcircled{<} \textcircled{0}^* [\textcircled{\rightarrow_0} \textcircled{\leftarrow_0} \textcircled{\leftarrow_1} \textcircled{\rightarrow_1}] \textcircled{0}^* \textcircled{>}$. By Lemma 36, we can enforce this using penalty terms. Independent of the configuration of any other track, a $\textcircled{\ominus}$ arrow (either $\textcircled{\rightarrow_0}$ or $\textcircled{\rightarrow_1}$) sweeps right along the chain until it reaches the end, whereupon it turns




Figure 8: Evolution of the Track 1 clock oscillator.

around and becomes a \ominus . (When a \ominus reaches the end and turns around, the label on the arrow steps through the sequence $(0), (1), \dots, (K)$. This is used later on to correctly initialise the other tracks. The restriction to $L \geq K + 3$ in Theorem 33 ensures that there is enough space for such sequences of transitions.) This \ominus arrow sweeps left until it reaches the beginning of the chain, at which point it turns around and becomes a \ominus again. We call this entire sequence an *oscillator cycle*. (One complete cycle is illustrated in Figure 8 for a chain of length six.) On a chain of length L , one complete oscillator cycle takes $2(L - 2)$ steps. The following transition rules on Track 1 enforce this:

$$\begin{aligned}
 \begin{array}{|c|c|} \hline \ominus \\ \hline 0 \\ \hline \end{array} \begin{array}{|c|c|} \hline \ominus \\ \hline 0 \\ \hline \end{array} &\longrightarrow \begin{array}{|c|c|} \hline \ominus \\ \hline 0 \\ \hline \end{array} \begin{array}{|c|c|} \hline \ominus \\ \hline 0 \\ \hline \end{array}, & \begin{array}{|c|c|} \hline \ominus \\ \hline 0 \\ \hline \end{array} \begin{array}{|c|c|} \hline \ominus \\ \hline > \\ \hline \end{array} &\longrightarrow \begin{array}{|c|c|} \hline \ominus \\ \hline 0 \\ \hline \end{array} \begin{array}{|c|c|} \hline \ominus \\ \hline > \\ \hline \end{array}, & \begin{array}{|c|c|} \hline \ominus \\ \hline (i) \\ \hline \end{array} \begin{array}{|c|c|} \hline \ominus \\ \hline (i) \\ \hline \end{array} &\longrightarrow \begin{array}{|c|c|} \hline \ominus \\ \hline (i+1) \\ \hline \end{array} \begin{array}{|c|c|} \hline \ominus \\ \hline (i) \\ \hline \end{array}, & \begin{array}{|c|c|} \hline \ominus \\ \hline (K) \\ \hline \end{array} \begin{array}{|c|c|} \hline \ominus \\ \hline (K) \\ \hline \end{array} &\longrightarrow \begin{array}{|c|c|} \hline \ominus \\ \hline 0 \\ \hline \end{array} \begin{array}{|c|c|} \hline \ominus \\ \hline (K) \\ \hline \end{array}, \\
 \begin{array}{|c|c|} \hline \ominus \\ \hline 0 \\ \hline \end{array} \begin{array}{|c|c|} \hline \ominus \\ \hline 0 \\ \hline \end{array} &\longrightarrow \begin{array}{|c|c|} \hline \ominus \\ \hline 0 \\ \hline \end{array} \begin{array}{|c|c|} \hline \ominus \\ \hline 0 \\ \hline \end{array}, & \begin{array}{|c|c|} \hline \ominus \\ \hline \bar{1} \\ \hline \end{array} \begin{array}{|c|c|} \hline \ominus \\ \hline \bar{1} \\ \hline \end{array} &\longrightarrow \begin{array}{|c|c|} \hline \ominus \\ \hline \bar{1} \\ \hline \end{array} \begin{array}{|c|c|} \hline \ominus \\ \hline \bar{1} \\ \hline \end{array}, & \begin{array}{|c|c|} \hline \ominus \\ \hline \bar{1} \\ \hline \end{array} \begin{array}{|c|c|} \hline \ominus \\ \hline > \\ \hline \end{array} &\longrightarrow \begin{array}{|c|c|} \hline \ominus \\ \hline \bar{1} \\ \hline \end{array} \begin{array}{|c|c|} \hline \ominus \\ \hline > \\ \hline \end{array}, & \begin{array}{|c|c|} \hline \ominus \\ \hline \bar{1} \\ \hline \end{array} \begin{array}{|c|c|} \hline \ominus \\ \hline \bar{1} \\ \hline \end{array} &\longrightarrow \begin{array}{|c|c|} \hline \ominus \\ \hline \bar{1} \\ \hline \end{array} \begin{array}{|c|c|} \hline \ominus \\ \hline \bar{1} \\ \hline \end{array}.
 \end{aligned} \tag{68}$$

When a left-moving \ominus arrow returns to the beginning of the chain, it turns around and becomes a right-moving \ominus arrow again. However, the label 0 on the arrow may change to a 1, depending on the Track 2 state. The \ominus arrow transitions to $\bar{1}$ if the Track 2 state is p_α (the initial state of the counter TM). Whereas $\bar{1}$ always transitions to $\bar{1}$ unless the Track 2 state is p_α . The following transition rules implement this:

$$\begin{array}{|c|c|} \hline \ominus \\ \hline 0 \\ \hline \end{array} \begin{array}{|c|c|} \hline p_\alpha \\ \hline p_\alpha \\ \hline \end{array} \longrightarrow \begin{array}{|c|c|} \hline \ominus \\ \hline \bar{1} \\ \hline \end{array} \begin{array}{|c|c|} \hline p_\alpha \\ \hline p_\alpha \\ \hline \end{array}, \quad \begin{array}{|c|c|} \hline \ominus \\ \hline \bar{1} \\ \hline \end{array} \begin{array}{|c|c|} \hline \neg p_\alpha \\ \hline \neg p_\alpha \\ \hline \end{array} \longrightarrow \begin{array}{|c|c|} \hline \ominus \\ \hline \bar{1} \\ \hline \end{array} \begin{array}{|c|c|} \hline \neg p_\alpha \\ \hline \neg p_\alpha \\ \hline \end{array}, \tag{69}$$

where $\neg p_\alpha$ here denotes any Track 2 state *other than* p_α .

7.3 Initialisation sweep

During the initial sweep of the Track 1 $\overset{\rightarrow}{0}$ from left to right and back, Track 1 contains a $\overset{\rightarrow}{0}$ or $\overset{\leftarrow}{0}$. We call this the *initialisation sweep*, and we want to use it to force Tracks 2 and 3 to be in the initial configurations

$$\langle \ominus p_0 \ominus^* \rangle \text{ and } \langle \ominus \#^* \rangle, \quad (70)$$

respectively. We do this by adding illegal pairs to forbid Track 2 from being anything other than \ominus when a $\overset{\rightarrow}{0}$ is over it (except at the beginning of the chain):

$$\begin{array}{|c|c|} \hline \neg \langle & \overset{\rightarrow}{0} \\ \hline \neg \ominus & \ominus \\ \hline \end{array}, \quad (71)$$

and similarly to forbid Track 3 from being anything other than $\#$ when a $\overset{\rightarrow}{0}$ is over it (except at the beginning of the chain):

$$\begin{array}{|c|c|} \hline \neg \langle & \overset{\rightarrow}{0} \\ \hline \neg \# & \# \\ \hline \end{array}, \quad (72)$$

where $\neg \odot$ again denotes anything *other than* the state \odot .

At the beginning of the chain, we use illegal pairs to forbid Track 2 from being anything other than p_α when $\overset{\leftarrow}{0}$ is over it. We also forbid it from being p_α when a $\overset{\leftarrow}{\vdash}$ is at the beginning of the chain:

$$\begin{array}{|c|c|} \hline \langle & \overset{\leftarrow}{0} \\ \hline p_\alpha & \neg p_\alpha \\ \hline \end{array}, \quad \begin{array}{|c|c|} \hline \langle & \overset{\leftarrow}{\vdash} \\ \hline p_\alpha & p_\alpha \\ \hline \end{array}. \quad (73)$$

Thus any standard basis state containing a $\overset{\rightarrow}{0}$, $\overset{\leftarrow}{0}$ or $\overset{\leftarrow}{\vdash}$ on Track 1 that matches the regular expressions (70), and is *not* in the initial configuration on Tracks 2 and 3, will evolve (either forwards or backwards) under the oscillator transition rules into an illegal configuration, in at most L steps.

If Tracks 2 and 3 are in their initial configurations, then a standard basis state containing a $\overset{\rightarrow}{0}$, $\overset{\leftarrow}{0}$ or $\overset{\leftarrow}{\vdash}$ will evolve backwards until it reaches the initial configuration $|\phi_0\rangle$, the standard basis state for Tracks 1 to 3 which has the form:

$$\begin{array}{|c|c|c|c|c|} \hline \langle & \overset{\rightarrow}{0} & \ominus & \dots & \ominus \\ \hline \langle & p_\alpha & \ominus & \dots & \ominus \\ \hline \langle & \vdash & \# & \dots & \# \\ \hline \end{array} \rangle. \quad (74)$$

7.4 Clock Counter

7.4.1 Counter TM construction

We construct the counter TM using a generalisation of the binary incrementing machine `INC` from Lemma 27, which will increment integers written in base- ζ instead of in binary. Constructing a base- ζ version of `INC` is largely a straightforward extension of Lemma 27, though ensuring reversibility is slightly subtle.

The following well-formed, normal-form, partial transition function implements a reversible TM that increments any little-endian base- ζ number written on the tape with a leading “start-of-tape” symbol \vdash , and returns the head to the starting cell without ever moving the head before the starting cell:

	\vdash	#	0	1	i	$\zeta - 1$
p_0	(\vdash, p_1, R)					
p_1		$(1, p'_2, R)$	$(1, p'_2, R)$	$(2, p_2, R)$	$(i + 1, p_2, R)$	$(0, p_1, R)$
p_2		$(\#, p_3, L)$	$(0, p_3, L)$	$(1, p_3, L)$	(i, p_3, L)	$(\zeta - 1, p_3, L)$
p'_2			$(0, p'_3, L)$	$(1, p'_3, L)$	(i, p'_3, L)	$(\zeta - 1, p'_3, L)$
p''_2		$(\#, p'_3, L)$				
p_3					(i, p_4, L)	$(\zeta - 1, p_4, L)$
p'_3				$(1, p_4, L)$		
p_4	(\vdash, p_f, N)		$(0, p_4, L)$			
p_f		$(\#, p_0, R)$	$(0, p_0, R)$	$(1, p_0, R)$	(i, p_0, R)	$(\zeta - 1, p_0, R)$

$$\forall i \in \{2, \dots, \zeta - 2\}$$

(75)

We now modify this machine in order to make it loop forever. We must be careful to ensure that we reenter the initial state reversibly. To do this, we introduce new initial and final states p_α and p_ω , and modify the p_0 and p_f transitions such that the TM reversibly transitions back into the state p_0 instead of halting, exploiting the fact that the symbol in the second tape cell will be blank in the first iteration,

and non-blank thereafter. The following set of transition rules accomplishes this:

	\vdash	$\#$	0	1	i	$\zeta - 1$
p_α	(\vdash, p_{-1}, R)					
p_{-1}		$(\#, p_0, L)$				
p_0	(\vdash, p_1, R)					
p_1		$(1, p'_2, R)$	$(1, p'_2, R)$	$(2, p_2, R)$	$(i + 1, p_2, R)$	$(0, p_1, R)$
p_2		$(\#, p_3, L)$	$(0, p_3, L)$	$(1, p_3, L)$	(i, p_3, L)	$(\zeta - 1, p_3, L)$
p'_2			$(0, p'_3, L)$	$(1, p'_3, L)$	(i, p'_3, L)	$(\zeta - 1, p'_3, L)$
p'_2		$(\#, p'_3, L)$				
p_3					(i, p_4, L)	$(\zeta - 1, p_4, L)$
p'_3				$(1, p_4, L)$		
p_4	(\vdash, p_f, N)		$(0, p_4, L)$			
p_f	(\vdash, p'_f, R)					
p'_f			$(0, p_0, L)$	$(1, p_0, L)$	(i, p_0, L)	$(\zeta - 1, p_0, L)$
p_ω	(\vdash, p_α, R)	$(\#, p_\alpha, R)$	$(0, p_\alpha, R)$	$(1, p_\alpha, R)$	(i, p_α, R)	$(\zeta - 1, p_\alpha, R)$

$\forall i \in \{2, \dots, \zeta - 2\}$

(76)

These transition rules implement a reversible base- ζ counter TM. When started from the tape configuration consisting of a \vdash symbol in the first cell followed by the all-blank tape (representing the number 0), this TM will loop indefinitely, incrementing the number written on the tape by 1 in each complete iteration. For brevity, we will refer to this particular initial tape configuration as the *standard input* to the counter TM tape.

We now declare certain configurations of the counter TM to be “illegal”. We will choose these illegal configurations to be such that they are never entered by the counter TM started from the standard input. Later on, when we come to encode the counter TM in a local Hamiltonian, we will use penalty terms to give energy penalties to all the configurations we declare “illegal” here, so that the corresponding standard basis states of the spin chain are indeed illegal in the sense of Definition 35. For now, however, the “illegal configurations” simply define a particular subset \mathcal{I} of the complete set of counter TM head and tape configurations.

Anticipating the later use of penalty terms, when defining illegal configurations we will often make use of the illegal pair notation introduced above, where the top row denotes the counter TM head position and internal state p , and the bottom row denotes the tape symbols ab in the section of tape near the head:

$$\begin{array}{|c|c|} \hline p & \cdot \\ \hline a & b \\ \hline \end{array} \quad \text{or} \quad \begin{array}{|c|c|} \hline \cdot & p \\ \hline a & b \\ \hline \end{array} . \quad (77)$$

A single row always denotes a section of tape:

$$\begin{array}{|c|c|} \hline a & b \\ \hline \end{array} . \quad (78)$$

Certain configurations that we declare to be illegal will specify that the counter TM head is or is not located at the starting cell. Again anticipating later use of penalty terms, we denote the starting cell by placing a \ominus symbol on the left. We denote tape cells *other than* the starting cell with $\neg\ominus$ on the left:

$$\begin{array}{|c|c|} \hline \ominus & p \\ \hline & a \\ \hline \end{array} \quad \begin{array}{|c|c|} \hline \neg\ominus & p \\ \hline & a \\ \hline \end{array} . \quad (79)$$

Note that the counter TM started from the standard input never moves its head before the starting cell, never reenters its initial state p_α , and never enters its final state p_ω . We therefore declare any counter TM configuration which is in state p_α with the head anywhere other than the starting cell to be illegal:

$$\begin{array}{|c|c|} \hline \neg\ominus & p_\alpha \\ \hline \cdot & \cdot \\ \hline \end{array} . \quad (80)$$

We also declare any configuration in state p_ω to be illegal:

$$\begin{array}{|c|c|} \hline p_\omega & \cdot \\ \hline \cdot & \cdot \\ \hline \end{array} . \quad (81)$$

Similarly, we declare any configuration in state p_α with the head adjacent to a non-blank symbol to be illegal, i.e. any configuration containing the illegal pair

$$\begin{array}{|c|c|} \hline p_\alpha & \cdot \\ \hline \cdot & \neg\# \\ \hline \end{array} . \quad (82)$$

The transition table in (76) is partial; some transitions are never used, so were not defined. For each (p, τ) for which no transition rule is defined in (76), we declare configurations in which the counter TM is in state p and the head is reading τ to be illegal, i.e. any configuration containing the illegal pair

$$\begin{array}{|c|} \hline p \\ \hline \tau \\ \hline \end{array} . \quad (83)$$

For each configuration (p_L, τ, L) , (p_R, τ, R) or (p_N, τ, N) which is not *entered* by any transition rule in (76), we declare the corresponding TM head and tape configurations to be illegal, i.e. any configurations containing one of the following illegal pairs:

$$\begin{array}{|c|c|} \hline \cdot & p_R \\ \hline \tau & \cdot \\ \hline \end{array} , \quad \begin{array}{|c|c|} \hline p_L & \cdot \\ \hline \cdot & \tau \\ \hline \end{array} , \quad \begin{array}{|c|c|} \hline p_N & \cdot \\ \hline \tau & \cdot \\ \hline \end{array} . \quad (84)$$

Note that during the evolution of the counter TM started from a blank tape, the head is never more than one cell to the right of a non-blank tape symbol. In fact,

the only moment at which it is over a blank symbol at all is when the incrementer TM needs to carry a digit, and moves the head to the next blank cell in order to write the carry. We therefore declare configurations in which the head is more than one cell to the right of a non-blank tape symbol to be illegal, i.e. all configurations containing the illegal pair

$$\begin{array}{|c|c|} \hline \cdot & p \\ \hline \# & \# \\ \hline \end{array} \cdot \quad (85)$$

Note also that, when started from the standard input, the counter TM never modifies the \vdash in the first cell, never writes a \vdash anywhere *other than* the first cell, and never creates an embedded blank symbol on the tape. We therefore define all tape configurations *without* a \vdash in the starting cell, all tape configuration with a \vdash anywhere other than the first cell, and all tape configurations with an embedded blank, to be illegal. These illegal TM configurations correspond exactly to tape configurations matching the regular expression

$$\vdash i^* \#, \quad i \in \{0 \dots \zeta - 1\}. \quad (86)$$

The little-endian numbers written by the counter TM are never padded with leading 0's, so the tape never has a 0 to the left of a $\#$. We therefore define this combination to be illegal, i.e. all tape configurations containing the illegal pair

$$\boxed{0 \#} \cdot \quad (87)$$

Later on, we will only be interested in configurations of the first L cells of the tape, and will stop the counter TM just before it tries to apply a transition rule that maps out of this portion of tape. Therefore, a 0 will never appear in the L 'th tape cell, and we declare configurations containing a 0 in this cell to be illegal:

$$\boxed{0 \triangleright} \quad (88)$$

Lemma 37 (Evolve-to-illegal) *Any counter TM configuration that is reached starting from the standard input is legal. All other configurations with the head in position $r \in [0, L]$ are either illegal, or evolve forwards or backwards to an illegal configuration within $O(L)$ time steps in such a way that the head never leaves the portion of tape $[0, L]$.*

Proof The first part of the Lemma is true by construction, since the partial transition rules of (76) implement the counter TM without using any of the undefined transitions, and the tape configurations defined to be illegal never occur when the counter TM is started from the standard input.

Recall that the alphabet of the counter TM is $\Xi := \{\vdash, \#\} \cup \{0, \dots, \zeta - 1\}$. Any tape configuration that does not start with a \vdash , or which contains a \vdash anywhere

other than the first cell, or which contains an embedded blank, or contains a 0 to the left of a #, or with a 0 in the L -th cell is illegal. So legal tape configurations match the regular expression

$$\vdash (\neg\#)^* \neg[0, \#]\#^*. \quad (89)$$

Configurations in which the head is more than one cell away from a non-blank symbol are illegal by (85), so all legal configurations have the head located in, or immediately adjacent to, the non-blank portion of the tape. All other configurations are illegal.

We divide the legal configurations into separate cases, according to the internal state of the counter TM:

State p_α : All configurations in state p_α are illegal by (80), except those with the head in the starting cell. The latter are illegal by (82) unless the tape cell to the right of the head is blank. The only such tape configuration which is legal is the standard input. Thus the only legal configuration in state p_α is the initial configuration of the counter TM started from the standard input.

State p_{-1} : There are no transitions out of (p_{-1}, x) where $x \neq \#$ in (76), so by (83) the only legal p_{-1} configurations have their head over a #. Moreover, there is only one transition into p_{-1} in (76). Hence, evolving any such configuration one step *backwards* according to (76) either enters a configuration that is illegal due to (83), or steps the head left and transitions to p_α . But we have shown that the only legal p_α configuration is the standard input. Thus any configuration in state p_{-1} not reachable starting from the standard input will evolve backwards to an illegal configuration in one time step.

State p_1 : Every TM configuration of the form

$$\begin{array}{c} p_1 \\ \boxed{\vdash \mid 0 \mid \cdots \mid 0 \mid \tau \mid \cdots} \\ \underbrace{\hspace{1.5cm}}_n \\ \tau \in \Xi / \{\vdash\}, \quad n \geq 0 \end{array}$$

with the rest of the tape compatible with (89) is reached when evolving the counter TM from the standard input; this configuration is produced whilst the counter TM is incrementing the number with little-endian base- ζ expansion

$$\underbrace{(\zeta - 1) \dots (\zeta - 1)}_n \tau \dots$$

The remaining legal configurations have at least one tape symbol x to the left of the head that is neither 0, nor #, nor \vdash :

$$\begin{array}{|c|c|c|c|c|c|c|c|c|c|} \hline & \vdash & i_1 & \cdots & i_m & x & 0 & \cdots & 0 & \tau & \cdots \\ \hline & & & & & & \underbrace{\hspace{2cm}}_n & & & & \\ \hline \end{array}$$

$$x \in \Xi / \{\vdash, \#, 0\}, \quad i_1, \dots, i_m \in \Xi / \{\vdash, \#\}, \quad \tau \in \Xi / \{\vdash\}, \quad m, n \geq 0.$$

Evolving any such configuration *backwards* according to (76) either enters into an illegal configuration due to (83), or rewinds the head left replacing 0's with $\zeta - 1$'s, eventually reaching the configuration with the head immediately to the right of the x :

$$\begin{array}{|c|c|c|c|c|c|c|c|c|c|} \hline & \vdash & i_1 & \cdots & i_m & x & \zeta - 1 & \cdots & \zeta - 1 & \tau & \cdots \\ \hline & & & & & & \underbrace{\hspace{2cm}}_n & & & & \\ \hline \end{array}$$

$$x \in \Xi / \{\vdash, \#, 0, 1\}, \quad i_1, \dots, i_m \in \Xi / \{\vdash, \#\}, \quad \tau \in \Xi / \{\vdash, \#\}, \quad m, n \geq 0.$$

The state p_1 can only be entered by moving right. But the configuration (p_1, x, R) is not entered by any transition rule in (76), so this configuration is illegal by (84).

State p_2 : The state p_2 is entered when incrementing any digit other than 0, # or $\zeta - 1$, so every TM configuration of the form

$$\begin{array}{|c|c|c|c|c|c|c|} \hline & \vdash & 0 & \cdots & 0 & x & \tau & \cdots \\ \hline & & \underbrace{\hspace{2cm}}_n & & & & & \\ \hline \end{array}$$

$$x \in \Xi / \{\vdash, \#, 0, 1\}, \quad \tau \in \Xi / \{\vdash\}, \quad n \geq 0.$$

with the rest of the tape compatible with (89) is reached when evolving the counter TM from the standard input. This configuration is produced whilst the counter TM is incrementing the number with little-edian base- ζ expansion

$$\underbrace{(\zeta - 1) \dots (\zeta - 1)}_n (x - 1) \tau \dots$$

The state p_2 can only be entered by moving right. Configurations (p_2, x, R) where $x \in \{\vdash, \#, 0, 1\}$ are not entered by any transition rule in (76), so these configurations are illegal by (84).

The remaining legal configurations have at least one tape symbol y to the left of the head that is neither 0 , $\#$, nor \vdash :

$$\begin{array}{|c|c|c|c|c|c|c|c|c|c|c|} \hline & & & & & & & & & & p_2 \\ \hline \vdash & i_1 & \dots & i_m & y & 0 & \dots & 0 & x & \tau & \dots \\ \hline & & & & & \underbrace{\hspace{2cm}}_n & & & & & \\ \hline \end{array}$$

$$x \in \Xi/\{\vdash, \#, 0, 1\}, \quad y \in \Xi/\{\vdash, \#, 0\}, \quad \tau \in \Xi/\{\vdash\}, \quad i_1, \dots, i_m \in \Xi/\{\vdash, \#\}, \quad m, n \geq 0.$$

Evolving any such configuration *backwards* one time step according to (76) either produces an illegal configuration due to (83), or produces the configuration:

$$\begin{array}{|c|c|c|c|c|c|c|c|c|c|c|} \hline & & & & & & & & & & p_1 \\ \hline \vdash & i_1 & \dots & i_m & y & 0 & \dots & 0 & x-1 & \tau & \dots \\ \hline & & & & & \underbrace{\hspace{2cm}}_n & & & & & \\ \hline \end{array}$$

$$x \in \Xi/\{\vdash, \#, 0, 1\}, \quad y \in \Xi/\{\vdash, \#, 0\}, \quad \tau \in \Xi/\{\vdash\}, \quad i_1, \dots, i_m \in \Xi/\{\vdash, \#\}, \quad m, n \geq 0.$$

But we have already shown above that all such p_1 configurations evolve backwards in time to an illegal configuration.

States p'_2, p''_2 : The argument for states p'_2 and p''_2 is very similar to that for p_2 , except that the legal value of x is now 1.

State p_3 : The state p_3 is entered by stepping left from p_2 , so every TM configuration of the form

$$\begin{array}{|c|c|c|c|c|c|} \hline & & & & & p_3 \\ \hline \vdash & 0 & \dots & 0 & \tau & \dots \\ \hline & & & \underbrace{\hspace{1cm}}_n & & \\ \hline \end{array}$$

$$\tau \in \Xi/\{\vdash, \#, 0, 1\}, \quad n \geq 0.$$

with the rest of the tape compatible with (89) is reached when evolving the counter TM from the standard input. This configuration is produced whilst the counter TM is incrementing the number with little-edian base- ζ expansion

$$\underbrace{(\zeta - 1) \dots (\zeta - 1)}_n (\tau - 1) \dots$$

There is no transition out of (p_3, x) where $\tau \in \{\vdash, \#, 0, 1\}$ in (76), so all such configurations are illegal by (83). The remaining legal configurations have at least

one tape symbol x to the left of the head that is neither 0, #, nor \vdash :

$$\begin{array}{c}
 \begin{array}{|c|c|c|c|c|c|c|c|c|}
 \hline
 \vdash & i_1 & \dots & i_m & x & 0 & \dots & 0 & \tau & \dots \\
 \hline
 \end{array} \\
 \underbrace{\hspace{10em}}_n \\
 \tau \in \Xi/\{\vdash, \#, 0, 1\}, x \in \Xi/\{\vdash, \#, 0\} \quad i_1, \dots, i_m \in \Xi/\{\vdash, \#\}, \quad m, n \geq 0.
 \end{array}$$

Evolving any such configuration forwards one time step according to (76) either produces an illegal configuration (for example if $n = 0$), or produces the configuration:

$$\begin{array}{c}
 \begin{array}{|c|c|c|c|c|c|c|c|c|}
 \hline
 \vdash & i_1 & \dots & i_m & x & 0 & \dots & 0 & \tau & \dots \\
 \hline
 \end{array} \\
 \underbrace{\hspace{10em}}_n \\
 \tau \in \Xi/\{\vdash, \#, 0, 1\}, x \in \Xi/\{\vdash, \#, 0\} \quad i_1, \dots, i_m \in \Xi/\{\vdash, \#\}, \quad m, n \geq 0.
 \end{array}$$

We show below that any such p_4 configuration evolves to an illegal configuration.

State p'_3 : The argument for state p'_3 is very similar to that for p_3 , except that $\tau = 1$ in this case.

State p_4 : In a legal state p_4 configuration, the head must be over a 0 or a \vdash , since there is no transition out of (p_4, x) for $x \notin \{0, \vdash\}$. We first analyse the case in which it is over a 0.

Any TM configuration of the form

$$\begin{array}{c}
 \begin{array}{|c|c|c|c|c|}
 \hline
 \vdash & 0 & \dots & 0 & \dots \\
 \hline
 \end{array} \\
 \underbrace{\hspace{4em}}_n \\
 n \geq 1
 \end{array}$$

with the rest of the tape compatible with (89) is reached by evolving the counter TM from the standard input.

The remaining legal configurations have at least one tape symbol x to the left of the head that is neither 0, #, nor \vdash :

$$\begin{array}{c}
 \begin{array}{|c|c|c|c|c|c|c|c|c|}
 \hline
 \vdash & i_1 & \dots & i_m & x & 0 & \dots & 0 & \dots \\
 \hline
 \end{array} \\
 \underbrace{\hspace{10em}}_n \\
 x \in \Xi/\{\vdash, \#, 0\}, \quad i_1, \dots, i_m \in \Xi/\{\vdash, \#\}, \quad m \geq 0, n \geq 1
 \end{array}$$

Evolving any such configuration according to (76) steps the head left over the 0's until the head is over the x :

$$\begin{array}{c}
 p_4 \\
 \boxed{\vdash \mid i_1 \mid \dots \mid i_m \mid x \mid 0 \mid \dots \mid 0 \mid \dots} \\
 \underbrace{\hspace{10em}}_n \\
 x \in \Xi / \{\vdash, \#\}, \quad i_1, \dots, i_m \in \Xi / \{\vdash, \#\}, \quad m \geq 0, n \geq 1.
 \end{array}$$

But there is no transition out of (p_4, x) in (76), so this configuration is illegal by (83).

We now analyse the case in which p_4 is over \vdash . Evolving one step backwards according to (76) steps the head right and either enters an illegal configuration, transitions to state p_4 with the had over a 0, or transitions to states p_3 or p'_3 . In all three cases, we have already shown that all such configurations not reachable from the standard input evolve to illegal configurations.

State p_f : There are no transitions out of (p_f, x) where $x \neq \vdash$ in (76), so the only legal p_f configurations have the head over a \vdash . The latter can only appear in the starting cell in legal configurations. Evolving any such configuration backwards according to (76) for one time step transitions to the state p_4 , and we have already shown that all p_4 configurations not reachable from the standard input evolve to illegal configurations.

State p'_f : There are no transitions out of (p'_f, x) where $x \in \{\vdash, \#\}$ in (76), so the only legal p'_f configurations have the head over a symbol in the set $\{0, \dots, \zeta - 1\}$. Evolving any such configuration backwards according to (76) for one time step transitions to the state p_f , and we have already shown that all p_f configurations not reachable from the standard input evolve to illegal configurations.

State p_0 : There are no transitions out of (p_0, x) where $x \neq \vdash$ in (76), so the only legal p_0 configurations have the head over a \vdash . The latter can only appear at the beginning of the tape in legal configurations. Since the only transitions in (76) into p_0 are from p_1 or p'_f , evolving any such configuration one step backwards according to (76) either enters an illegal configuration, or steps the head right into state p_{-1} or p'_f . In both cases we have already shown that configurations not reachable from the standard input evolve to illegal ones.

State p_ω : The counter TM never enters the state p_ω , and all such configurations are illegal. \square

7.4.2 Counter TM Hamiltonian

We want all legal Track 2 configurations to contain a single state $p \in P'$ marking the location of the counter TM head, and blanks everywhere else. I.e. the Track 2 configuration should match the regular expression $\langle \circ^* p \circ^* \rangle$ where $p \in P'$. By Lemma 36, we can enforce this using penalty terms. This Track 2 regular expression allows us to identify Tracks 2 and 3 with the configurations of the counter TM (with internal state p , its position being the head location, and Track 3 the tape). This will be explicitly or implicitly exploited in all that follows.

To encode the evolution of this counter TM in our local Hamiltonian, we encode its transition rules in transition terms. However, the transition rules will only apply when a $\overleftrightarrow{\tau}$ arrow on Track 1 sweeps past the counter TM head encoded on Track 2, so that the counter TM advances exactly one step for each complete left-to-right sweep of the $\overleftrightarrow{\tau}$ arrow. Since we are restricting to two-body interactions and the arrow is moving to the right, transitions in which the head moves left must be triggered when the $\overleftrightarrow{\tau}$ is to the left of the TM head, in order to both move the arrow and update the head location. This way of implementing the left transitions means that

$$\begin{array}{|c|} \hline \overleftrightarrow{\tau} \\ \hline p \\ \hline \sigma \\ \hline \end{array} \quad (90)$$

never occurs during the evolution of the counter TM if $\delta(p, \sigma) = (\tau, p_L, L)$. We add a penalty term to make such configurations illegal.

Transitions in which the head moves to the right or stays still will be triggered when the $\overleftrightarrow{\tau}$ is on top of the head. In order to avoid two transitions in which the head moves right from being triggered during the same sweep of the $\overleftrightarrow{\tau}$, we must split these latter transitions into two stages, in which we first transition into an auxiliary state $p'_R \in P'_R$, and then transition into the correct state $p_R \in P_R$ in the following time step (adapting the construction of [GI09]). Thus, when the $\overleftrightarrow{\tau}$ is on top of the the head and the head will move right, we will update the tape, move the head, and step the $\overleftrightarrow{\tau}$ to the right. But we transition into the auxiliary state $p'_R \in P'_R$ instead of $p_R \in P_R$. In the next step, with the $\overleftrightarrow{\tau}$ now on top of the p'_R state, we transition to the correct state p_R , and step the $\overleftrightarrow{\tau}$ to the right once more so that it is no longer on top of the head. This way of carrying out the right-moving transitions means that p'_R only ever appears below a $\overleftrightarrow{\tau}$, so we add a penalty term to make all other p'_R configurations illegal:

$$\begin{array}{|c|} \hline \neg \overleftrightarrow{\tau} \\ \hline p'_R \\ \hline \end{array} \quad (91)$$

For TM transition rules $\delta(p, \sigma) = (\tau, p_N, N)$, $\delta(p, \sigma) = (\tau, p_L, L)$ and $\delta(p, \sigma) = (\tau, p_R, R)$ (where $p \in P \setminus \{p_\sigma\}$, $p_{L,N,R} \in P_{L,N,R}$ and $\sigma, \tau \in \Xi$), the following local transition terms on Tracks 1 to 3 implement the desired transformations (cells

marked \cdot can be in any state, and are left unchanged by the transition):

$$\begin{array}{|c|c|} \hline \textcircled{\uparrow} & \textcircled{\circ} \\ \hline p & \textcircled{\circ} \\ \hline \sigma & \cdot \\ \hline \end{array} \rightarrow \begin{array}{|c|c|} \hline \textcircled{\circ} & \textcircled{\uparrow} \\ \hline p_N & \textcircled{\circ} \\ \hline \tau & \cdot \\ \hline \end{array}, \quad \begin{array}{|c|} \hline \textcircled{\uparrow} \\ \hline p \\ \hline \sigma \\ \hline \end{array} \rightarrow \begin{array}{|c|} \hline \textcircled{\uparrow} \\ \hline p_N \\ \hline \tau \\ \hline \end{array}, \quad (92a)$$

$$\begin{array}{|c|c|} \hline \textcircled{\uparrow} & \textcircled{\circ} \\ \hline \textcircled{\circ} & p \\ \hline \cdot & \sigma \\ \hline \end{array} \rightarrow \begin{array}{|c|c|} \hline \textcircled{\circ} & \textcircled{\uparrow} \\ \hline p_L & \textcircled{\circ} \\ \hline \cdot & \tau \\ \hline \end{array}, \quad \begin{array}{|c|c|} \hline \textcircled{\uparrow} & \textcircled{\circ} \\ \hline p & \textcircled{\circ} \\ \hline \sigma & \cdot \\ \hline \end{array} \rightarrow \begin{array}{|c|c|} \hline \textcircled{\circ} & \textcircled{\uparrow} \\ \hline \textcircled{\circ} & p'_R \\ \hline \tau & \cdot \\ \hline \end{array}, \quad (92b)$$

$$\begin{array}{|c|c|} \hline \textcircled{\uparrow} & \textcircled{\circ} \\ \hline p'_R & \textcircled{\circ} \\ \hline \cdot & \cdot \\ \hline \end{array} \rightarrow \begin{array}{|c|c|} \hline \textcircled{\circ} & \textcircled{\uparrow} \\ \hline p_R & \textcircled{\circ} \\ \hline \cdot & \cdot \\ \hline \end{array}, \quad \begin{array}{|c|} \hline \textcircled{\uparrow} \\ \hline p'_R \\ \hline \cdot \\ \hline \end{array} \rightarrow \begin{array}{|c|} \hline \textcircled{\uparrow} \\ \hline p_R \\ \hline \cdot \\ \hline \end{array}. \quad (92c)$$

Transitions in which the head moves left from the end of the chain are already covered by the left-moving transition in (92b). Transitions in which the head would move right off the end of the chain are not implemented. (Transitions in which the head would move left off the beginning of the chain are not implemented either, but in fact the counter TM construction of Section 7.4 will never attempt such a transition when initialised properly. Moreover, we just declared such configurations illegal in (90).)

When started from the blank input (which represents the number 0), the encoded counter TM will loop ζ^{L-3} times before it exceeds the $L - 2$ tape space available on Track 2 (the space overhead of 1 is due to the way the incremter TM in implemented in (75)). The incremter machine takes $\Theta(\log x)$ steps to increment the number x (written on the tape in base- ζ). Thus the counter TM will run for at least $\Omega(\zeta^L)$ (and at most $O(\log L \zeta^L)$) time-steps before it exceeds the available tape space.

The transition function for the counter TM defines a transition rule for each pair $(p, \tau) \in P \times \Xi$. So for any legal configuration containing a $p \in P$ on Track 2, exactly one of the transition rules in (92a) and (92b) will apply during each $\textcircled{\uparrow}$ sweep. Similarly, for any configuration containing a $p'_R \in P'_R$, a (92c) rule will apply. It is easy to verify that if any of the transition rules from (92a)–(92c) is applied to any legal configuration, none of the (92a) and (92b) rules can apply again to the resulting configuration. If the final (92b) rule applies, the (92c) rule will apply exactly once in the following step, after which none of the rules can apply again to the resulting configuration. The transition rules in (92a)–(92c) therefore implement a single step of the TM during each left-to-right sweep of the $\textcircled{\uparrow}$ state on Track 1, as required. We call this phase of the evolution, in which Track 1 contains a $\textcircled{\uparrow}$ or $\textcircled{\circ}$, the *computation phase*.

We also modify the clock oscillator rules from Section 7.2 involving a right-moving $\overrightarrow{\tau}$ arrow to only apply when there is no TM head under the $\overrightarrow{\tau}$, since the $\overrightarrow{\tau}$ movement is already taken care of by the above transition rules when there is a TM head present:

$$\begin{array}{|c|c|} \hline \overrightarrow{\tau} & \odot \\ \hline \circ & \circ \\ \hline \end{array} \rightarrow \begin{array}{|c|c|} \hline \circ & \overrightarrow{\tau} \\ \hline \circ & \circ \\ \hline \end{array}, \quad \begin{array}{|c|c|} \hline \overrightarrow{\tau} & \odot \\ \hline \odot & \odot \\ \hline \end{array} \rightarrow \begin{array}{|c|c|} \hline \circ & \overrightarrow{\tau} \\ \hline \odot & \odot \\ \hline \end{array}, \quad \begin{array}{|c|} \hline \overrightarrow{\tau} \\ \hline \odot > \\ \hline \end{array} \rightarrow \begin{array}{|c|} \hline \overrightarrow{\tau} \\ \hline \odot > \\ \hline \end{array}. \quad (93)$$

We must also enforce illegality of all the configurations declared to be illegal in Section 7.4.1. To do this, we simply introduce local penalty terms corresponding to the combinations declared illegal in (80)–(85), (87) and (88), and enforce the Track 3 regular expression of (86) using penalty terms and Lemma 36.

The complete set of Track 1 to 3 transition rules is summarised in Table 1. The regular expressions on Tracks 1 to 3 enforced by penalty terms, together with all the additional illegal pairs defined so far, are summarised in Table 2.

Table 1: All transition rules for Tracks 1 to 3.

Track 1 rules		
$\overrightarrow{\tau}_0, \overleftarrow{\tau}_0$ rules	$\overrightarrow{\tau}$ rules	$\overleftarrow{\tau}$ rules
$\begin{array}{ c c } \hline \overrightarrow{\tau}_0 & \odot \\ \hline \circ & \circ \\ \hline \end{array} \rightarrow \begin{array}{ c c } \hline \circ & \overrightarrow{\tau}_0 \\ \hline \circ & \circ \\ \hline \end{array}$		$\begin{array}{ c c } \hline \circ & \overleftarrow{\tau} \\ \hline \circ & \odot \\ \hline \end{array} \rightarrow \begin{array}{ c c } \hline \overleftarrow{\tau} & \odot \\ \hline \circ & \circ \\ \hline \end{array} \quad (*)^1$
$\begin{array}{ c c } \hline \overrightarrow{\tau}_0 & > \\ \hline \circ & \circ \\ \hline \end{array} \rightarrow \begin{array}{ c c } \hline \overleftarrow{\tau}_0 & > \\ \hline \circ & \circ \\ \hline \end{array}$		
$\begin{array}{ c c } \hline \circ & \overleftarrow{\tau}_i \\ \hline \circ & \circ \\ \hline \end{array} \rightarrow \begin{array}{ c c } \hline \overleftarrow{\tau}_{i+1} & \odot \\ \hline \circ & \circ \\ \hline \end{array}$		
$\begin{array}{ c c } \hline \circ & \overleftarrow{\tau}_k \\ \hline \circ & \circ \\ \hline \end{array} \rightarrow \begin{array}{ c c } \hline \overleftarrow{\tau}_0 & \odot \\ \hline \circ & \circ \\ \hline \end{array}$		






¹Transition rules marked with an (*) will be replaced in Section 7.6 by a set of rules that also act on Tracks 4, 5 and 6.

Track 1 and 2 rules

$\overset{\leftarrow}{0}, \overset{\leftarrow}{0}$ rules

$\overset{\leftarrow}{1}$ rules

$\overset{\leftarrow}{1}$ rules



Track 1, 2 and 3 rules

$\overset{\rightarrow}{\circ}, \overset{\leftarrow}{\circ}$ rules

$\overset{\uparrow}{\tau}$ rules

$\overset{\downarrow}{\tau}$ rules










Table 2: All illegal pairs and regular expressions enforced by illegal pairs for Tracks 1 to 3.

Regular expressions			
Track 1	Track 2	Track 3	
$\langle \circ \circ^* [\overset{\rightarrow}{0}, \overset{\leftarrow}{0}, \overset{\leftarrow}{i}, \overset{\rightarrow}{1}, \overset{\leftarrow}{1}] \circ \circ^* \rangle$			
$\langle \circ \circ^* p \circ \circ^* \rangle$			
$\langle \circ \vdash i^* \# \circ \rangle$			
Illegal pairs			
Tracks 1 and 2			
Illegal pairs			
Tracks 1 and 3	Track 2	Tracks 2 and 3	Track 3
Illegal pairs			
Tracks 2 and 3 for undefined transitions in (76)			

Illegal pairs

Tracks 1, 2 and 3 if $\delta(p, \sigma) = (p_L, \tau, L)$ (90)



7.5 Clock Hamiltonian

Before considering the remaining tracks, it is helpful to analyse in depth the Hamiltonian defined so far on the first three tracks. For that we start with the definition of well-formed states.

Definition 38 (Well-formed state) *We say that a standard basis state on Tracks 1 to 3 is well-formed if it is a bracketed state, its Track 1 configuration matches the regular expression $\langle \circ \circ^* [\overset{\leftarrow}{\circ} \overset{\leftarrow}{\circ} \overset{\rightarrow}{\circ} \overset{\leftarrow}{\tau} \overset{\leftarrow}{\tau}] \circ \circ^* \rangle$, and its Track 2 configuration matches the regular expression $\langle \circ \circ^* p \circ \circ^* \rangle$ where $p \in P'$.*

The penalty terms defined previously give an energy penalty of at least 1 to any standard basis state in S_{br} that does not match the desired regular expressions on Tracks 1–3. Thus only well-formed states can have zero energy. The following result shows that only one transition rule can apply to each well-formed state.

Lemma 39 (Well-formed transitions) *For any well-formed standard basis state, at most one transition rule applies in the forward direction, and at most one in the backwards direction. Furthermore, the set of well-formed states is closed under the transition rules.*

Proof A well-formed standard basis state contains exactly one of $\overset{\rightarrow}{\circ}$, $\overset{\leftarrow}{\circ}$, $\overset{\leftarrow}{\circ}$, $\overset{\leftarrow}{\tau}$ or $\overset{\leftarrow}{\tau}$ on Track 1. Thus clearly only transition rules from one of the sets in Table 1 (the set of $\overset{\rightarrow}{\circ}$, $\overset{\leftarrow}{\circ}$ rules, $\overset{\leftarrow}{\tau}$ rules, or $\overset{\leftarrow}{\tau}$ rules) can apply in the forward direction. The left hand sides of the rules within the $\overset{\rightarrow}{\circ}$, $\overset{\leftarrow}{\circ}$ and $\overset{\leftarrow}{\tau}$ sets are manifestly mutually exclusive. Since the counter TM is deterministic, there is a unique TM transition that applies at any step, so the left-hand-sides of the $\overset{\leftarrow}{\tau}$ rules are also mutually exclusive.

The counter TM is reversible, so there is also a unique backwards TM transition at every step. Thus the same argument applies in the backwards direction to the right hand sides of the rules in each set, with the exception of the $\overset{\leftarrow}{\circ}$ turning rule that changes a $\overset{\leftarrow}{\tau}$ into a $\overset{\leftarrow}{\circ}$ (in the backwards direction). But this rule clearly

cannot apply at the same time as the right hand side of any rule from the $\textcircled{\uparrow}$ set, which concludes the proof of the first part of the lemma.

It is straightforward to verify that all the transition rules in Table 1 preserve well-formedness, implying the second part. \square

Recall that the state $|\phi_0\rangle$ is the standard basis state with $\textcircled{\rightarrow}_0 \textcircled{\circ}^* \textcircled{\triangleright}$ on Track 1, $\textcircled{\leftarrow} p_\alpha \textcircled{\circ}^* \textcircled{\triangleright}$ on Track 2, and $\textcircled{\leftarrow} \vdash \#^* \textcircled{\triangleright}$ on Track 3. This corresponds to having the clock oscillator in the first state of the Track 1 sequence (see Figure 8), and the counter TM Tracks 1 and 2 in their initial configurations. There is no backwards transition out of $|\phi_0\rangle$, so we will refer to it as the *initial clock state*. Let $|\phi_t\rangle$ be the standard basis state obtained by applying t transitions (in the forwards direction) to $|\phi_0\rangle$, which is well-defined thanks to Lemma 39.

Let us consider the form of the states $|\phi_t\rangle$. Note that $|\phi_0\rangle$ contains a $\textcircled{\rightarrow}_0$ on Track 1. The only transition rules that apply when Track 1 contains a $\textcircled{\rightarrow}_0$, $\textcircled{\leftarrow}_0$ or $\textcircled{\leftarrow}_{(i)}$ (see Table 1) are those that sweep the arrow all the way to the right and back again, without affecting Tracks 2 or 3. We saw in Section 7.2 that one complete oscillator cycle takes $2(L-2)$ steps. Thus $t \leq 2(L-2)$ corresponds to the initialisation sweep, in which the state $|\phi_t\rangle$ contains a $\textcircled{\rightarrow}_0$, $\textcircled{\leftarrow}_0$ or $\textcircled{\leftarrow}_{(i)}$ on Track 1, and Tracks 2 and 3 remain in their initial configurations.

At $t = 2(L-2) + 1$, the $\textcircled{\leftarrow}_0$ turns around and (because Track 2 contains a p_α at the beginning of the chain) becomes a $\textcircled{\uparrow}$. This $\textcircled{\uparrow}$ is now over a p_α on Track 2. Thus the next step implements the first step of the counter TM on Tracks 2 and 3, causing Track 2 to transition to p_α .

Recall from Section 7.4 that the transition rules implement one step of the counter TM during each left-to-right sweep of the $\textcircled{\uparrow}$ on Track 1, and then move the $\textcircled{\uparrow}$ back to the beginning of the chain in the second half of the oscillator cycle. Note that the counter TM construction of Section 7.4 never transitions back to its initial state p_α . So when the $\textcircled{\uparrow}$ reaches the beginning of the chain, it will never find a p_α on Track 3. Thus it will always transition into $\textcircled{\uparrow}$, never into $\textcircled{\rightarrow}_0$ (see Table 1). Therefore, for all $t > 2(L-2)$ the state $|\phi_t\rangle$ contains a $\textcircled{\uparrow}$ or $\textcircled{\leftarrow}_{(i)}$ on Track 1. At $t = 2(L-2)(n+1)$ the $\textcircled{\uparrow}$ in $|\phi_t\rangle$ is back at the beginning of the chain, and Tracks 2 and 3 are in the configuration corresponding to the n 'th step of the counter TM.

Eventually, the number written on Track 3 will be incremented until it reaches ζ^{L-4} , the maximum number that can be written in base- ζ in $L-4$ digits (the -4 accounts for the space-overhead of 2 required by the Inc TM implementation, plus the two bracket states at the ends of the chain). When the counter TM tries to increment this number, the head will move right until it reaches the cell adjacent to the $\textcircled{\triangleright}$ at the end of the chain. At this point, it will be in an internal state that would normally step right in order to write the carry. But the site to the right of the head contains $\textcircled{\triangleright}$ instead of a blank Tape 2 cell. When the $\textcircled{\uparrow}$ on Track 1 reaches the head at the end of on Track 2, there is no further forwards transition out of

this configuration (see Table 1). Let T be the number of transitions required to reach this configuration $|\phi_T\rangle$ starting from the initial state $|\phi_0\rangle$. We will refer to $|\phi_T\rangle$ as the *final clock state*. Since this configuration occurs when the counter TM is about to exceed L tape space, which occurs after at most $O(\log L\zeta^L)$ steps of the TM (see Section 7.4) each of which requires $2(L-2)$ transitions, we have that $T = O(L \log L\zeta^L)$. (Though this bound on the run-time will not be important for our purposes.)

The following result will be important later when we come to analyse the Hamiltonian.

Lemma 40 (Evolve-to-illegal) *Evolving any $|\phi_i\rangle$ forwards or backwards in time according to the transition rules will never reach an illegal configuration. All other well-formed standard basis states will evolve either forwards or backwards to an illegal configuration after $O(L^2)$ transitions.*

Proof As in Lemma 37 the first part of the Lemma is true by construction. For the second part, let us analyse all possible well-formed standard basis states, which we divide into two cases depending on the type of arrow state we have on Track 1:

Case 1: Track 1 contains $\overset{\rightarrow}{\circ}$, $\overset{\leftarrow}{\circ}$ or $\overset{\leftarrow}{\ominus}$. If Tracks 2 and 3 are in the initial configuration, then we are in one of the $|\phi_t\rangle$ states with $0 \leq t \leq 2(L-2)$.

If Tracks 2 and 3 are not in the initial configuration, then within at most $O(L)$ steps the $\overset{\rightarrow}{\circ}$ or $\overset{\leftarrow}{\circ}$, $\overset{\leftarrow}{\ominus}$ will move forwards or backwards along the chain due to the initialisation sweep clock oscillator transition rules, until it is over a site containing the wrong initial Track 2 or Track 3 state. But this is illegal due to the Track 1–3 illegal pairs from Table 2.

Case 2: Track 1 contains $\overset{\uparrow}{\ominus}$ or $\overset{\downarrow}{\ominus}$. Recall that there is a one-to-one correspondence between configurations of the counter TM and well-formed standard basis states that do not contain a $p'_R \in P'_R$ on Track 2.

Well-formed configurations containing a $p'_R \in P'_R$ on Track 2 evolve by a single backwards transition from (92b) to a configuration containing a $p \notin P'_R$ on Track 2, and evolve by a single forwards transition from (92c) to a configuration containing a $p_R \in P_R$. A configuration containing $p'_R \in P'_R$ is therefore reachable (resp. unreachable) by transition rules from the initial configuration iff the corresponding counter TM configuration with p'_R replaced by the corresponding p_R is reachable (resp. unreachable).

For any well-formed configuration, each sweep of $\overset{\uparrow}{\ominus}$, the transition rules implements exactly one step of the counter TM encoded in Tracks 2 and 3. The only counter TM transitions that are not implemented are those that would move the head beyond the L sites between the \ominus and \ominus . Therefore, if the configuration

of Tracks 2 and 3 is not reachable from the initial configuration, then Lemma 37 guarantees that evolving it forwards or backwards by the transition rules will reach an illegal configuration in $O(L)$ steps of the counter TM, which takes $O(L^2)$ transitions. (Note that it is crucial for this argument that to reach an illegal configuration in Lemma 37 never requires moving the head outside the portion of tape $[0, L]$.)

It only remains to consider well-formed states with Track 2 and 3 configurations that are reachable from the initial configuration. For the initial Tracks 2–3 configuration itself, if $\textcircled{\uparrow}$ is in the starting cell then we have the state $|\phi_t\rangle$ for $t = 2(L - 2) + 1$. If not, either we have $\textcircled{\uparrow}$ to the right of the head or $\textcircled{\uparrow}$ on Track 1, with p_α at the start of Track 2. In both cases, evolving such a state forwards in time will reach a state with a $\textcircled{\uparrow}$ over the p_α within $O(L)$ transitions, which is illegal by the second Track 1 and 2 illegal pair from Table 2.

If the Track 2–3 configuration is not the initial one, then all positions of $\textcircled{\uparrow}$ and $\textcircled{\downarrow}$ give rise to elements in the set $\{|\phi_t\rangle\}_t$ except exactly those declared illegal in (90) and (91).

This completes the proof of the Lemma. \square

7.6 QTM Hamiltonian

With the clock in place, encoding an arbitrary quantum Turing Machine in the Hamiltonian is now possible using ideas from [GI09] (and is similar to the construction used in Section 7.4 to encode the reversible counter TM). The configurations are now quantum states, so there can be arbitrary superpositions over standard basis states. But it suffices to verify that the construction does the right thing on standard basis states; this then extends to arbitrary superpositions by linearity and well-formedness.

When we later come to analyse the spectrum, we will require our Hamiltonian to be *standard-form* in the following sense:

Definition 41 (Standard-form Hamiltonian) *We say that a Hamiltonian $H = H_{\text{trans}} + H_{\text{pen}}$ acting on a Hilbert space $\mathcal{H} = (\mathbb{C}^C \otimes \mathbb{C}^Q)^{\otimes L} = (\mathbb{C}^C)^{\otimes L} \otimes (\mathbb{C}^Q)^{\otimes L} =: \mathcal{H}_C \otimes \mathcal{H}_Q$ is of standard form if $H_{\text{trans,pen}} = \sum_{i=1}^{L-1} h_{\text{trans,pen}}^{(i,i+1)}$, and $h_{\text{trans,pen}}$ satisfy the following conditions:*

(i). $h_{\text{trans}} \in \mathcal{B}((\mathbb{C}^C \otimes \mathbb{C}^Q)^{\otimes 2})$ is a sum of transition rule terms, where all the transition rules act diagonally on $\mathbb{C}^C \otimes \mathbb{C}^C$ in the following sense. Given standard basis states $a, b, c, d \in \mathbb{C}^C$, exactly one of the following holds:

- there is no transition from ab to cd at all; or
- $a, b, c, d \in \mathbb{C}^C$ and there exists a unitary U_{abcd} acting on $\mathbb{C}^Q \otimes \mathbb{C}^Q$ together with an orthonormal basis $\{|\psi_{abcd}^i\rangle\}_i$ for $\mathbb{C}^Q \otimes \mathbb{C}^Q$, both depending only on a, b, c, d , such that the transition rules from ab to cd appearing in h_{trans} are exactly $|ab\rangle |\psi_{abcd}^i\rangle \rightarrow |cd\rangle U_{abcd} |\psi_{abcd}^i\rangle$ for all i .

(ii). $h_{pen} \in \mathcal{B}((\mathbb{C}^C \otimes \mathbb{C}^Q)^{\otimes 2})$ is a sum of penalty terms which act non-trivially only on $(\mathbb{C}^C)^{\otimes 2}$ and are diagonal in the standard basis.

In our case, \mathbb{C}^C corresponds to Tracks 1–3 and \mathbb{C}^Q to Tracks 4–6. Transitions from ab to cd will correspond exactly to the transitions shown for Tracks 1–3 in Table 1. While adding new transitions involving Tracks 4–6, we will need to remove some of the transitions in Tracks 1–3, but they will be recovered as restrictions of the new rules to those tracks. (In Table 1, transitions that will be replaced later are marked with an asterisk.) We will define the unitaries U_{abcd} partially, and then complete them to a full unitary. All we have to take care of is that (i) the new transition involving the rest of the tracks, when restricted to Tracks 1–3, recover exactly the transitions removed from Table 1; and that (ii) orthogonality is preserved in the construction.

By the standard form,

$$h_{\text{trans}} = \sum_{ab \rightarrow cd} (|ab\rangle - |cd\rangle)(\langle ab| - \langle cd|) \otimes (2\mathbb{1} - U_{abcd} - U_{abcd}^\dagger). \quad (94)$$

Now, if we start with a family of QTM P_n which satisfies part (iii) of Theorem 11, it will be immediate from our construction that the partial definition of all U_{abcd} will only involve elements in the set

$$\mathcal{S} = \left\{ 0, 1, \pm \frac{1}{\sqrt{2}}, e^{i\pi\varphi}, e^{i\pi 2^{-|\varphi|}} \right\}. \quad (95)$$

As argued above, the same will hold for the full U_{abcd} , which gives part (vi) in Theorem 33.

7.6.1 QTM transition rules

We first analyse how to incorporate the QTM transition rules into the Hamiltonian. We will simulate the QTM transition rules $\delta(p, \sigma) = \sum_{\tau, q, D} \delta(p, \sigma, \tau, q, D) |\tau, q, D\rangle$ during the second half of the clock oscillator cycle, when the Track 1 \oplus state sweeps from right to left. This is done by including the following transition rule terms (cf. (92a)–(92c)), where q is any state in $Q/\{q_f\}$ (i.e. we exclude all the

transitions out of the QTM's final state):

$$\begin{aligned}
 \left| \begin{array}{cc} \circ & \ominus \\ q & \circ \\ \sigma & \cdot \end{array} \right\rangle &\rightarrow \sum_{\tau, q_R} \delta(q, \sigma, \tau, q_R, R) \left| \begin{array}{cc} \ominus & \circ \\ \circ & q_R \\ \tau & \cdot \end{array} \right\rangle + \sum_{\tau, q_N} \delta(q, \sigma, \tau, q_N, N) \left| \begin{array}{cc} \ominus & \circ \\ q_N & \circ \\ \tau & \cdot \end{array} \right\rangle \\
 &+ \sum_{\tau, q_L} \delta(q, \sigma, \tau, q_L, L) \left| \begin{array}{cc} \ominus & \circ \\ q'_L & \circ \\ \tau & \cdot \end{array} \right\rangle,
 \end{aligned} \tag{96a}$$

$$\left| \begin{array}{cc} \circ & \ominus \\ \circ & q'_L \\ \cdot & \cdot \end{array} \right\rangle \rightarrow \left| \begin{array}{cc} \ominus & \circ \\ q_L & \circ \\ \cdot & \cdot \end{array} \right\rangle. \tag{96b}$$

As in the implementation of the reversible counter TM (Section 7.4), these transition rules implement the **R**eft-moving, **N**on-moving and **L**ight-moving transitions separately. This time, since the \ominus arrow is sweeping to the left, it is the *left*-moving transitions that are implemented in two stages, first transitioning to an auxiliary state $q'_L \in Q'_L$, then in a second step transitioning to the corresponding q_L state.

Note that if the QTM head ever ends up at the end of the chain, the transition rules of (96) can never apply, because they require the \ominus to be to the right of the QTM head. Thus the final Track 5 tape cell is not used, and the effective tape length available to the QTM for a chain of length $L + 3$ is only L rather than¹ $L + 1$. If the Track 4 QTM head ends up next to the \ominus at the very end of the chain, this indicates that the head has stepped off the usable portion of the tape, regardless of the internal state.

We also include the following transition rules involving a left-moving \ominus arrow to cover the case where there is no QTM head present:

$$\begin{array}{ccc}
 \begin{array}{cc} \circ & \ominus \\ \circ & \circ \end{array} & \rightarrow & \begin{array}{cc} \ominus & \circ \\ \circ & \circ \end{array}, & \begin{array}{cc} \circ & \ominus \\ \circ & \circ \end{array} & \rightarrow & \begin{array}{cc} \ominus & \circ \\ \circ & \circ \end{array}, & \begin{array}{cc} \circ & \ominus \\ \circ & q \end{array} & \rightarrow & \begin{array}{cc} \ominus & \circ \\ \circ & q \end{array},
 \end{array} \tag{97}$$

7.6.2 QTM Unitarity

With the set of transition rule terms defined so far in (96), there is no transition out of the final state q_f of the QTM. There is also no transition out of a QTM configuration in which the QTM head is at the end of the chain, and the next step would move the head to the right, off the end of the chain. The same is true of

¹Two cells are always used by the \ominus , \otimes markers.

configurations trying to move the head left when it is located at the beginning of the chain.¹

Ideally, we would like it if, when the QTM reaches one of these configurations, the clock continued ticking but the encoded QTM did nothing for the remaining time steps. However, simply omitting transitions out of these configurations causes a subtle but important issue relating to unitarity of the quantum transition rules, which will be crucial to the analysis of the resulting Hamiltonian in Section 7.7.

If we complete the partial U (associated with the Track 1–3 transitions $\begin{array}{|c|c|} \hline \circ & \text{⌞} \\ \hline \end{array} \rightarrow \begin{array}{|c|c|} \hline \text{⌞} & \circ \\ \hline \end{array}$) as defined so far by (96)–(98) to a unitary, this will add additional transitions out of the final state q_f , and also out of configurations with the head at the beginning or end of the chain. But this means that, instead of the encoded QTM doing nothing for the remaining time steps after reaching such a configuration, it will continue to evolve in some arbitrary way, depending on the particular choice of completion of U . On the other hand, if we omit transitions out of these configurations, the resulting U will only be a partial isometry, not a full unitary.

We need to use up any remaining time steps in a controlled way, such that the Track 5 QTM tape is left unaltered after the QTM enters q_f or runs out of tape. To this end, we deliberately add additional transitions out of the final state q_f , and also out of configurations with the head at the beginning or end of the chain, which cause the QTM to switch over to running some other, inconsequential computation to use up the remaining time. This time-wasting computation will use Track 6 as its tape, leaving the configuration of the Track 5 QTM tape untouched.

The only requirement on this time-wasting computation is that it must be guaranteed not to halt or run out of tape before it has used up all the remaining time steps of the clock (otherwise we face the same unitarity issue once again). The natural choice is simply to run the same base- ζ counter computation as the clock (but running on Tracks 4 and 6). The counter TM never halts; it increments the number written on its tape forever, or – when encoded on a finite chain – until it runs out of tape space. Furthermore, since the clock will already have “ticked” for some positive number of time steps before this time-wasting counter TM starts running, the clock counter TM is guaranteed to run out of tape (and hence the clock stop ticking) first.

Transitioning from q_f to the time-wasting counter TM can be accomplished simply by dovetailing the counter TM after the QTM, and encoding this dovetailed machine instead of the original QTM. Concretely, this means we must add to the Hamiltonian more transition rule terms of the form (96) that encode all the transition rules of the counter TM from (76).² Note that Track 6 contains additional

¹Although the QTMs we will encode will in fact never do this.

²Since these are the transition rules of a classical reversible TM, they have a particularly simple form, the right hand side of each (96a) rule containing a single term.

subscripted variants \vdash_q of the \vdash tape symbol; the time-wasting counter TM transition rules simply ignore the subscript, and treat all of these as a \vdash . (I.e. the counter TM rules that read a \vdash symbol in (76) are duplicated for each \vdash_q variant.) These time-wasting counter TM transition rules involve a disjoint set of Track 4 internal states $q \in P'$ to those $q \in Q'$ used by the QTM. Thus the left- and right-hand sides of all time-wasting counter TM (96) transition rules are orthogonal to those encoding the QTM transition rules.

We also add the following transition out of the final state q_f of the original QTM into the initial state p_α of the counter TM, which dovetails the counter TM after the QTM. This transition rule acts on Tracks 1, 4 and 6:

$$\begin{array}{|c|c|} \hline \bigcirc & \textcircled{\vdash} \\ \hline q_f & \bigcirc \\ \hline \vdash & \cdot \\ \hline \end{array} \longrightarrow \begin{array}{|c|c|} \hline \textcircled{\vdash} & \bigcirc \\ \hline p_\alpha & \bigcirc \\ \hline \vdash_{q_f} & \cdot \\ \hline \end{array} . \quad (98)$$

We must also add transitions that switch to running the time-wasting counter TM if the QTM runs out of tape, i.e. if it enters a configuration in which the head is at the beginning of the tape and would move left in the next time step, or one with the head at the end of the chain that would move right. (Though in the case of the proper machines considered in parts (iv) to (vi) of Theorem 33, the former will never occur.) Moving left off the beginning of the tape is particularly easy, both because the head is already at the beginning of the chain, and because left-moving QTM transitions are implemented by first transitioning to an auxiliary $q'_L \in Q'_L$ state without moving the head. The following transition rule acting on Tracks 1, 2, 4 and 6 accomplishes this:

$$\begin{array}{|c|c|} \hline \textcircled{\vdash} \\ \hline \leftarrow \neg p_\alpha \\ \hline q'_L \\ \hline \vdash \\ \hline \end{array} \longrightarrow \begin{array}{|c|c|} \hline \textcircled{\vdash} \\ \hline \leftarrow \neg p_\alpha \\ \hline p_\alpha \\ \hline \vdash_{q'_L} \\ \hline \end{array} . \quad (99)$$

Note that, to preserve reversibility (unitarity), we must keep a record of which internal state q'_L led to the QTM running out of tape. We record this in the \vdash_q symbol written to the Track 6 time-wasting TM tape.

Configurations in which the head step right off the end of the chain involve slightly more effort, as the head is at the wrong end of the chain to start running the time-wasting counter TM. From Section 7.6, a configuration in which the Track 4 QTM head is at the very end of the chain indicates that the head has stepped right off the usable portion of the Track 5 tape. We first transition from any such configuration into an auxiliary Track 4 head reset state r , which moves the head all the way back to the beginning of the chain, before transitioning into the initial state of the counter TM. This requires three additional sets of transition rules.

We want the first new rule to act on Tracks 1 and 4 at the very end of the chain, and transition into the auxiliary reset state r :

$$\begin{array}{|c|c|} \hline \textcircled{\tau} & \textcircled{>} \\ \hline q & \\ \hline \end{array} \rightarrow \begin{array}{|c|c|} \hline \textcircled{\tau} & \textcircled{>} \\ \hline r_q & \\ \hline \end{array}. \quad (100)$$

However, we have to make it compatible with the transitions already defined for Tracks 1–3. Hence, instead of (100), we include the following transitions acting on Tracks 1,2,4 and 1,2,3,4.

$$\begin{array}{|c|c|} \hline \textcircled{\tau} \\ \hline \textcircled{\circ} \\ \hline q \\ \hline \end{array} \rightarrow \begin{array}{|c|c|} \hline \textcircled{\tau} \\ \hline \textcircled{\circ} \\ \hline r_q \\ \hline \end{array}, \quad \begin{array}{|c|c|} \hline \textcircled{\tau} \\ \hline p \\ \hline \sigma \\ \hline q \\ \hline \end{array} \rightarrow \begin{array}{|c|c|} \hline \textcircled{\tau} \\ \hline p_N \\ \hline \tau \\ \hline r_q \\ \hline \end{array}, \quad \begin{array}{|c|c|} \hline \textcircled{\tau} \\ \hline p'_R \\ \hline \cdot \\ \hline q \\ \hline \end{array} \rightarrow \begin{array}{|c|c|} \hline \textcircled{\tau} \\ \hline p_R \\ \hline \cdot \\ \hline r_q \\ \hline \end{array}. \quad (101)$$

The second new rule also acts on Tracks 1 and 4, and steps the auxiliary r state to the left along with the $\textcircled{\tau}$:

$$\begin{array}{|c|c|} \hline \textcircled{\circ} & \textcircled{\tau} \\ \hline \textcircled{\circ} & r_q \\ \hline \end{array} \rightarrow \begin{array}{|c|c|} \hline \textcircled{\tau} & \textcircled{\circ} \\ \hline r_q & \textcircled{\circ} \\ \hline \end{array}. \quad (102)$$

The third acts on on Tracks 1, 2 4 and 6 at the very beginning of the chain, and transitions into the initial configuration of the time-wasting counter TM:

$$\begin{array}{|c|c|} \hline \textcircled{\tau} \\ \hline \neg p_\alpha \\ \hline r_q \\ \hline \vdash \\ \hline \end{array} \rightarrow \begin{array}{|c|c|} \hline \textcircled{\tau} \\ \hline \neg p_\alpha \\ \hline p_\alpha \\ \hline \vdash_q \\ \hline \end{array}. \quad (103)$$

Again, to preserve reversibility, we record in the \vdash_q symbol which internal state caused the QTM to run out of tape.

Table 3: All transition rules for Tracks 4, 5 and 6.

Track 1 and 4 rules

Track 1, 4 and 5 rules

\longrightarrow

$\sum_{\tau, q_R} \delta(p, \sigma, \tau, q_R, R)$

$+$


$\sum_{\tau, q_N} \delta(p, \sigma, \tau, q_N, N)$

$+$


$\sum_{\tau, q_L} \delta(p, \sigma, \tau, q_L, L)$

\longrightarrow


Track 1, 4 and 6 rules




Track 1, 2 and 4 rules



Track 1, 2, 4 and 6 rules



Track 1, 2, 3 and 4 rules



The transition rules in Table 3 implement all the transitions that will take place for a properly initialised QTM. All of them have the desired form $|ab\rangle|ij\rangle \rightarrow |cd\rangle U_{abcd}|ij\rangle$ where $ab \rightarrow cd$ is a Track 1–3 transition and U_{abcd} are partial isometries defined by their action on a subset $|ij\rangle$ of the computational basis of Tracks 4–6. We can now complete the unitaries U_{abcd} at will. All we have to check is that they do indeed define a partial isometry.

The left hand sides of each transition rule in (96)–(99) and (101)–(103) are manifestly orthogonal. The right hand sides of all (96b), (97)–(99) and (101)–(103) rules are clearly mutually orthogonal and normalised, and are also orthogonal to the right hand sides of all (96a) rules.

It remains to show that (96a) defines an isometry. Recall that there are two sets of (96a) terms, one encoding the QTM transition rules, and the other the time-wasting counter TM. Furthermore, the left- and right-hand sides of all rules in the QTM set are orthogonal to those of the counter TM set, as they involve disjoint sets \mathbb{Q}' and P' (resp.) of internal states. Thus it suffices to prove that each set of transition rules considered separately defines an isometry. Since classical reversible TMs are special cases of QTMs, we prove the result for QTMs.

Let $|\psi(p, \sigma)\rangle$ denote the state on the RHS of (96a). Then

$$\begin{aligned} \|\psi(p, \sigma)\| &= \sum_{\tau, q_R} |\delta(p, \sigma, \tau, q_R, R)|^2 + \sum_{\tau, q_N} |\delta(p, \sigma, \tau, q_N, N)|^2 \\ &\quad + \sum_{\tau, q_L} |\delta(p, \sigma, \tau, q_L, L)|^2 \\ &= \|\delta(p, \sigma)\| = 1 \end{aligned} \tag{104}$$

by the normalisation condition of Theorem 19, so the right hand sides of (96a) are normalised.

Similarly, for $(p_1, \tau_1) \neq (p_2, \tau_2)$,

$$\begin{aligned} \langle \psi(p_1, \sigma_1) | \psi(p_2, \sigma_2) \rangle &= \sum_{\tau, q_R} \delta(p_1, \sigma_1, \tau, q_R, R)^* \delta(p_2, \sigma_2, \tau, q_R, R) \\ &\quad + \sum_{\tau, q_N} \delta(p_1, \sigma_1, \tau, q_N, N)^* \delta(p_2, \sigma_2, \tau, q_N, N) \\ &\quad + \sum_{\tau, q_L} \delta(p_1, \sigma_1, \tau, q_L, L)^* \delta(p_2, \sigma_2, \tau, q_L, L) \\ &= \langle \delta(p_1, \tau_1), \delta(p_2, \tau_2) \rangle = 0 \end{aligned} \tag{105}$$

by the orthogonality condition of Theorem 19, so the right hand sides of (96a) are mutually orthogonal. Thus the QTM and time-wasting counter TM transition rules from (96a) preserve orthonormality, as required.

7.6.3 QTM initialisation sweep

We use penalty terms that only apply during the initialisation sweep to enforce the correct initial configurations of Tracks 4, 5 and 6. For Track 4, we use the following penalty terms, which force a single q_0 at the very beginning of the track:

$$\begin{array}{|c|c|} \hline \neg \leftarrow & \begin{array}{c} \rightarrow \\ 0 \end{array} \\ \hline & \neg \odot \end{array}, \quad \begin{array}{|c|c|} \hline \leftarrow & \begin{array}{c} \rightarrow \\ 0 \end{array} \\ \hline & \neg q_0 \end{array}. \quad (106)$$

The $\begin{array}{c} \leftarrow \\ \odot \end{array}$ Track 1 states only ever occur during the initialisation sweep in the final K sites at the right end of the chain. We use these to force the initial configuration of Track 5 to contain K blank symbols at the very end, as required by Theorem 33.¹

$$\begin{array}{|c|} \hline \begin{array}{c} \leftarrow \\ 0 \end{array} \\ \hline \neg 1 \end{array}, \quad \begin{array}{|c|} \hline \begin{array}{c} \leftarrow \\ \odot \end{array} \\ \hline \neg \# \end{array}. \quad (107)$$

Track 6 is forced to be in the all-blank configuration, except for the very first cell which contains a \vdash :

$$\begin{array}{|c|c|} \hline \neg \leftarrow & \begin{array}{c} \rightarrow \\ 0 \end{array} \\ \hline & \neg \# \end{array}, \quad \begin{array}{|c|c|} \hline \leftarrow & \begin{array}{c} \rightarrow \\ 0 \end{array} \\ \hline & \neg \vdash \end{array}. \quad (108)$$

Table 4: Initialisation illegal pairs for Tracks 4–6.

Illegal pairs		
Tracks 1 and 4	Tracks 1 and 5	Tracks 1 and 6
$\begin{array}{ c c } \hline \neg \leftarrow & \begin{array}{c} \rightarrow \\ 0 \end{array} \\ \hline & \neg \odot \end{array}$	$\begin{array}{ c } \hline \begin{array}{c} \leftarrow \\ 0 \end{array} \\ \hline \neg 1 \end{array}$	$\begin{array}{ c c } \hline \neg \leftarrow & \begin{array}{c} \rightarrow \\ 0 \end{array} \\ \hline & \neg \# \end{array}$
$\begin{array}{ c c } \hline \leftarrow & \begin{array}{c} \rightarrow \\ 0 \end{array} \\ \hline & \neg q_0 \end{array}$	$\begin{array}{ c } \hline \begin{array}{c} \leftarrow \\ \odot \end{array} \\ \hline \neg \# \end{array}$	$\begin{array}{ c c } \hline \leftarrow & \begin{array}{c} \rightarrow \\ 0 \end{array} \\ \hline & \neg \vdash \end{array}$

7.7 Analysis

Lemmas 39 and 40 are the key results needed to prove the desired ground state properties of the Hamiltonian we have constructed, thanks to the ‘‘Clairvoyance Lemma’’

¹These additional blank symbols will be needed later to provide for the small space overhead of the QTM constructed in Theorem 11.

from Aharonov et al. [Aha+09, Lemma 4.2] (see also [GI09, Lemma 5.6]).¹ Since the Clairvoyance Lemma has previously only been stated and proven for specific constructions, we state and prove a very general version of the Lemma here, which applies to any standard-form Hamiltonian, as defined above:

Lemma 42 (Invariant subspaces) *Let H_{trans} and H_{pen} define a standard-form Hamiltonian as defined in Definition 41. Let $\mathcal{S} = \{S_i\}$ be a partition of the standard basis states of \mathcal{H}_C into minimal subsets S_i that are closed under the transition rules (where a transition rule $|ab\rangle_{CD} |\psi\rangle \rightarrow |cd\rangle_{CD} U_{abcd} |\psi\rangle$ acts on \mathcal{H}_C by restriction to $(\mathbb{C}^C)^{\otimes 2}$, i.e. it acts as $ab \rightarrow cd$).*

Then $\mathcal{H} = (\bigoplus_S \mathcal{K}_{S_i}) \otimes \mathcal{H}_Q$ decomposes into invariant subspaces $\mathcal{K}_{S_i} \otimes \mathcal{H}_Q$ of $H = H_{\text{pen}} + H_{\text{trans}}$ where \mathcal{K}_{S_i} is spanned by S_i .

Proof H_{pen} is diagonal in the standard basis by definition (part (ii) of Definition 41), so $\mathcal{K}_{S_i} \otimes \mathcal{H}_Q$ are trivially invariant under H_{pen} . But, by the form of the transition rule terms (part (i) of Definition 41), the image $H_{\text{trans}} |x\rangle_C |\varphi\rangle_Q$ of a standard basis state $|x\rangle_C$ under H_{trans} has support only on standard basis states of \mathcal{H}_C that are reachable by transition rules from $|x\rangle_C$. Closure of $\mathcal{K}_{S_i} \otimes \mathcal{H}_Q$ under H_{trans} is then immediate from the definition of S_i . \square

Lemma 43 (Clairvoyance Lemma) *Let $H = H_{\text{trans}} + H_{\text{pen}}$ be a standard-form Hamiltonian as defined in Definition 41, and let \mathcal{K}_S be defined as in Lemma 42. Let $\lambda_0(\mathcal{K}_S)$ denote the minimum eigenvalue of the restriction $H|_{\mathcal{K}_S \otimes \mathcal{H}_Q}$ of $H = H_{\text{trans}} + H_{\text{pen}}$ to the invariant subspace $\mathcal{K}_S \otimes \mathcal{H}_Q$.*

Assume that there exists a subset \mathcal{W} of standard basis states for \mathcal{H}_C with the following properties:

- (i). *All legal standard basis states for \mathcal{H}_C are contained in \mathcal{W} .*
- (ii). *\mathcal{W} is closed with respect to the transition rules.*
- (iii). *At most one transition rule applies in each direction to any state in \mathcal{W} .*
- (iv). *For any subset $S \subseteq \mathcal{W}$ that contains only legal states, there exists at least one state to which no backwards transition applies.*

Then each \mathcal{K}_S falls into one of the following categories:

- (1). *S contains only illegal states, and $\lambda_0(\mathcal{K}_S) \geq 1$.*
- (2). *S contains both legal and illegal states, and $\lambda_0(\mathcal{K}_S) = \Omega(1/|S|^3)$.*

¹Note that our definition of “legal” states follows that of [GI09], which differs from the definition in [Aha+09].

(3). S contains only legal states, and $\lambda_0(\mathcal{K}_S) = 0$. The corresponding eigenspace is

$$\ker(H_{\text{trans}} + H_{\text{pen}}) = \text{span} \left\{ \frac{1}{\sqrt{|S|}} \sum_{t=0}^{|S|-1} |t\rangle_C |\psi_t\rangle_Q \right\} \quad (109)$$

where $|t\rangle_C$ are the states in S , $|\psi_0\rangle$ is any state in \mathcal{H}_Q , and $|\psi_t\rangle := U_t \dots U_1 |\psi_0\rangle_Q$ where U_t is the unitary on \mathcal{H}_Q appearing in the transition rule that takes $|t-1\rangle_C$ to $|t\rangle_C$. All other states in $\mathcal{K}_S \otimes \mathcal{H}_Q$ have energy at least $\Omega(1/|S|^2)$.

To prove this, we will need Kitaev's geometrical lemma:

Lemma 44 (Geometrical Lemma – Lemma 14.4 in [KSV02]) *Let $A, B \geq 0$ be non-negative operators such that $\ker A \cap \ker B = \{0\}$, $\lambda_{\min}(A|_{\text{supp } A}) \geq \mu$ and $\lambda_{\min}(B|_{\text{supp } B}) \geq \mu$. Then*

$$A + B \geq 2\mu \sin^2 \left(\frac{\theta}{2} \right) \quad (110)$$

where

$$\cos \theta := \max_{\substack{|\psi\rangle \in \ker A \\ |\varphi\rangle \in \ker B}} |\langle \psi | \varphi \rangle|. \quad (111)$$

Proof (of Clairvoyance Lemma 43) Case (1) is trivial since $\langle x|_C \langle \psi|_Q H_{\text{pen}} |x\rangle_C |\psi\rangle_Q \geq 1$ for any illegal standard basis state $|x\rangle_C$.

Now consider cases (2) and (3). By assumption, all legal standard basis states of \mathcal{H}_C are contained in \mathcal{W} , which is closed under transition rules. Thus closure of S (Lemma 42) implies $S \subseteq \mathcal{W}$. Consider the directed graph of states in \mathcal{W} formed by adding a directed edge between pairs of states connected by transition rules. By assumption, only one transition rule applies in each direction to any state in \mathcal{W} , so the graph consists of a union of disjoint paths (which could be loops in case (2)). Minimality of S (Lemma 42) implies that S consists of a single such connected path.

Let $t = 0, \dots, |S| - 1$ denote the states in S enumerated in the order induced by the directed graph. H_{trans} then acts on the subspace $\mathcal{K}_S \otimes \mathcal{H}_Q$ as

$$H_{\text{trans}}|_{\mathcal{K}_S \otimes \mathcal{H}_Q} = \sum_{t=0}^T \frac{1}{2} \left(|t\rangle\langle t| \otimes \mathbb{1} + |t+1\rangle\langle t+1| \otimes \mathbb{1} - |t+1\rangle\langle t| \otimes U_t - |t\rangle\langle t+1| \otimes U_t^\dagger \right), \quad (112)$$

where U_t is the unitary on \mathcal{H}_Q appearing in the transition rule that takes $|t\rangle_C$ to $|t+1\rangle_C$. $T = |S| - 1$ if the path in S is a loop, otherwise $T = |S| - 2$.

Consider a single term $H_t := |t\rangle\langle t| \otimes \mathbb{1} + |t+1\rangle\langle t+1| \otimes \mathbb{1} + |t+1\rangle\langle t| \otimes U_t + |t\rangle\langle t+1| \otimes U_t^\dagger$ from (112). Defining the unitary $W_t = |t+1\rangle\langle t+1| \otimes U_t^\dagger + (\mathbb{1} - |t+1\rangle\langle t+1|) \otimes \mathbb{1}$, we have

$$W_t H_t W_t^\dagger = (|t\rangle - |t+1\rangle) (\langle t| - \langle t+1|) \otimes \mathbb{1} \geq 0. \quad (113)$$

Thus, whether the path in S forms a loop or not, we have¹

$$\begin{aligned} H_{\text{trans}}|_{\mathcal{K}_S \otimes \mathcal{H}_Q} &\geq \sum_{t=0}^{|S|-2} \frac{1}{2} \left(|t\rangle\langle t| \otimes \mathbb{1} + |t+1\rangle\langle t+1| \otimes \mathbb{1} - |t+1\rangle\langle t| \otimes U_t - |t\rangle\langle t+1| \otimes U_t^\dagger \right) \\ &=: H_{\text{path}}, \end{aligned} \quad (114)$$

with equality if the path is not a loop.

Defining the unitary

$$W = \sum_{t=0}^{|S|-2} |t\rangle\langle t| \otimes \prod_{i=0}^t U_i^\dagger + (|S|-1)\langle S-1| \otimes \mathbb{1}, \quad (115)$$

we have $H_{\text{path}} \simeq WH_{\text{path}}W^\dagger = E \otimes \mathbb{1}$ where

$$E = \begin{pmatrix} \frac{1}{2} & -\frac{1}{2} & 0 & \dots & \dots & 0 & 0 \\ -\frac{1}{2} & 1 & -\frac{1}{2} & \ddots & & & 0 \\ 0 & -\frac{1}{2} & 1 & -\frac{1}{2} & & & \vdots \\ \vdots & \ddots & -\frac{1}{2} & 1 & \ddots & & \\ & & & \ddots & \ddots & \ddots & \vdots \\ \vdots & & & & & \ddots & \ddots & 0 \\ 0 & & & & \ddots & \ddots & 1 & -\frac{1}{2} \\ 0 & 0 & \dots & \dots & 0 & -\frac{1}{2} & \frac{1}{2} \end{pmatrix}. \quad (116)$$

The matrix E is the Laplacian of the random walk on a line, and it is well known that its eigenvalues are given by $\lambda_k = 1 - \cos q_k$ where $q_k = k\pi/|S|$, $k = 0, \dots, |S|-1$, with corresponding eigenvectors $|\phi_k\rangle \propto \sum_{j=0}^{|S|-1} \cos(q_k(j + \frac{1}{2})) |j\rangle$ [KSV02]. Thus

$$\ker H_{\text{path}} = \text{span} \left\{ W^\dagger \frac{1}{\sqrt{|S|}} \sum_{t=0}^{|S|-1} |t\rangle_C |\psi_0\rangle_Q \right\} = \text{span} \left\{ \frac{1}{\sqrt{|S|}} \sum_{t=0}^{|S|-1} |t\rangle_C |\psi_t\rangle_Q \right\} \quad (117)$$

where $|\psi_t\rangle := U_t \dots U_0 |\psi_0\rangle$ for any $|\psi_0\rangle$.

First consider case (2). Take $A = H_{\text{path}}$ and $B = H_{\text{pen}}|_{\mathcal{K}_S \otimes \mathcal{H}_Q}$ in Lemma 44. We have $\lambda_{\min}(A|_{\text{supp } A}) = 1 - \cos q_1 = \Omega(1/|S|^2)$ and $\lambda_{\min}(B|_{\text{supp } B}) = 1$. Furthermore, since S contains at least one illegal state in case (2), we have $\ker A \cap \ker B = \{0\}$ and

$$\cos^2 \theta = \max_{|\psi\rangle} \langle \phi_0 |_C \langle \psi |_Q W^\dagger \Pi_{\ker B} W |\phi_0\rangle_C |\psi\rangle_Q \leq 1 - 1/|S|, \quad (118)$$

¹This operator has the same form as Kitaev's Hamiltonian, and the analysis from here on is similar to that in Kitaev, Shen, and Vyalyi [KSV02].

where $\Pi_{\ker B}$ is the projector onto $\ker B$ and we have used the fact that $B = H_{\text{pen}}|_{\mathcal{K}_S \otimes \mathcal{H}_Q}$ is diagonal in the standard basis. Invoking Lemma 44, we obtain

$$(H_{\text{trans}} + H_{\text{pen}})|_{\mathcal{K} \otimes \mathcal{H}_Q} \geq H_{\text{path}} + H_{\text{pen}}|_{\mathcal{K} \otimes \mathcal{H}_Q} = \Omega(1/|S|^3), \quad (119)$$

thus $\lambda_0(\mathcal{K}_S) = \Omega(1/|S|^3)$ as claimed.

Finally, consider case (3). Since S contains only legal states in this case, $H_{\text{pen}}|_{\mathcal{K}_S \otimes \mathcal{H}_Q} = 0$. Furthermore, by condition (iv) the states in S cannot form a loop, so $H|_{\mathcal{K}_S \otimes \mathcal{H}_Q} = H_{\text{trans}}|_{\mathcal{K}_S \otimes \mathcal{H}_Q} = H_{\text{path}} \simeq E \otimes \mathbb{1}$. The claim in the Lemma now follows from the form of the eigenvalues and kernel of E , given above. \square

Note that the proof of the Clairvoyance Lemma relied crucially on unitarity of the operators U_{abcd} appearing in the transition rules (see Lemma 42). In particular, it is *not* sufficient for the U_{abcd} to be partial isometries. That would not allow us to unitarily transform $H_{\text{trans}}|_{\mathcal{K}_S \otimes \mathcal{H}_Q}$ into $E \otimes \mathbb{1}$ using the unitary W in (115) and (116). Thus, in order to apply the Clairvoyance Lemma 43, we *must* complete the quantum parts of the transition rules to a full unitary. This necessitates the time-wasting construction in Section 7.6.2 (or similar).

We are finally in a position to prove Theorem 33, the main result of this section. As commented above, let $\mathcal{H}_C = (\mathbb{C}^C)^{\otimes L}$ and $\mathcal{H}_Q = (\mathbb{C}^Q)^{\otimes L}$ be the Track 1–3 and Track 4–5 Hilbert spaces, respectively, for a chain of length L , where $\mathbb{C}^C \simeq \text{span}\{|s, p, \sigma\rangle\} \oplus \text{span}\{|\otimes, \otimes\rangle\}$ and $\mathbb{C}^Q \simeq \text{span}\{|q, \tau\rangle\} \oplus \text{span}\{|\otimes, \otimes\rangle\}$ (cf. (67)). Let $h_{\text{trans}} \in \mathcal{B}((\mathbb{C}^C \otimes \mathbb{C}^Q)^{\otimes 2})$ be the sum of all the transition rule terms defined in Tables 1 and 3 (omitting those marked in Table 1) after completing to a full unitary so that the resulting Hamiltonian is standard-form. Let $h_{\text{pen}} \in \mathcal{B}((\mathbb{C}^C \otimes \mathbb{C}^Q)^{\otimes 2})$ be the sum of all penalty terms defined in Table 2, and $h_{\text{init}} \in \mathcal{B}((\mathbb{C}^C \otimes \mathbb{C}^Q)^{\otimes 2})$ be the sum of all Track 4–6 initialisation penalty terms defined in Table 4.

Define the standard-form (Definition 41) Hamiltonian $H(L) = H_{\text{trans}}(L) + H_{\text{pen}}(L) \in \mathcal{H}_C \otimes \mathcal{H}_Q$ on a chain of length L , where $H_{\text{trans,pen}} = \sum_{i=1}^{L-1} h_{\text{trans,pen}}^{(i,i+1)}$, and similarly define $H_{\text{init}}(L) := \sum_{i=1}^{L-1} h_{\text{init}}^{(i,i+1)}$.

Proposition 45 (Unique ground state) *The unique 0-energy eigenstate of $(H_{\text{trans}}(L) + H_{\text{pen}}(L) + H_{\text{init}}(L))|_{S_{br}}$ is the computational history state*

$$\frac{1}{\sqrt{T}} \sum_{t=0}^T |\phi_t\rangle_C |\psi_t\rangle_Q, \quad (120)$$

where $T = \Omega(L|\Sigma \times Q|^L)$, $|\psi_0\rangle_Q \in \mathcal{H}_Q$ is the initial configuration of the QTM with the unary representation of the number $L - K - 3$ as input, and $|\psi_t\rangle_Q \in \mathcal{H}_Q$ is the sequence of states produced by evolving $|\psi_0\rangle$ under the QTM and time-wasting counter TM transition rules (where each such state is duplicated $O(L)$ times in succession in the sequence).

Proof Let us first apply the Clairvoyance Lemma to $(H_{\text{trans}}(L) + H_{\text{pen}}(L))|_{S_{\text{br}}}$. We have to check that h_{trans} and h_{pen} fulfil the requirements of the Lemma.

Take \mathcal{H}_C to be the Hilbert space of Tracks 1–3, and \mathcal{H}_Q the Hilbert space of Tracks 4–5. Define \mathcal{W} to be the set of all well-formed Track 1–3 standard basis states. Any state that is not well-formed violates an illegal pair enforcing a regular expression, so all legal Track 1–3 states are well-formed. \mathcal{W} therefore fulfils condition (i) of Lemma 43. By Lemma 39, the set of well-formed states is closed under the transition rules and at most one transition rule applies in each direction to any well-formed state, so \mathcal{W} fulfils conditions (ii) and (iii) of Lemma 43. Finally, by Lemma 40, the only subset of legal Track 1–3 standard basis states that is closed under the transition rules is the set of clock states $\{|\phi_t\rangle\}$, which has one state $|\phi_0\rangle$ to which no backward transition applies, and one state $|\phi_T\rangle$ to which no forward transition applies. So \mathcal{W} also fulfils condition (iv) of Lemma 43.

All penalty terms are diagonal in the standard basis, as required by Lemmas 42 and 43, and all transition rules are of the form required in Lemma 43, by construction. h_{trans} and h_{pen} therefore fulfil all the requirements of Lemmas 42 and 43.

Invoking the Clairvoyance Lemma 43, the 0-energy eigenspace of $(H_{\text{trans}} + H_{\text{pen}})|_{S_{\text{br}}}$ has the form

$$\ker(H_{\text{trans}} + H_{\text{pen}})|_{S_{\text{br}}} = \text{span} \left\{ \frac{1}{\sqrt{|S|}} \sum_{t=0}^{|S|-1} |\phi_t\rangle_C |\psi_t\rangle_Q \right\}, \quad (121)$$

where $|\psi_t\rangle_Q = U_t \dots U_1 |\psi_0\rangle_Q$ and $|\psi_0\rangle_Q$ is any state of \mathcal{H}_Q .

Now we include the h_{init} terms, and show that $|\psi_0\rangle_Q$ must be as claimed in Proposition 45. For $0 \leq t \leq 2L - 1$, the clock states $|\phi_t\rangle$ consist of an \odot that sweeps from left to right and back along the chain (where for the time steps $t = L+1, \dots, L+K$ of this sweep, the arrow is in the states $\odot_{(0)}$, \dots , $\odot_{(K)}$, respectively). The transition rules that apply during the initialisation sweep act trivially on Tracks 4–5, so $|\psi_t\rangle_Q = |\psi_0\rangle_Q$ for $t \leq 2L - 1$. By construction, the h_{init} penalty terms from Table 4 give an additional energy penalty to any Track 4–5 states that are not in the desired initial QTM configuration when the \odot sweeps past.

Consider any Track 4–5 state $|\psi_0\rangle_Q$ that is not the desired initial QTM configuration. Then there is some $0 \leq t \leq 2L - 1$ for which $\langle \phi_t |_C \langle \psi_t |_Q H_{\text{init}} |\phi_t\rangle_C |\psi_t\rangle_Q > 0$. Noting that the overall Hamiltonian $H \geq 0$ and the $|\phi_t\rangle$ are mutually orthogonal, this immediately implies that the unique 0-energy state is that given in (120).

The initial state $|\psi_0\rangle_Q$ on Track 4 matches the regular expression $\odot \circ^* q \odot^* \odot$. The transition rules U_t preserve this form – which allows us to identify Tracks 4 and 5 with the head and tape configuration of the QTM – and by construction implement one step of the QTM or time-wasting counter TM in each \odot sweep (all other unitaries applied to Tracks 4 and 5 during each sweep being identities).

Finally, the clock counter TM advances through $\Omega(|Q \times L|^L)$ configurations before it runs out of tape, and each step of this TM requires one complete sweep of the Track 1 arrow which takes $\Omega(L)$ steps. Thus $T = \Omega(L|Q \times L|^L)$ as claimed. We are done. \square

The local Hilbert space dimension in (67) manifestly depends only on the alphabet size and internal states of the counter TM encoded on Tracks 2 and 3, the QTM encoded on Tracks 4 and 5, and the time-wasting counter TM encoded on Tracks 4 and 6. However, the alphabet size and internal state space of both counter TMs depends only on the base ζ we choose, and this is completely determined by the alphabet size and internal state space of the QTM. Thus the local Hilbert space dimension d depends only on the alphabet size and number of internal states of the QTM, as required by part (i) of Theorem 33.

Since the local interaction h of our Hamiltonian is constructed by summing transition rule terms and penalty terms, and $h_{\text{trans}}, h_{\text{pen}}, h_{\text{init}} \geq 0$, we have $h \geq 0$. But we have shown that the overall Hamiltonian has a 0-energy eigenstate, hence it is frustration-free, as required in part (ii) of Theorem 33.

Proposition 45 implies that the unique ground state of $H|_{S_{\text{br}}}$ has the computational history state form claimed in part (iii) of Theorem 33. Moreover, since the history state encodes $\Omega(|Q \times L|^L)$ complete sweeps of the Track 1 arrow, and the QTM is advanced by exactly one step in each right-to-left sweep, the computational history state encodes $\Omega(|Q \times L|^L)$ steps of the computation. If the QTM halts before this, it transitions by the q_f rule of (98) to running the time-wasting counter TM, which never alters the QTM tape encoded on Track 5. Track 5 is therefore only in the state q_f for one time step. Thus part (iv) of Theorem 33 is also satisfied.

Now imagine that the QTM reaches a configuration in which the next step would move the head before the starting cell. (Though in the case of the proper machines considered in parts (iv) to (vi) of Theorem 33, the former will never occur.) Consider the corresponding $|\psi_t\rangle$ in the computational history state. The QTM transition rules do not alter Track 6, so Track 6 is still in its initial configuration. Since the QTM has deterministic head movement, then at no point during the evolution is its head ever in a superposition of locations. $|\psi_t\rangle$ is therefore of the form:

$$|\psi_t\rangle = |\otimes\rangle \left(\sum_q |q \otimes \dots \otimes\rangle_{\text{Track 4}} \otimes |\varphi_q\rangle_{\text{Track 5}} \otimes | \vdash_q \# \dots \# \rangle_{\text{Track 6}} \right) |\otimes\rangle, \quad (122)$$

where $|\varphi_q\rangle$ are unnormalised Track 5 QTM tape states, and all configurations in the superposition would step the QTM head left in the next time-step. The q_f transition rule from (98) therefore applies to all states in the superposition, and the next state

must have the form:

$$|\psi_{t+1}\rangle = |\langle\rangle\rangle \left(\sum_q |p_\alpha \odot \dots \odot\rangle_{\text{Track 4}} \otimes |\varphi_q\rangle_{\text{Track 5}} \otimes |{-}_q \# \dots \#\rangle_{\text{Track 6}} \right) |\rangle\rangle. \quad (123)$$

Thenceforth, only time-wasting counter TM transition rules apply.

A very similar argument applies if the QTM reaches a configuration in which the head moves beyond cell L . Since the time-wasting counter TM transitions never alter Track 5, part (v) of Theorem 33 follows.

This concludes the proof of Theorem 33, and this section.

8 Quasi-periodic tilings

In order to ultimately prove the required spectral properties of our final Hamiltonian, we require significantly stronger properties of quasi-periodic tilings than those used to prove tiling results. In particular, we will need strong “rigidity” results, which show that the quasi-periodic structure is in a sense robust to errors. As far as we are aware, these rigidity results are new. All the tiling results used in the proof of our main results are gathered in this section.

We exploit the very particular properties of a quasi-periodic tiling due to Robinson, which we briefly review in the following section. The Robinson tiling has a hierarchical geometric structure that lends itself to inductive proofs of the required rigidity results.

8.1 Robinson’s tiling

We now describe Robinson’s quasi-periodic tiling. It was discovered by Robinson in 1971 [Rob71] as a tool to simplify Berger’s proof of the undecidability of the tiling problem [Ber66].




Figure 9: The five basic tiles of Robinson’s tiling.

Robinson’s tiling is based on the five basic tiles showed in Figure 9 and all those that can be obtained from them by rotation or reflection. Tile (a) is called a *cross*, whereas tiles (b)–(e) are called *arms*. In a valid tiling, arrow heads must meet arrow tails. As drawn, cross (a) is said to face up/right. An arm is said to point in the direction of its only complete central arrow.




Figure 10: Parity markings to force the position of the crosses.

For each one of the five basic tiles of Figure 9, we introduce two possible colours, red or green, to each of the side arrows (note that there is one tile without

colour) with the following restrictions: at most one colour may be used horizontally and at most one colour may be used vertically; for the cross, the same colour must be used in both directions; for tile (b) one colour must be used horizontally and the other colour vertically. Moreover, we impose the restriction that green crosses must appear in alternate positions in alternate rows (and maybe in other positions too). This can be achieved by adding parity markings to the above tiles: the additional arrows showed in Figure 10, which should match properly in any valid tiling, with the following rules: parity marking (a) is associated with green crosses, parity marking (b) with horizontal green arrows, and (c) with vertical green arrows. Parity marking (d) is associated with all tiles. It is important to remark here that parity markings “live in a different layer” than the five basic tiles. That is, arrows from the parity markings must match only those from the parity markings and arrows from the basic tiles must match only those from the basic tiles.




Figure 11: A possible Robinson tiling of the plane.

If we now draw only the red collared lines, one possible tiling of the plane looks like Figure 11, which has the crucial (for our purposes) quasi-periodic structure consisting of squares of increasing size.

Let us try to explain briefly (full details can be found in [Rob71]) how Figure 11 emerges, and the remaining freedom it allows in constructing valid tilings of the




Figure 12: *Left*: A 3×3 square of the Robinson tiling, with the choice of a right/up facing cross in the middle tile. *Right*: A 7×7 square of the Robinson tiling, with the choice of a right/down facing cross in the middle tile. (To avoid confusion, we have only drawn the coloured lines of the tiles in this second figure.)

plane. Since crosses must appear in alternate rows in alternate positions, together with the form of the five basic tiles, if we take any cross, this completely determines the structure of the 3×3 square constructed in the direction it faces. For example, assume we start with a green cross facing left/down. Then we are forced to complete the 3×3 square having it as the top/right corner as shown in Figure 12, where in the middle tile we must have a red cross, with the only freedom of choosing where it faces.

Once this is fixed, the 7×7 square obtained by growing the 3×3 square in the direction this red crosses faces is fixed as shown in Figure 12, where there is a green cross in the central position, with the only freedom of choosing where it faces. Keeping this procedure gives a tiling like the one shown in Figure 11. It consists of a quasi-periodic structure of squares (called borders) of sizes 4^n (which means following Robinson's notation that the distance between two of the facing crosses delimiting the border is $4^n + 1$) repeated with period 2^{n+1} , for all $n \in \mathbb{N}$. Such a tiling can tile the whole plane, a half-plane or a quarter-plane, depending on how we sequentially choose the orientation of the central crosses. Lines between valid half or quarter planes must consist only of arms. If tiling divides the plane into two half-planes, the patterns in these half-planes may be shifted relative to one another by an arbitrary even number of cells. In this case, we call the line separating the two half-planes a *fault*. However, if a half-plane is further divided into two quarter-planes, no further shift is possible between these quarter-planes.

8.2 Rigidity of the Robinson tiling

We have just seen that with the original Robinson tiles, the quasi-periodic pattern of borders does not necessarily extend throughout the entire plane; there can be a fault line between two half-planes, with “slippage” between the patterns on either side of the fault (see Section 8.1 and [Rob71, § 8]). It will be convenient to slightly modify the Robinson tiles in order to prevent fault lines. To this end, wherever one of the five basic tiles does not have side-arrows, we add dashed side-arrows parallel to the solid central arrows (see Figure 13). We then extend the five basic tiles to the full tile set by rotation and reflection, adding parity markings, and adding red and green colourings to the solid side arrows, exactly as in Section 8.1. We refer to these as the “modified Robinson tiles”. (Similar modifications to the Robinson tile set appear elsewhere in the literature, e.g. in [Mie97].)




Figure 13: The five basic tiles of the modified Robinson tiling.

Since the all original tile markings are still present, any tiling using the modified tiles must also be a valid tiling using the original tiles if the dashed side-arrows are ignored. It is easy to verify that, for any tiling in which the 2^n -borders repeat periodically throughout the plane for all n (i.e. any tiling that does not contain a fault line), dashed side-arrows can be added to make it into a valid tiling using the modified tiles.

However, if the tiling contains a fault line, there is no way to consistently add dashed side-arrows. To see this, assume without loss of generality that the fault line is vertical. Then there must exist two back-to-back crosses on either side of the fault that are not aligned, so that one points up and the other down. Consider the cross to the left of the fault line. The dashed side-arrow on this cross that points to the right forces a dashed arm to extend in that direction, towards the other cross. Similarly, the left-pointing dashed side-arrow on the cross to the right of the fault forces a dashed arm to extend to the left. These dashed arms must meet somewhere between the two crosses. But, since one cross points up and the other down, one of these arms has the dashed side-arrow on top, the other has it on the bottom, so there is no way to consistently join up the arms.

The crucial property we will use is that the Robinson tiling has a very “rigid” structure. Even if we start from a tile configuration that contains a defect (a pair of non-matching adjacent tiles), outside of a ball around the defect, the defect has no

effect on the pattern of borders that appears in the rest of the tiling. The following Lemmas make this rigorous.¹

Lemma 46 (Cross rigidity) *In any tiling of a connected region with Robinson tiles, crosses must occur in alternate columns and in alternate rows.*

Proof Immediate from the observation that, for each “parity tile” in Figure 10, its neighbour in each direction is uniquely defined. □

Lemma 47 (Robinson rigidity) *Consider a connected region of a 2D square grid made up of square blocks of size $2^{n+1} \times 2^{n+1}$. Any tiling of such a region with modified Robinson tiles must contain a periodic pattern of 2^n -borders (some of which may be incomplete due to boundaries), repeating horizontally and vertically with period 2^{n+1} . That is, the tiling must contain the same periodic pattern of 2^n -borders as a section of the tiling of the infinite plane, up to translation and defects that do not affect the 2^n borders.*

Proof By Lemma 46, crosses must occur in alternate rows and columns throughout the region being tiled. Following Robinson [Rob71], we refer to these crosses as “1-squares”. (Note that there may be additional crosses in the tiling that do not constitute 1-squares.) Pick an orientation for one of these 1-square crosses. Its solid side-arrows force a solid arm to extend in the directions it faces. Unless there is insufficient space to the boundary of the region, there must be a cross two cells away in the direction of each arm. Just as in Robinson [Rob71, §3], these solid arms force the orientation of any two facing crosses to be mirror images.

Meanwhile, the dashed side-arrows force a dashed arm to extend in the directions that the cross faces *away* from. Again, these arms force the orientations of any back-to-back crosses to be mirror images. Thus, as soon as we pick the orientation of a single 1-square cross, the orientations of all adjacent crosses are forced, and hence the orientations of all 1-square crosses throughout the connected region.

Base case Consider first the case $n = 1$ with 4×4 blocks. Since orientations of adjacent pairs of 1-square crosses are mirror images of each other, they are forced to form groups of four facing towards the centre of a 3×3 block. As in Robinson [Rob71, §3], these blocks must be completed to form “3-squares” with 2-borders running around the edges and crosses in the centre. (The orientation of the central crosses is not fixed.) In our modified Robinson tiles, the dashed side-arrows also force the orientations of *back-to-back* crosses to be mirror images, which forces

¹Our “rigidity” results for the Robinson tiling could certainly be strengthened, but they suffice for our purposes.

adjacent 3-squares to be aligned. Thus the pattern of complete 2-borders must repeat horizontally and vertically with period 4 throughout the region.

However, in general some of the 1-squares will be too close to the boundaries of the region to form complete 3-squares. We need to show that these still force partial 2-borders. If a 1-square faces a boundary and is located at that boundary, then there is no further section of 2-border in that direction in any case, and we have nothing to show. If a 1-square faces a boundary and is more than two cells away from it, there must be another 1-square facing it before reaching the boundary, and the section of 2-border between them must be completed as usual. The only interesting case is if a 1-square faces a boundary one cell away from it. In this case, the solid side-arrows still force the cell between the cross and the boundary form an arm extending to the boundary. The only difference is that there is no facing cross in that direction to force the side-arrows of the arm to point inwards; they could instead point towards the boundary. Thus any of the arm tiles with solid side-arrows (in a suitable orientation) can be used to form this piece of the 2-border. For our purposes, we consider any of these to still constitute a section of 2-border. So partial 2-borders are still forced at the boundaries.

This completes the proof for the base case $n = 1$. It will be important for later to note that, if a partial 2-border is large enough to contain the cell that would normally be at the centre of the border, then the central cross is still forced in that cell. To see this, note that because our region is made up of adjacent square blocks, there must be arms on at least two sides of the central cell (not necessarily opposite sides, if the partial 2-square is located in a corner of the region). Regardless of which tile is used to form these arms, the central arrows perpendicular to the arms must face outwards as usual, because the solid side-arrows on the arms must be on the side away from the central cell. But the only tile that has outward-facing central arrows on adjacent sides is the cross, so it is forced.

$n = 2$ case Before proving the general case, it is instructive to consider the next case $n = 2$ (with 8×8 blocks). We can always divide each 8×8 block into four 4×4 blocks and apply the previous argument, so the 2-borders must repeat horizontally and vertically with period 4, throughout the region, with a cross at the centre of each 2-border (whether partial or complete). Each 8×8 block must therefore contain at least one complete 2-border, and all 2-borders must be aligned (hence their central crosses are also aligned).

Pick an orientation for one of these central crosses. The side-arrows of this cross force solid arms extending in the directions that the cross faces, and dashed arms extending in the directions it faces away from. Now, the only tile at which an arm can terminate is the cross. But there can be no crosses between adjacent 2-borders. Thus the arms must extend all the way from a central cross to its neighbours, again

forcing the orientation of both facing and back-to-back crosses to be mirror images of each other, so that they form groups of four facing towards the centre of a 5×5 block. As in Robinson [Rob71, §3], these blocks must be completed to form “5-squares” with 4-borders running around the edges and crosses in the centre. As before, the dashed arms force adjacent 5-squares to be aligned. So the pattern of complete 4-borders repeats horizontally and vertically throughout the region, with period 8.

It remains to consider crosses that are too close to the boundaries to form complete 4-borders. If the central cross of a 2-border faces a boundary (rather than another cross), its solid side-arrows still force an arm that extends towards the boundary. The only tile that can terminate an arm is a cross. However, we cannot place a cross anywhere within the 3-square surrounding the central cross, as it would break the tiling within the 3-square. The only remaining possibility is to place a cross in the one-cell-wide corridor that runs along each side of a 3-square. However, where an arm intersects a 2-border it must point outwards, which prevents placing a cross in the corridor. Thus the arm must be extended all the way to the boundary to form a partial 4-border (with some additional freedom in which arm tiles are used to form this section of 4-border).

This proves the Lemma for $n = 2$. As in [Rob71, § 3], a cross is forced at the centre of all complete 4-borders. If a partial 4-border is large enough to surround the cell that would be at the centre of the complete border, a cross is forced there too. To see this, observe that for partial borders of this size, the centre cell has partial 4-borders on at least two sides. The central arrows perpendicular to a border necessarily have arrow tails on the inside of the border. These force an arm extending back towards the central cell. That cell therefore has arrow tails on at least two sides, and the central cross is forced.

General case The argument for general n is very similar to the argument for the $n = 2$ case, and proceeds by induction on n . Assume for induction that the pattern of 2^{n-1} -borders repeats horizontally and vertically with period 2^n in any connected region made up of $2^n \times 2^n$ blocks, and that a cross is forced at the centre of each of those 2^{n-1} -borders (whether partial or complete).

Now consider a connected region made up of $2^{n+1} \times 2^{n+1}$ blocks. We can always divide each block into four $2^n \times 2^n$ blocks. So, by assumption, the pattern of 2^{n-1} -borders must repeat horizontally and vertically with period 2^n throughout the region, with a cross at the centre of each (partial or complete) 2^{n-1} -border. Each $2^{n+1} \times 2^{n+1}$ block therefore contains at least one complete 2^{n-1} -border, and adjacent borders (hence also their central crosses) are aligned. The solid and dashed side-arrows of these central crosses force arms extending between adjacent crosses, which in turn force the orientations of adjacent crosses to be mirror images.

Thus the arms extending between facing crosses form 2^n -borders, and adjacent 2^n -borders are aligned. The pattern of 2^n -borders therefore repeats horizontally and vertically throughout the region, with period 2^{n+1} .

For central crosses that face a boundary instead of an adjacent cross, the solid side-arrows on the cross still force an arm extending towards the boundary. To show that this arm necessarily extends all the way to the boundary, we must show that it is impossible to place a cross anywhere along the arm. Within the $(2^n - 1)$ -square surrounding the central cross, we cannot place a cross along the arm without breaking the tiling within $2^n - 1$ -square (cf. [Rob71, § 3]). Thus the arm must extend to the boundary of this $(2^n - 1)$ -square. Similarly, we cannot place a cross along the arm anywhere within the adjacent $(2^n - 1)$ -square. So, if the arm reaches the adjacent $(2^n - 1)$ -square, it must continue onwards to the boundary. The only remaining possibility is to place a cross in the one-cell-wide corridor between the adjacent $(2^n - 1)$ -squares. However, the arm necessarily points out of the $(2^n - 1)$ -square, which prevents placing a cross in the corridor.

Finally, for the induction step we must show that a cross is forced at the centre of each (complete or partial) 2^n -border. The argument applies to any border that extends sufficiently far to surround its centre cell (which includes complete borders). The central cell has (complete or partial) 2^n -borders on at least two sides. The central arrows perpendicular to a border necessarily have arrow tails pointing inwards, which force arms extending back towards the central cell. That cell therefore has arrow tails on at least two sides, forcing the central cross.

Together with the base case $n = 1$, this completes the proof. \square

We refer to the top edge of a 2^n -border as a 2^n -segment.¹

Lemma 48 (Segment bound) *The minimum number of 2^n -segments in a tiling of an $L \times H$ rectangle (width L , height H) using modified Robinson tiles is $\lfloor H/2^{n+1} \rfloor (\lfloor L/2^{n+1} \rfloor - 1)$, and this minimum can be attained simultaneously for all n .*

Proof Lemma 47 implies that the tiling must contain the usual periodic pattern of 2^n -borders, up to translation. Since the borders repeat vertically with period 2^{n+1} , the minimum number of rows of 2^n -segments is $\lfloor H/2^{n+1} \rfloor$ (obtained by translating the pattern vertically so that there is a row of incomplete borders along the top edge of the region). Similarly, since the borders repeat horizontally with period 2^{n+1} , the minimum number of 2^n -segments in each row of segments is $\lfloor L/2^{n+1} \rfloor - 1$ (obtained by translating the pattern horizontally so that there are incomplete segments at the left and/or right edges of the region).

¹It is on these 2^n -segments that the QTMs will ultimately “run” in the final construction.

Since any 2^n -border in the Robinson tiling contains 2^{n-1} borders inside it, the translations can be chosen such that there are incomplete segments of all sizes at the boundaries of the region. The result is now trivial. \square

The following Lemma is the key rigidity result that we will need later. It shows that a defect in the tiling (i.e. a non-matching pair of adjacent tiles) cannot affect the pattern of 2^n -borders in the tiling outside a ball of size 2^{n+1} centred on the defect.

Lemma 49 (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 2^n -segments is at least $\lfloor H/2^{n+1} \rfloor (\lfloor L/2^{n+1} \rfloor - 1) - 2d$.*

Proof Divide the $L \times H$ rectangle into contiguous square blocks of size $2^{n+1} \times 2^{n+1}$, with $\lfloor L/2^{n+1} \rfloor$ blocks in each row and $\lfloor H/2^{n+1} \rfloor$ in each column. (Any leftover cells play no further role in the argument.) Delete any block that contains at least one defect (arbitrarily assigning defects that occur at block boundaries to the bottom/left block). Consider any horizontal row of blocks. If the row contains d_i deleted blocks, those missing blocks divide it into at most $d_i + 1$ strips of height 2^{n+1} .

Since each strip is a connected, defect-free region made up of $2^{n+1} \times 2^{n+1}$ blocks, Lemma 47 implies that it must be tiled with the usual periodic pattern of 2^n -borders, up to translations. The strip may be connected to strips in adjacent rows, in which case we are not free to translate the tiling in each strip independently. Nonetheless, we can lower-bound the number of 2^n -segments by allowing the tiling within each strip to be translated independently, and minimising over all such translations.

Each strip has height 2^{n+1} , so, regardless of how the pattern is translated vertically, it must contain exactly one row of 2^n -segments. The minimum number of segments within a strip of block-length l is $l - 1$ (obtained by translating the pattern horizontally so that there are incomplete segments at the beginning and end of the strip). In other words, each block in the strip effectively contributes one segment to the total, except for the block at the right end of the strip, which does not contribute (we assign the segment between adjacent blocks to the left one).

If the row contains d_i deleted blocks, those missing blocks divide it into at most $d_i + 1$ strips.¹ Thus the row contains $\lfloor L/2^{n+1} \rfloor - d_i$ non-deleted blocks, with at most $d_i + 1$ of them located at ends of strips. The whole row therefore contains at least $\lfloor L/2^{n+1} \rfloor - d_i - (d_i + 1) = \lfloor L/2^{n+1} \rfloor - 2d_i - 1$ segments.

Summing over all $\lfloor H/2^{n+1} \rfloor$ rows, and noting that the total number of deleted blocks cannot be more than the total number of defects, we obtain the claimed lower bound on the total number of 2^n -segments. \square

¹For $d_i > n/2$, this substantially over-counts the number of strips, but it is sufficient for our purposes.

9 Putting it all together

9.1 Undecidability of the g.s. energy density

To prove our first main result, Theorem 6, we will prove a sequence of lemmas which allow us to combine together the Hamiltonian constructions from the previous sections, and progressively build up the final Hamiltonian.

We will repeatedly need to refer to 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 50 (Gottesman-Irani Hamiltonian) *Let \mathbb{C}^Q be a finite-dimensional Hilbert space with two distinguished orthogonal states labelled $|\otimes\rangle, |\oslash\rangle$. A Gottesman-Irani Hamiltonian is a 1D, translationally-invariant, nearest-neighbour Hamiltonian $H_q(r)$ on a chain of length r 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). $[h_q, |\otimes\rangle\langle\otimes| \otimes |\otimes\rangle\langle\otimes|] = [h_q, |\otimes\rangle\langle\otimes| \otimes |\oslash\rangle\langle\oslash|] = [h_q, |\oslash\rangle\langle\oslash| \otimes |\otimes\rangle\langle\otimes|] = [h_q, |\oslash\rangle\langle\oslash| \otimes |\oslash\rangle\langle\oslash|] = 0$.
- (iii). $\lambda_0(r) := \lambda_0(H_q(r)|_{S_{br}}) < 1$, where S_{br} is the subspace of states with fixed boundary conditions $|\otimes\rangle, |\oslash\rangle$ 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 $|\otimes\rangle$ at the left end of the chain or a $|\oslash\rangle$ at the right end of the chain.

Lemma 51 (Tiling + quantum layers) *Let $h_c^{\text{row}}, h_c^{\text{col}} \in \mathcal{B}(\mathbb{C}^C \otimes \mathbb{C}^C)$ be the local interactions of a 2D tiling Hamiltonian H_c , with two distinguished states (tiles) $|L\rangle, |R\rangle \in \mathbb{C}^C$. 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 50. Then there is a Hamiltonian on a 2D square lattice with nearest-neighbour interactions $h^{\text{row}}, h^{\text{col}} \in \mathcal{B}(\mathbb{C}^{C+Q+1} \otimes \mathbb{C}^{C+Q+1})$ with the following properties: For any region of the lattice, the restriction of the Hamiltonian to that region has an eigenbasis of the form $|T\rangle_c \otimes |\psi\rangle_q$, where $|T\rangle_c$ is a product state representing a classical configuration of tiles. Furthermore, for any given $|T\rangle_c$, the lowest energy choice for $|\psi\rangle_q$ consists of ground states of $H_q(r)$ on segments between sites in which $|T\rangle_c$ contains an $|L\rangle$ and an $|R\rangle$, a 0-energy eigenstate on segments between an $|L\rangle$ or $|R\rangle$ and the boundary of the region, and $|0\rangle$'s everywhere else.*

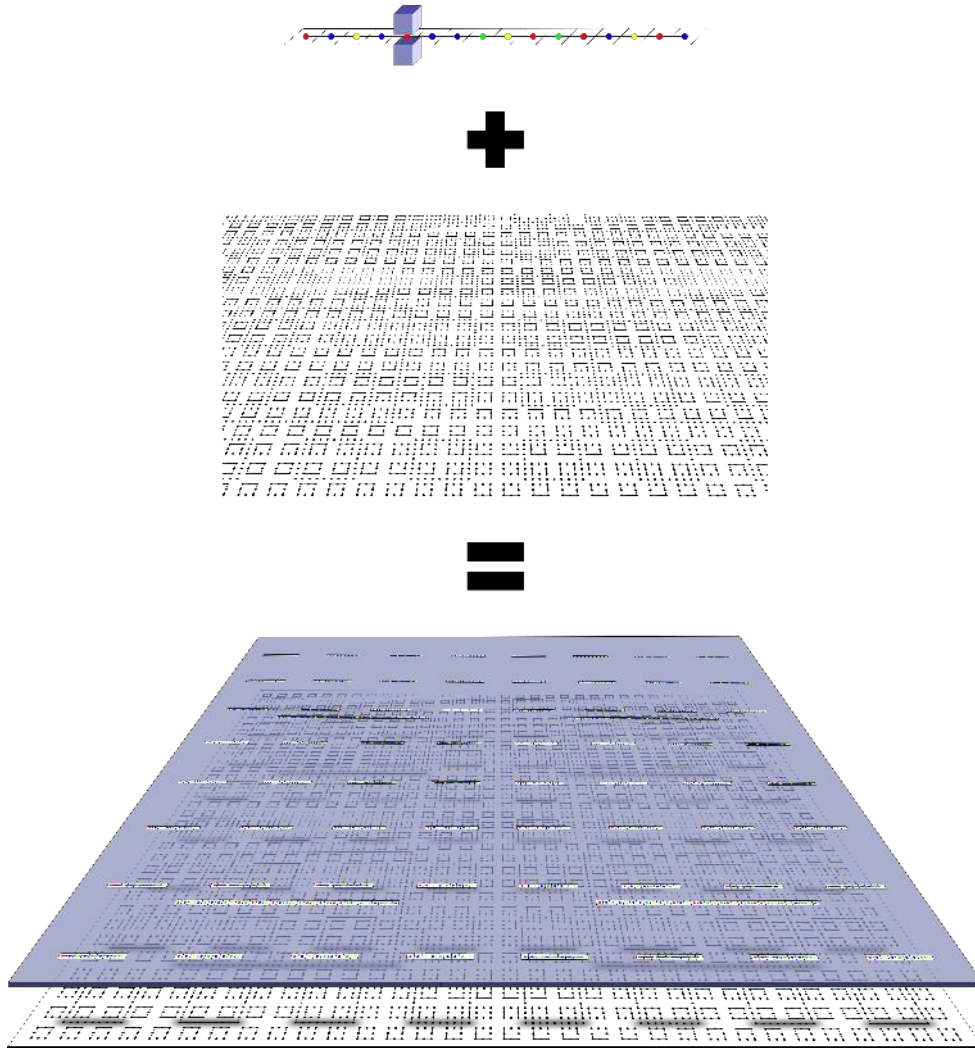


Figure 14: Intuition behind construction in Lemma 51: we “sandwich” together a Gottesman-Irani Hamiltonian and the Robinson tiling Hamiltonian, so that computational history states are forced to appear in the “top” layer above every every square border in the Robinson tiling of the “bottom” layer.

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 $|L\rangle$ and $|R\rangle$ states appearing in the same row of the c -layer, where the q -layer acts like H_q on that line segment. (See Figure 14.)

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:

$$h_{j,j+1}^{\text{col}} = h_c^{\text{col}} \otimes \mathbb{1}_{eq}^{(j)} \otimes \mathbb{1}_{eq}^{(j+1)} \quad (124a)$$

$$h_{i,i+1}^{\text{row}} = h_c^{\text{row}} \otimes \mathbb{1}_{eq}^{(i)} \otimes \mathbb{1}_{eq}^{(i+1)} \quad (124b)$$

$$+ \mathbb{1}_c^{(i)} \otimes \mathbb{1}_c^{(i+1)} \otimes h_q \quad (124c)$$

$$+ |L\rangle\langle L|_c^{(i)} \otimes (\mathbb{1}_{eq} - |\otimes\rangle\langle\otimes|)^{(i)} \otimes \mathbb{1}_{ceq}^{(i+1)} \quad (124d)$$

$$+ (\mathbb{1}_c - |L\rangle\langle L|_c)^{(i)} \otimes |\otimes\rangle\langle\otimes|^{(i)} \otimes \mathbb{1}_{ceq}^{(i+1)} \quad (124e)$$

$$+ \mathbb{1}_{ceq}^{(i)} \otimes |R\rangle\langle R|_c^{(i+1)} \otimes (\mathbb{1}_{eq} - |\oslash\rangle\langle\oslash|)^{(i+1)} \quad (124f)$$

$$+ \mathbb{1}_{ceq}^{(i)} \otimes (\mathbb{1}_c - |R\rangle\langle R|_c)^{(i+1)} \otimes |\oslash\rangle\langle\oslash|^{(i+1)} \quad (124g)$$

$$+ \mathbb{1}_c^{(i)} \otimes |0\rangle\langle 0|_e^{(i)} \otimes |R\rangle\langle R|_c^{(i+1)} \otimes \mathbb{1}_{eq}^{(i+1)} \quad (124h)$$

$$+ |L\rangle\langle L|_c^{(i)} \otimes \mathbb{1}_{eq}^{(i)} \otimes \mathbb{1}_c^{(i+1)} \otimes |0\rangle\langle 0|_e^{(i+1)} \quad (124i)$$

$$+ \mathbb{1}_c^{(i)} \otimes |0\rangle\langle 0|_e^{(i)} \otimes (\mathbb{1}_c - |L\rangle\langle L|_c)^{(i+1)} \otimes (\mathbb{1}_{eq} - |0\rangle\langle 0|_e)^{(i+1)} \quad (124j)$$

$$+ (\mathbb{1}_c - |R\rangle\langle R|_c)^{(i)} \otimes (\mathbb{1}_{eq} - |0\rangle\langle 0|_e)^{(i)} \otimes \mathbb{1}_c^{(i+1)} \otimes |0\rangle\langle 0|_e^{(i+1)}, \quad (124k)$$

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. (124d) and (124e) force a $|\otimes\rangle$ in the q -layer whenever there is an $|L\rangle$ in the c -layer. (124f) and (124g) do the same with $|\oslash\rangle$ and $|R\rangle$. (124h) and (124i) force non-blank to the left and right of an $|R\rangle$ or $|L\rangle$, respectively. Finally, (124j) and (124k) 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 $|L\rangle$ or $|R\rangle$ in the c -layer.

Since h_c is a tiling Hamiltonian, (124b) is by assumption diagonal in the canonical product basis on $\mathcal{H}_c^{\otimes 2}$ (and acts trivially on $(\mathcal{H}_q \oplus \mathcal{H}_e)^{\otimes 2}$). Meanwhile, (124c) is block-diagonal with respect to the four one-dimensional subspaces of $\mathcal{H}_q^{\otimes 2}$ spanned by products of $|\otimes\rangle$ and $|\oslash\rangle$, together with the orthogonal complement thereof. ((124c) acts trivially on $\mathcal{H}_c^{\otimes 2}$.) The remaining terms are manifestly block-diagonal with respect to both of these decompositions simultaneously. The overall Hamiltonian is therefore block-diagonal w.r.t. the product basis on the c -layer. Hence 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. This proves the first claim of the

Lemma.

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 $|L\rangle$ and an $|R\rangle$ (inclusive) in the classical configuration $|T\rangle_c$. Let \mathcal{L}_L denote the set of all horizontal line segments between an L and the right boundary of the region, and similarly \mathcal{L}_R the segments between the left boundary and an R .

Consider first a 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 $\ell \in \mathcal{L}_R \cup \mathcal{L}_L$, and $|0\rangle$ everywhere else in the q -layer. The associated energy $\langle T|_c \langle \psi_0|_q |H| |T\rangle_c |\psi_0\rangle_q$ is clearly $\langle T|H_c|T\rangle + \sum_{\ell \in \mathcal{L}} \lambda_0(|\ell\rangle)$. Indeed $\langle T|H_c|T\rangle$ is the contribution of (124a) and (124b), and $\sum_{\ell \in \mathcal{L}} \lambda_0(|\ell\rangle)$ the contribution of (124c), the rest of the terms being 0 as $|T\rangle_c |\psi_0\rangle_q$ satisfies all constraints imposed by (124d) to (124k).

To see that this is the lowest energy state for a given classical configuration on the c -layer, we define a signature σ for each state in the canonical basis of the q -layer. The local Hilbert space at site i in the q -layer is $\mathcal{H}_e \oplus \mathcal{H}_q$, so we assign $\sigma_i = 0$ for states in \mathcal{H}_e and $\sigma_i = 1$ for states in \mathcal{H}_q . By collecting computational basis states with the same signature, we decompose \mathcal{H}_q into a direct sum of subspaces with given signature,

$$\bigotimes_{i \in \Lambda} \mathcal{H}_q^{(i)} \simeq \bigoplus_{\sigma} \mathcal{H}_{\sigma}. \quad (125)$$

The overall Hamiltonian H is block-diagonal with respect to this decomposition, so all eigenstates $|T\rangle_c |\psi\rangle_q$ have a $|\psi\rangle_q$ part with well-defined signature. We distinguish two cases:

Case 1: $\sigma_i = 1$ for all $i \in \ell \in \mathcal{L}$. Consider a q -layer state $|\psi_{\sigma}\rangle_q$ with this signature; that is, a state supported on the corresponding Hilbert space \mathcal{H}_{σ} . Since all Hamiltonian terms (124d) to (124k) are ≥ 0 , and for all $\ell \in \mathcal{L}$ the (124c) contribution is $\geq \lambda_0(|\ell\rangle)$, we easily get that $\langle T|_c \langle \psi_{\sigma}|_q H |T\rangle_c |\psi_{\sigma}\rangle_q \geq \langle T|H_c|T\rangle + \sum_{\ell \in \mathcal{L}} \lambda_0(|\ell\rangle)$, with equality iff $|\psi_{\sigma}\rangle_q = |\psi_0\rangle_q$.

Case 2: There exists $i \in \ell \in \mathcal{L}$ such that $\sigma_i = 0$. Let j be the rightmost position in ℓ such that $\sigma_j = 0$. If this is the right or left end of ℓ , then the state picks up an energy contribution of 1 from (124d) or (124f). If it is the position next to the right or left end, it picks up a contribution of 1 from (124h) or (124i). Finally, if it is not one of the above cases, it acquires a contribution of 1 from (124k). In all cases, the contribution is $\geq \lambda_0(|\ell\rangle)$ since $\lambda_0(|\ell\rangle) < 1$ by assumption for a Gottesman-Irani Hamiltonian (see Definition 50). As all terms in the Hamiltonian apart from (124a) and (124b) are positive-semidefinite, and the contribution from (124a) and (124b) is $\langle T|H_c|T\rangle$, we have that $\langle T|_c \langle \psi_{\sigma}|_q H |T\rangle_c |\psi_{\sigma}\rangle_q \geq \langle T|H_c|T\rangle + \sum_{\ell \in \mathcal{L}} \lambda_0(|\ell\rangle)$, which completes the proof of the Lemma. \square

With this, we can now prove the following:

Lemma 52 (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 horizontal line segments of the lattice that lie between down/right-facing and down/left-facing red crosses (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 50.

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

$$\lambda_0(H^{\wedge(L)}) = \min_{\mathcal{L} \subset \wedge(L)} \sum_{\ell \in \mathcal{L}} \lambda_0(|\ell|), \quad (126)$$

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

Proof Construct the Hamiltonian H as in Lemma 51, with the red down/right- and down/left-facing crosses from the modified Robinson tile set as the designated $|L\rangle$ and $|R\rangle$ states.

The tiling and QTM Hamiltonians satisfy the requirements of Lemma 51, implying that the lowest energy for a given c -layer configuration is attained by putting $|0\rangle$'s in the q -layer everywhere except in the segments between an $|L\rangle$ and an $|R\rangle$, inclusive. In the modified Robinson tiling, these are exactly the 4^n -segments. Therefore, to bound the ground state energy, we can restrict our attention to classical configurations of Robinson tiles (not necessarily valid tilings) with an eigenstate of h_q along each 4^n -segment.

By Lemma 48, the minimum energy of an eigenstate with a defect-free Robinson tiling on the $L \times H$ rectangle is

$$E(0 \text{ defects}) = \sum_{\ell \in \mathcal{L}} \lambda_0(|\ell|) = \sum_{n=1}^{\lfloor \log_4(L/2) \rfloor} \left(\left\lfloor \frac{H}{2^{2n+1}} \right\rfloor \left(\left\lfloor \frac{L}{2^{2n+1}} \right\rfloor - 1 \right) \right) \lambda_0(4^n), \quad (127)$$

where the final sum is over all segment lengths 4^n . (Recall that we only use the red 4^n -segments to define line-segments on which h_q acts non-trivially.)

On the other hand, since each defect in the classical tile configuration contributes energy at least 1 from the h_c term, Lemma 49 implies that the energy of an eigenstate with d defects on the $L \times H$ rectangle is at least

$$E(d \text{ defects}) = d + \sum_{\ell \in \mathcal{L}} \lambda_0(|\ell|) \geq d + \sum_{n=1}^{\lfloor \log_4(L/2) \rfloor} \left(\left\lfloor \frac{H}{2^{2n+1}} \right\rfloor \left(\left\lfloor \frac{L}{2^{2n+1}} \right\rfloor - 1 \right) - 2d \right) \lambda_0(4^n). \quad (128)$$

Since $\sum_r \lambda_0(r) < 1/2$ by assumption for a Gottesman-Irani Hamiltonian (see Definition 50), we have for all $d > 0$

$$E(d \text{ defects}) - E(0 \text{ defects}) \geq d \left(1 - 2 \sum_{n=1}^{\lfloor \log_4(L/2) \rfloor} \lambda_0(4^n) \right) > 0. \quad (129)$$

Therefore, the lowest energy eigenstate always contains a defect-free tiling, and the Lemma follows. \square

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 53 (Diverging g.s. energy)

There exists 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 $\leq \frac{1}{2}$ 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 $\lambda_0(H_u^{\wedge(L)}(n)) = -L\alpha_1^l(n) + \alpha_0(n)$, or $\lambda_0(H_u^{\wedge(L)}(n)) = L^2\delta_2(n) - L[\alpha_1^u(n) + \delta_1(n)]$, 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 computable number, $h_u^{row}(n)$ $\{0, \beta\}$ -valued and independent of n and

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

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 . Recall that φ is defined as the rational number whose binary decimal expansion contains the digits of n in reverse order after the decimal.

Proof Let h_{q0} be the Hamiltonian obtained by applying Theorem 33 with $K = 3$ to the QTM from Theorem 11 with a proper reversible universal TM dovetailed after it. The Hamiltonian $h_q(n)$ in Lemma 52 will then be $h_q(n) = h_{q0}(n) + |\top\rangle\langle\top| \otimes \mathbb{1} + \mathbb{1} \otimes |q_f\rangle\langle q_f|$, where $|\top\rangle$ is the halting state of the universal TM. h_q has clearly the form given in part (vi) of Theorem 33. We claim that this Hamiltonian is a Gottesman-Irani Hamiltonian according to Definition 50.

The only requirements in Definition 50 that are not immediate are those concerning the minimum eigenvalues restricted to various subspaces. Recall the construction of the computational history state Hamiltonian from Theorem 33. In the initialisation sweep, the $\overset{\curvearrowright}{\circ}$ normally sweeps once from left-to-right along the chain, turns around at the $\overset{\curvearrowleft}{\circ}$ to become a $\overset{\curvearrowright}{\circ}$, which sweeps once right-to-left back along the chain. The superpositions over the standard basis states in these

sequences contain no illegal pairs as long as the other tracks are initialised correctly. However, if the \circlearrowright is missing, then when the \circlearrowright_0 reaches the right end of the chain, there is no forward transition out of the resulting state. Thus the uniform superposition over the first part of the initialisation sweep, involving just the left-to-right sweep, is an eigenstate of h_q . Furthermore, it has 0 energy if the rest of the tracks are correctly initialised. Similarly, if the \circlearrowleft is missing, there is no further forward transition once the \circlearrowleft_0 reaches the left end of the chain, and the superposition over the full initialisation sweep is a 0-energy eigenstate of h_q . Thus condition (v) of Definition 50 is satisfied.

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 33. For $r > 2$, we have

$$\begin{aligned} \lambda_0(r) &\leq \langle \psi | H_q(r) | \psi \rangle = \langle \psi | \left(\sum_i h_{q0}^{(i,i+1)}(n) + |\top\rangle\langle\top|_i \otimes \mathbb{1}_{i+1} + \mathbb{1}_i \otimes |\top\rangle\langle\top|_{i+1} \right) | \psi \rangle \\ &= \sum_{t=1}^T \frac{1}{T} \langle \psi_t | \left(\sum_i |\top\rangle\langle\top|_i \otimes \mathbb{1}_{i+1} + \mathbb{1}_i \otimes |\top\rangle\langle\top|_{i+1} \right) | \psi_t \rangle \leq O\left(\frac{1}{|\Sigma \times Q|^r}\right), \end{aligned} \tag{130}$$

thus $\sum_{m=1}^{\infty} \lambda_0(4^m) < 1/2$. The remaining conditions (iii) and (iv) of Definition 50 are therefore also satisfied. Hence $h_q(n)$ is a Gottesman-Irani Hamiltonian, as claimed.

Let $\tilde{h}_u^{\text{row}}(n), \tilde{h}_u^{\text{column}}(n)$ be the local interactions obtained by applying Lemma 52 to $h_q(n)$. Let $N(n) := \max\{\|h^{\text{row}}(n)\|, \|h^{\text{col}}(n)\|\}$, 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 33 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 11 has sufficient tape and time to finish, and we can be sure that the reversible universal TM starts. 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 Lemma 48, 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 52 gives a ground state energy for

$H^{\wedge(L)}(n)$ of

$$\lambda_0(H^{\wedge(L)}) = \beta \min_{\mathcal{L} \subset \Lambda(L)} \sum_{\ell \in \mathcal{L}} \lambda_0(|\ell|) = \beta \sum_{\substack{1 \leq r \leq |n|+6 \\ r=4^m, m \in \mathbb{N}}} \left[\frac{L}{2r} \right] \left(\left[\frac{L}{2r} \right] - 1 \right) \lambda_0(r) \quad (131a)$$

$$\begin{aligned} &= L^2 \left[\beta \sum_{\substack{1 \leq r \leq |n|+6 \\ r=4^m, m \in \mathbb{N}}} \frac{\lambda_0(r)}{4r^2} \right] - L \left[\beta \sum_{\substack{1 \leq r \leq |n|+6 \\ r=4^m, m \in \mathbb{N}}} \frac{\lambda_0(r)}{2r} \left(2 \operatorname{frac} \left(\frac{L}{2r} \right) + 1 \right) \right] \\ &\quad + \left[\beta \sum_{\substack{1 \leq r \leq |n|+6 \\ r=4^m, m \in \mathbb{N}}} \lambda_0(r) \operatorname{frac} \left(\frac{L}{2r} \right) \left(\operatorname{frac} \left(\frac{L}{2r} \right) + 1 \right) \right] \end{aligned} \quad (131b)$$

$$=: L^2 \alpha_2(n) - L \alpha_1(n, L) + \alpha_0(n, L), \quad (131c)$$

where $\operatorname{frac}(x) := x - \lfloor x \rfloor$ 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 computable algebraic number. We also have

$$\alpha_1^l(n) := \beta \sum_{\substack{1 \leq r \leq |n|+6 \\ r=4^m, m \in \mathbb{N}}} \frac{\lambda_0(r)}{2r} \leq \alpha_1(n, L) < \beta \sum_{\substack{1 \leq r \leq |n|+6 \\ r=4^m, m \in \mathbb{N}}} \frac{3\lambda_0(r)}{2r} =: \alpha_1^u(n), \quad (132)$$

$$0 \leq \alpha_0(n, L) \leq \beta \sum_{\substack{1 \leq r \leq |n|+6 \\ r=4^m, m \in \mathbb{N}}} \frac{\lambda_0(r)}{r} =: \alpha_0(n). \quad (133)$$

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 52, the ground state energy of

$H^{\Lambda(L)}$ is

$$\lambda_0(H^{\Lambda(L)}) = \beta \min_{\mathcal{L} \subset \Lambda(L)} \sum_{\ell \in \mathcal{L}} \lambda_0(|\ell|) \quad (134a)$$

$$= \beta \sum_{\substack{1 \leq r \leq |n|+6 \\ r=4^m, m \in \mathbb{N}}} \left\lfloor \frac{L}{2r} \right\rfloor \left(\left\lfloor \frac{L}{2r} \right\rfloor - 1 \right) \lambda_0(r) + \beta \sum_{\substack{r \geq r_1(n) \\ r=4^m, m \in \mathbb{N}}} \left\lfloor \frac{L}{2r} \right\rfloor \left(\left\lfloor \frac{L}{2r} \right\rfloor - 1 \right) \lambda_0(r) \quad (134b)$$

$$= L^2 \left(\alpha_2(n) + \beta \sum_{\substack{r \geq r_1(n) \\ r=4^m, m \in \mathbb{N}}} \frac{\lambda_0(r)}{4r^2} \right) - L \left(\alpha_1(n) + \beta \sum_{\substack{r \geq r_1(n) \\ r=4^m, m \in \mathbb{N}}} \frac{\lambda_0(r)}{2r} \left(2 \operatorname{frac} \left(\frac{L}{2r} \right) + 1 \right) \right) \\ + \left(\alpha_0(n) + \beta \sum_{\substack{r \geq r_1(n) \\ r=4^m, m \in \mathbb{N}}} \lambda_0(r) \operatorname{frac} \left(\frac{L}{2r} \right) \left(\operatorname{frac} \left(\frac{L}{2r} \right) + 1 \right) \right) \quad (134c)$$

$$=: L^2[\alpha_2(n) + \delta_2(n)] - L[\alpha_1(n, L) + \delta_1(n, L)] + \alpha_0(n, L) + \delta_0(n, L). \quad (134d)$$

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

$$\beta \sum_{\substack{r \geq r_1(n) \\ r=4^m, m \in \mathbb{N}}} \frac{\lambda_0(r)}{2r} \leq \delta_1(n, L) < \beta \sum_{\substack{r \geq r_1(n) \\ r=4^m, m \in \mathbb{N}}} \frac{3\lambda_0(r)}{2r} =: \delta_1(n), \quad (135)$$

$$0 \leq \delta_0(n, L) \leq \beta \sum_{\substack{r \geq r_1(n) \\ r=4^m, m \in \mathbb{N}}} \frac{\lambda_0(r)}{r} := \delta_0(n). \quad (136)$$

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

$$\lambda_0(H^{\Lambda(L)}) = -L\alpha_1(n, L) + \alpha_0(n, L) \leq -L\alpha_1^l(n) + \alpha_0(n). \quad (137)$$

In the halting case, we have

$$\lambda_0(H^{\Lambda(L)}) = L^2\delta_2(n) - L[\alpha_1(n, L) + \delta_1(n, L)] + \alpha_0(n, L) + \delta_0(n, L) \\ \geq L^2\delta_2(n) - L[\alpha_1^u(n) + \delta_1(n)]. \quad (138)$$

The Proposition follows from undecidability of the Halting Problem. \square

The following corollary follows immediately by letting $L(n)$ be the minimal L in Proposition 53 such that $-L\alpha_1^l(n) + \alpha_0(n) \leq 0$ and $L^2\delta_2(n) - L[\alpha_1^u(n) + \delta_1(n)] \geq 1$. This will be useful shortly.

Corollary 54 (Undecidability of g.s. energy with promise) *There exists a family of interactions $h_u^{\text{row}}(n), h_u^{\text{col}}(n) \in \mathcal{B}(\mathbb{C}^U \otimes \mathbb{C}^U)$ and $h_u^{(1)}(n) \in \mathcal{B}(\mathbb{C}^U)$ with operator norm $\leq \frac{1}{2}$ and algebraic matrix entries, and an (uncomputable) function $L(n)$, such that either $\lambda_0(H_u^{\Lambda(L)}(n)) \leq 0$ for all $L \geq L(n)$, or $\lambda_0(H_u^{\Lambda(L)}(n)) \geq 1$ for all $L \geq L(n)$, but determining which is undecidable. Moreover, the interactions can be taken to have the same form as in Proposition 53.*

Undecidability of the ground state energy density is now immediate from Proposition 53 by the definition $E_\rho := \lim_{L \rightarrow \infty} \lambda_0(H^{\Lambda(L)})/L^2$ of the ground state energy density. We restate this result here for convenience:

Theorem 6 (restated) 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 strengths bounded by 1. Then determining whether $E_\rho = 0$ or $E_\rho > 0$ is an undecidable problem. \square

9.2 From g.s. energy density to spectral gap

We are finally in a position to prove our main result: undecidability of the spectral gap, which we restate here for convenience. Recall that for each natural number n , we define $\varphi = \varphi(n)$ to be the rational number whose binary decimal expansion contains the digits of n in reverse order after the decimal.

Theorem 3 (restated) We construct explicitly a dimension d , $d^2 \times d^2$ matrices A, B, C, D and a rational number β so that

- (i). A is Hermitian and with coefficients in $\mathbb{Z} + \beta\mathbb{Z} + \frac{\beta}{\sqrt{2}}\mathbb{Z}$,
- (ii). B, C have integer coefficients,
- (iii). D is Hermitian and with coefficients in $\{0, 1, \beta\}$.

For each natural number n , define:

$$\begin{aligned} h_1(n) &= \alpha_2(n)\mathbb{1} \quad (\alpha_2(n) \text{ an algebraic number}) \\ h_{\text{row}}(n) &= D \quad (\text{independent of } n) \\ h_{\text{col}}(n) &= A + \beta \left(e^{i\pi\varphi} B + e^{-i\pi\varphi} B^\dagger + e^{i\pi 2^{-|\varphi|}} C + e^{-i\pi 2^{-|\varphi|}} C^\dagger \right). \end{aligned}$$

Then:

- (i). The local interaction strength is ≤ 1 . I.e. all terms $h_1(n), h_{\text{row}}(n), h_{\text{col}}(n)$ have operator norm bounded by 1.

- (ii). If UTM halts on input n , then the associated family of Hamiltonians $\{H^{\Lambda(L)}(n)\}$ is gapped in the strong sense of Definition 1 and, moreover, the gap $\gamma \geq 1$.
- (iii). If UTM does not halt on input n , then the associated family of Hamiltonians $\{H^{\Lambda(L)}(n)\}$ is gapless in the strong sense of Definition 2. \square

Proof We assign a Hilbert space $\mathcal{H}^{(i)} := |0\rangle \oplus \mathcal{H}_u \otimes \mathcal{H}_d \simeq \mathbb{C} \oplus \mathbb{C}^U \otimes \mathbb{C}^D$ to each site $i \in \Lambda$, where $|0\rangle$ denotes the generating element of \mathbb{C} . Let h_u be the two-body interaction obtained in Corollary 54 (keeping the one-body term $h_u^{(1)}(n)$ separate, for later convenience).

Let h_d be the normalised (so that $\|h_d\| \leq \frac{1}{2}$) two-body interaction of any nearest-neighbour Hamiltonian H_d with 0 ground state energy whose spectrum becomes dense in the thermodynamic limit (cf. Definition 2): $\lim_{L \rightarrow \infty} \text{spec } H_d^{\Lambda(L)} \rightarrow [0, \infty)$. (For example [LSM61], we can take $D = 2$ and h_d to be the critical XY-model $h_{\text{row}} = \sigma_x \otimes \sigma_x + \sigma_y \otimes \sigma_y + \sigma_z \otimes \mathbb{1} + \mathbb{1} \otimes \sigma_z$ along the rows, where $\sigma_{x,y,z}$ are the Pauli matrices, with no interactions along the columns.)

Define the Hamiltonian $H(n)$ in terms of its two-body and one-body interactions $h(n)$ as follows:

$$h^{(i,j)}(n) := |0\rangle\langle 0|^{(i)} \otimes \Pi_{ud}^{(j)} \quad (139a)$$

$$+ h_u^{(i,j)}(n) \otimes \mathbb{1}_d^{(i)} \otimes \mathbb{1}_d^{(j)} \quad (139b)$$

$$+ \mathbb{1}_u^{(i)} \otimes \mathbb{1}_u^{(j)} \otimes h_d^{(i,j)}, \quad (139c)$$

$$h^{(i)}(n) := -\alpha_2(n) \Pi_{ud}^{(i)}. \quad (139d)$$

Clearly, both have norm bounded by 1.

As in Lemma 51, $\mathbb{1}_u, \mathbb{1}_d$ denote the identity operators on $\mathcal{H}_u, \mathcal{H}_d$, and Π_{ud} denotes the projection of \mathcal{H} onto its $\mathcal{H}_u \otimes \mathcal{H}_d$ subspace. We decompose the global Hamiltonian in the square $\Lambda(L)$ as $H^{\Lambda(L)} =: \tilde{H}_0 + \tilde{H}_u + \tilde{H}_d - \alpha_2 \mathbb{1}_{ud}$, where $\tilde{H}_0, \tilde{H}_u, \tilde{H}_d$ are defined by taking the sum over sites separately for the expressions in (139a) to (139c), respectively. Once again, we assign a *signature* $\sigma \in \{0, 1\}^{L^2}$ to every state in the computational basis, by assigning a 0 to the $|0\rangle$ state, and a 1 to any state in the ud subspace, and decompose the Hilbert space into subspaces of given signature,

$$\bigotimes_{i \in \Lambda} \mathcal{H}^{(i)} \simeq \bigoplus_{\sigma} \mathcal{H}_{\sigma}. \quad (140)$$

Since (139) is manifestly block-diagonal with respect to this direct sum decomposition, we have

$$\text{spec } H^{\Lambda(L)} = \bigcup_{\sigma} \text{spec } H_{\sigma}. \quad (141)$$

Therefore, it is sufficient to analyse the spectra for each signature separately. We distinguish three cases:

- (i). $\sigma = (0, \dots, 0)$: In this sector the Hamiltonian is 0.
- (ii). $\sigma = (1, \dots, 1)$: In this sector $\tilde{H}_0 = 0$ and $\tilde{H}_u + \tilde{H}_d - \alpha_2(n) \sum_i \mathbb{1}_{uq}^{(i)} = H_u^{\Lambda(L)} \otimes \mathbb{1} + \mathbb{1} \otimes H_d^{\Lambda(L)}$, where $H_u^{\Lambda(L)} = \sum_{i,j \in \Lambda} h_u^{(i,j)}(n) - \alpha_2(n) \sum_i \mathbb{1}_u^{(i)}$ is the Hamiltonian from Corollary 54. Thus the spectrum stemming from this subspace is $\text{spec } H_u^{\Lambda(L)} + \text{spec } H_d^{\Lambda(L)}$.

It only remains to analyse the case of mixed signatures:

- (iii). $\sigma \neq (0, \dots, 0)$ but $\sigma_i = 0$ for some i .

Let $|\psi\rangle$ be any state with this signature. Since all $\sigma_i = 0$ sites must be in the state $|0\rangle$, the overall state is of the form $|\varphi\rangle_{\{i:\sigma_i=1\}} \otimes_{j:\sigma_j=0} |0\rangle_j$. We must show that any such state has higher energy than either the $\sigma = (0, \dots, 0)$ or $\sigma = (1, \dots, 1)$ state (whichever is the lower).

Intuitively, 0-signature regions of the lattice reduce the number of QTMs “running” in the lattice in two ways: by reducing the total area that contains QTMs, and by breaking QTMs at boundaries between signature-1 and signature-0 regions. The former reduces the energy contribution from the \tilde{H}_u (139b) by something that scales as $\alpha_2(n) \times \text{area}$. But it also misses out on the negative $-\alpha_2(n) \times \text{area}$ contribution from the $-\alpha_2(n) \mathbb{1}_{ud}$ (139d). So the net effect of the reduction in area occupied by QTMs cancels out. Meanwhile, breaking the QTMs at the boundaries reduces the \tilde{H}_u energy contribution by something that scales as $\sum_r \lambda_0(r) \times \text{length}$, since each length of boundary can break one QTM of each size r . But the boundaries also pick up a positive $1 \times \text{length}$ energy penalty from \tilde{H}_0 ((139a) term). Since $\sum_r \lambda_0(r) < 1$, the energy penalty outweighs the energy reduction. Thus, overall, signature-0 sites within signature-1 region always increase the energy, and the ground state is never a mixed-signature eigenstate.

We now make this rigorous, analysing the halting and non-halting cases separately.¹

Non-halting case: Let M denote the total number of signature-0 (“marked”) sites, U the total number of signature-1 (“unmarked”) sites, and B the total number of boundaries between signature-0 and signature-1 sites. Now divide the $L \times L$ lattice into blocks of size $2r \times 2r$ with $\lfloor L/2r \rfloor$ blocks in each row and column,² where $r = 4^m$, $m \in \mathbb{N}$. We classify blocks according to the signature within the block: *unmarked* blocks, whose signature contains only 1’s; *completely marked* blocks, whose signature contains only 0’s; and *partially marked* blocks, whose signature contains both 0’s and 1’s.

¹The argument is reminiscent of Lemma 49.

²The remaining $L - 2r \lfloor L/2r \rfloor$ columns will play essentially no role.

We first analyse the contribution of the $\tilde{H}_u + \tilde{H}_d$ terms. Consider any horizontal row i of blocks, and let c_i and p_i denote, respectively, the number of completely and partially marked blocks in that row. The number of unmarked blocks in the row is then $\lfloor L/2r \rfloor - c_i - p_i$. Let b_i denote the number of unmarked blocks immediately to the left of a marked block (complete or partial). The marked blocks divide the row into strips of contiguous unmarked blocks. Within each such strip, the Hamiltonian acts as $\tilde{H}_u^{\text{strip}} + \tilde{H}_d^{\text{strip}} \geq H_u^{\text{strip}}$. By Lemma 52, the lowest energy eigenstate of $\tilde{H}_u^{\text{strip}}$ has energy $\min_{\mathcal{L} \subset \Lambda(L)} \sum_{\ell \in \mathcal{L}} \lambda_0(|\ell|)$, where the minimisation is over all valid tilings of the strip. By Lemma 48, the minimum number of r -segments in a strip of $2r \times 2r$ blocks of length l is $l - 1$; in other words, each unmarked block contributes one r -segment, except for blocks at the ends of strips. Thus the energy contribution from the r -segments in the row is at least

$$\left(\left\lfloor \frac{L}{2r} \right\rfloor - c_i - p_i - b_i - 1 \right) \lambda_0(r). \quad (142)$$

Now, each boundary between an unmarked and a marked block (complete or partial) contributes at least one additional vertical boundary between a signature-0 and a signature-1 site. Meanwhile, each partially marked block contributes at least one additional 0/1 signature boundary. Thus, summing over all rows, we have

$$\sum_i p_i + b_i \leq B. \quad (143)$$

Also, each completely marked block must contain at least $4r^2$ signature-0 sites, thus

$$\sum_i c_i \leq \frac{M}{4r^2}. \quad (144)$$

Summing (142) over all $\lfloor L/2r \rfloor$ rows of blocks and over all segment lengths r , we obtain the following lower bound on the total contribution of the $\tilde{H}_u + \tilde{H}_d$ terms:

$$\langle \psi | \tilde{H}_u + \tilde{H}_d | \psi \rangle \geq \sum_{\substack{r=4^m \\ m \in \mathbb{N}}} \sum_i \left(\left\lfloor \frac{L}{2r} \right\rfloor - c_i - p_i - b_i - 1 \right) \lambda_0(r) \quad (145a)$$

$$\geq \sum_r \left\lfloor \frac{L}{2r} \right\rfloor \left(\left\lfloor \frac{L}{2r} \right\rfloor - 1 \right) \lambda_0(r) - B \lambda_0(r) - M \frac{\lambda_0(r)}{4r^2}. \quad (145b)$$

Now consider the energy contribution of the $\tilde{H}_0 - \alpha_2(n) \mathbb{1}_{uq}$ terms. \tilde{H}_0 gives an energy contribution of 1 to every boundary between signature-0 and signature-1 sites, whereas $-\alpha_2(n) \mathbb{1}_{uq}$ gives a negative energy contribution of $-\alpha_2(n)$ to every signature-1 site. Thus the total energy contribution from these terms is

$$\langle \psi | \tilde{H}_0 - \alpha_2(n) \mathbb{1}_{uq} | \psi \rangle = B - \alpha_2(n)(L^2 - M). \quad (146)$$

Putting (145) and (146) together, the total energy of any mixed-signature state is at least

$$\begin{aligned} \langle \psi | H^{\Lambda(L)} | \psi \rangle &= \langle \psi | \tilde{H}_u^{\Lambda(L)} + \tilde{H}_d^{\Lambda(L)} + \tilde{H}_0^{\Lambda(L)} - \alpha_2(n) \mathbb{1} | \psi \rangle \\ &\geq \sum_r \left(\left\lfloor \frac{L}{2r} \right\rfloor \left(\left\lfloor \frac{L}{2r} \right\rfloor - 1 \right) \lambda_0(r) - B \lambda_0(r) - M \frac{\lambda_0(r)}{4r^2} \right) \\ &\quad + B - \alpha_2(n)(L^2 - M) \end{aligned} \quad (147a)$$

$$\begin{aligned} &= \left[\sum_r \left\lfloor \frac{L}{2r} \right\rfloor \left(\left\lfloor \frac{L}{2r} \right\rfloor - 1 \right) \lambda_0(r) - L^2 \alpha_2(n) \right] + B \left(1 - \sum_r \lambda_0(r) \right) \\ &\quad - M \left(\sum_r \frac{\lambda_0(r)}{4r^2} - \alpha_2(n) \right) \end{aligned} \quad (147b)$$

$$> \lambda_0(H_u^{\Lambda(L)}(n)) + \frac{B}{2} - M \sum_{r \geq |n|+6} \frac{\lambda_0(r)}{4r^2}. \quad (147c)$$

In the final inequality we have used the fact that $\alpha_2(n) := \sum_{r < |n|+6} \lambda_0(r)/4r^2$ and $\sum_r \lambda_0(r) < 1/2$, from Proposition 53 and Definition 50.

In the non-halting case, $\lambda_0(r) = 0$ for all $r \geq |n| + 6$. Furthermore, for any mixed-signature state, the total number of 0/1 signature boundaries $B \geq 2$. Thus (147) implies the bound $\langle \psi | H^{\Lambda(L)} | \psi \rangle > \lambda_0(H_u^{\Lambda(L)}) + 1$.

Halting case: The argument is similar to the non-halting case, but we count things slightly differently. Let U denote the total number of signature-1 (“unmarked”) sites, and B the total number of boundaries between signature-0 and signature-1 sites. We again divide the $L \times L$ lattice into blocks of size $2r \times 2r$, with $\lfloor L/2r \rfloor$ blocks in each row and column. Let \mathcal{R}_r be the set of blocked rows that contain only signature-1 sites, with $R_r = |\mathcal{R}_r|$ denoting the size of this set. As before, we categorise the blocks into *unmarked*, *completely marked* and *partially marked*, letting P_r denote the total number of partially marked blocks.

We first analyse the contribution of the $\tilde{H}_u + \tilde{H}_d$ terms. Consider any horizontal row i of blocks. Let u_i denote the number of unmarked blocks in row i , and b_i the number of unmarked blocks which do *not* have an unmarked block to their right (either because they are adjacent to a completely- or partially-marked block, or because they are the last block in the row). By the same argument as above, Lemmas 48 and 52 imply that the total number of r -segments in the i -th row must be at least $u_i - b_i$. Thus the energy contribution from the r -segments in the row is at least

$$(u_i - b_i) \lambda_0(r). \quad (148)$$

Now we count the number of 0/1 signature boundaries. An unmarked block with a marked block on its right contributes at least one boundary between a

signature-1 site to the left of a signature-0 site. An unmarked block at the end of a row does not contribute this boundary. However, in this case, if the row contains at least one marked block, there must exist at least one unmarked block that has a marked block on its *left*, which contributes an additional boundary between a signature-1 site to the *right* of a signature-0 site. Therefore, in any row i that contains at least one marked block, there are at least b_i 0/1 signature boundaries. Finally, in a completely unmarked row we have $b_i = 1$ from the unmarked block at the end of the row.

The number of vertical 0/1 signature boundaries is 0 if a row belongs to \mathcal{R}_r and at least 1 if it does not. Thus rows in \mathcal{R}_r contribute one less 0/1 boundary than the others. Meanwhile, each of the P_r partially marked blocks contributes at least one additional 0/1 boundary. Therefore, summing over all rows i , we have

$$P_r + \sum_i b_i - R_r \leq B. \quad (149)$$

Each unmarked block contains $4r^2$ signature-1 sites, and each of the P_r partially marked blocks contains at most $4r^2 - 1$, out of a total of U . Thus

$$\sum_i u_i > \frac{U - 4r^2 P_r}{4r^2} = \frac{U}{4r^2} - P_r. \quad (150)$$

Finally, each row in \mathcal{R}_r contains $2rL$ signature-1 sites, so

$$R_r \leq \frac{U}{2rL}. \quad (151)$$

Summing (148) over all rows i and over all segment lengths r , and using (149)–(151), we obtain the following lower bound on the total contribution of the $\tilde{H}_u + \tilde{H}_d$ terms:

$$\langle \psi | \tilde{H}_u + \tilde{H}_d | \psi \rangle \geq \sum_{\substack{r=4^m \\ m \in \mathbb{N}}} \sum_i (u_i - b_i) \lambda_0(r) \quad (152a)$$

$$\geq \sum_r \left(\frac{U}{4r^2} - P_r - (B + R_r - P_r) \right) \lambda_0(r) \quad (152b)$$

$$\geq \sum_r \left(\frac{U}{4r^2} - B - \frac{U}{2rL} \right) \lambda_0(r). \quad (152c)$$

Now consider the energy contribution of the $\tilde{H}_0 - \alpha_2(n) \mathbb{1}_{uq}$ terms. As before, \tilde{H}_0 gives an energy contribution of 1 to every boundary between signature-0 and signature-1 sites, whereas $-\alpha_2(n) \mathbb{1}_{uq}$ gives a negative energy contribution of $-\alpha_2(n)$ to every signature-1 site. So the total energy contribution from these terms is

$$\langle \psi | \tilde{H}_0 - \alpha_2(n) \mathbb{1}_{uq} | \psi \rangle = B - \alpha_2(n)U. \quad (153)$$

Putting (152) and (153) together, the total energy of any mixed-signature state is at least

$$\langle \psi | H^{\Lambda(L)} | \psi \rangle = \langle \psi | \tilde{H}_u^{\Lambda(L)} + \tilde{H}_d^{\Lambda(L)} + \tilde{H}_0^{\Lambda(L)} - \alpha_2(n) \mathbb{1} | \psi \rangle \quad (154a)$$

$$\geq \sum_r \left(\frac{U}{4r^2} - B - \frac{U}{2rL} \right) \lambda_0(r) + B - \alpha_2(n)U \quad (154b)$$

$$= U \left(\sum_r \frac{\lambda_0(r)}{4r^2} - \alpha_2(n) \right) + B \left(1 - \sum_r \lambda_0(r) \right) - \frac{U}{L} \sum_r \frac{\lambda_0(r)}{2r} \quad (154c)$$

$$> U \left(\delta_2(n) - \frac{1}{L} \right) + \frac{B}{2}. \quad (154d)$$

In the final inequality, we have used $\alpha_2(n) := \sum_{1 \leq r < |n|+6} \lambda_0(r)/4r^2$, $\delta_2(n) := \sum_{r \geq |n|+6} \lambda_0(r)/4r^2$ and $\sum_r \lambda_0(r) < 1/2$, all from Proposition 53 and Definition 50.

For any mixed-signature state, the total number of 0/1 signature boundaries $B \geq 2$. Moreover, in the halting case, $\delta_2(n) > 0$ independent of L . Thus, for sufficiently large lattice sizes L , (154) implies the bound $\langle \psi | H^{\Lambda(L)} | \psi \rangle > 1$.

Denoting the part of the spectrum which corresponds to these mixed signatures by $S \subset \mathbb{R}$, we obtain that

$$\text{spec } H^{\Lambda(L)} = \{0\} \cup \left\{ \text{spec } H_u^{\Lambda(L)} + \text{spec } H_d^{\Lambda(L)} \right\} \cup S. \quad (155)$$

Recall that the spectrum of H_d becomes dense in the thermodynamic limit:

$$\lim_{L \rightarrow \infty} \text{spec } H_d^{\Lambda(L)} \rightarrow [0, \infty). \quad (156)$$

By taking $L \geq L(n)$ and using Corollary 54, one of two things may happen. If the universal TM does not halt on input n , then $S \geq \lambda_0(H_u^{\Lambda(L)}) + 1$, $\lambda_0(H_u^{\Lambda(L)}) \leq 0$, and $\text{spec } H^{\Lambda(L)}$ becomes dense in $[\lambda_0(H_u^{\Lambda(L)}), \infty)$ as L goes to infinity. If the universal TM *does* halt, then $S \geq 1$ and $\lambda_0(H_u^{\Lambda(L)}) \geq 1$, thus $\Delta(H^{\Lambda(L)}) \geq 1$. Determining which is clearly undecidable, thus concluding the proof of our main Theorem. \square

9.3 Periodic boundary conditions

The previous section proves Theorem 3 for open boundary conditions. This is arguably the most important type of boundary conditions in the context of the spectral gap problem, both physically, because periodic boundary conditions rarely occur in real physical systems, and mathematically, because the thermodynamic limit is better behaved. Nonetheless, our result can also be extended to other types of boundary condition. Extending it to fixed boundary conditions is trivial, so

we omit the argument. In this section, we consider the more interesting case of periodic boundary conditions.

For this, we must first address an issue that does not occur for open boundary conditions. For any local interaction $h \geq 0$ on a lattice with open boundary conditions, it is easy to see¹ that the ground state energy is non-decreasing with the lattice size: for any $L_2 \geq L_1$, $\lambda_0(H(L_1)) \leq \lambda_0(H(L_2))$. However, this is no longer necessarily the case for other types of boundary conditions. In particular, for periodic boundary conditions, there are trivial examples in which the ground state energy oscillates periodically with the lattice size. A simple example is the classical tiling Hamiltonian on a line, with N tile-types and the constraint that tile type n must be to the left of type $n + 1 \bmod N$. For periodic boundary conditions, this has ground state energy 0 for lines of length $L = 0 \bmod N$, but has ground state energy 1 otherwise.

Using a very similar construction to that in the proof of Theorem 3 (see Section 9.2), this oscillatory behaviour of the ground state energy translates into oscillatory behaviour of the spectral gap. If we restrict the lattice size L to the subsequence with $L \bmod N = 0$, then the system is gapless. If we restrict to any subsequence with $L \bmod N \neq 0$, then the system is gapped with spectral gap 1. Our Definitions 1 and 2 of gapped and gapless systems are therefore not adequate for periodic boundary condition.

Since our Definition 2 of gapless systems expresses convergence of the sequence of spectra to a dense subset, the most natural generalisation appropriate to other boundary conditions is to relax this requirement to a convergent *subsequence* of lattice sizes:

Definition 55 (Gapless – generalised) *We say that a family $\{H^{\Lambda(L)}\}$ of Hamiltonians characterises a gapless system if there is a constant $c > 0$ such that for all $\varepsilon > 0$ there is an $L_0 \in \mathbb{N}$ and a subsequence $\{L_i\}$ of \mathbb{N} such that for all $L_i > L_0$ every point in $[\lambda_0(H^{\Lambda(L_i)}), \lambda_0(H^{\Lambda(L_i)}) + c]$ is within distance ε from $\text{spec } H^{\Lambda(L_i)}$.*

With this generalised definition of gapless systems, Theorem 3 proving undecidability of the spectral gap generalises to the case of periodic boundary conditions. The proof involves slightly modifying the Robinson tiling part of our construction, using a tiling trick from [GI09]. We first restate and prove this trick here (in a slightly generalised form).

The construction works for any (unweighted) tile set. First, duplicate all the tiles to create black- and white-coloured variants of each tile, and impose the constraint that adjacent tiles must have different colours. Turn this into an equivalent weighted

¹If $|\psi\rangle$ is a ground state with energy $E_0 = \sum_{\langle i,j \rangle \in \Lambda(L_2) \times \Lambda(L_2)} \text{Tr}_{\Lambda(L_2)} (h^{(i,j)} |\psi\rangle\langle\psi|)$ for the lattice of size L_2 with open boundary conditions, then the mixed state $\rho = \text{Tr}_{i \notin \Lambda(L_1)} |\psi\rangle\langle\psi|$ on the lattice of size $L_1 \leq L_2$ has energy $\text{Tr}_{\Lambda(L_1)} (\sum_{\langle i,j \rangle \in \Lambda(L_1) \times \Lambda(L_1)} h^{(i,j)} \rho) = \text{Tr}_{\Lambda(L_2)} (\sum_{\langle i,j \rangle \in \Lambda(L_1) \times \Lambda(L_1)} h^{(i,j)} |\psi\rangle\langle\psi|) \leq E_0$.

tiling by defining the weight of adjacent pairs of matching tiles to be 0, and that of non-matching pairs to be +5. (Note that adjacent pairs can fail to match either because the original tile markings fail to match, or because their colours are the same.) Finally, add to the tile set the three new tile types shown in Figure 15 (which do not have additional colour markings).




Figure 15: Boundary tiles for periodic boundary condition construction.

The weights for adjacent pairs of tiles in this extended tile set are given in the following table, where \square and \blacksquare denote any tiles of the given colour¹ (the weight for adjacent pairs of opposite-colour tiles is 0 if the original tile markings match, +5 if they do not match):

		Right tile				
		\square	\blacksquare	$\begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}$	$\begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}$	$\begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}$
Left tile	\square	5	0/5	5	0	5
	\blacksquare	0/5	5	5	0	5
	$\begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}$	5	5	0	5	1
	$\begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}$	0	0	5	5	5
	$\begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}$	5	5	1	5	5

		Tile above				
		\square	\blacksquare	$\begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}$	$\begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}$	$\begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}$
Tile below	\square	5	0/5	0	5	5
	\blacksquare	0/5	5	0	5	5
	$\begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}$	0	0	5	5	5
	$\begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}$	5	5	5	0	1
	$\begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}$	5	5	5	1	5

(157)

The following Lemma encapsulates the construction given in [GI09, § 4.2]. It shows that these weights effectively turn the weighted tiling problem for an odd-sized square region with periodic boundary conditions (a torus), into the unweighted tiling problem for the original tile set on an even-sized square region with open boundary conditions.

Lemma 56 (Periodic to open b.c.) *Given any unweighted tile set and any odd $L \in \mathbb{N}$, all minimum weight tilings of an $L \times L$ torus using the extended tile set described above, with the weights given in (157), consist of a single $\begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}$ tile at an arbitrary location, a complete row of $\begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}$ tiles extending horizontally from the $\begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}$, a complete column of $\begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}$ tiles extending vertically from the $\begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}$, and a valid tiling of the remaining $(L - 1) \times (L - 1)$ region with the original tile set using alternating colours.*

¹The choice of weights is somewhat arbitrary. Here, we choose them slightly more symmetrically than in [GI09], which makes the proof of Lemma 56 marginally simpler to describe.

Proof The tilings described in the Lemma have total weight $+4$ (coming from the \square and \square tiles adjacent to the \square). We must show that all other tilings have higher weight. Since L is odd, the torus cannot be tiled consistently with alternating colours, thus any weight < 5 tiling must contain at least one \square , \square or \square .

If the tiling contains any tile other than \square or \square to the left or right of a \square , it has weight ≥ 5 . Thus, in any weight ≤ 5 tiling, \square tiles can only appear in complete rows of \square and \square tiles. The analogous argument for columns implies that \square tiles can only appear in complete rows of \square and \square tiles. A \square adjacent to a black or white tile has weight $+5$, so \square tiles can only appear where such rows and columns meet. Moreover, adjacent \square and \square tiles have weight $+5$, so *all* such rows and columns must meet at \square tiles.

Therefore, any tiling with weight < 5 must consist of some non-zero number of rows of \square tiles and columns of \square tiles meeting at \square tiles, with coloured tiles everywhere else. Each intersection of a \square row with a \square column at a \square contributes weight $+4$, so the weight is minimised by having just one such row and column.

Finally, the remaining region to be filled with coloured tiles is a square of size $(L - 1) \times (L - 1)$. $L - 1$ is even, so this can be tiled consistently with alternating colours. It contributes weight 0 iff, in addition, the tiles (ignoring the colour markings) form a valid tiling for the original tile set. \square

Applying Lemma 56 to the modified Robinson tiles of Section 8, we obtain a weighted tile set whose minimum weight tilings on a torus of size L have exactly the same quasi-periodic structure as the original Robinson tilings on a square of size $L - 1$, for any odd lattice size L . Using this tile set in Lemma 51, the proofs of Lemmas 51 and 52, Proposition 53, and Corollary 54 go through unchanged for odd-sized square lattices with periodic boundary conditions. Undecidability of the ground state energy density (Theorem 6), and undecidability of the spectral gap (Theorem 3) under the appropriate Definition 55 of gapless, follow immediately.

10 Acknowledgements

TSC would like to thank IBM T. J. Watson Laboratory for their hospitality during many visits, and Charlie Bennett in particular for insightful discussions about aperiodic tilings and reversible Turing Machines. Also, for pointing out that he [Charlie] had wrestled with implementation details of reversible Turing Machines thirty years ago, and it was time a younger generation took over! TSC also thanks Ashley and Catherine Montanaro for their kind hospitality in lodging him for two months in their house in Cambridge, where part of this work was carried out.

TSC, DPG and MMW thank the Isaac Newton Institute for Mathematical Sciences, Cambridge for their hospitality during the programme “Mathematical Challenges in Quantum Information”, where part of this work was carried out. TSC and DPG are both very grateful to MMW and the Technische Universität München for their hospitality over summer 2014, when another part of the work was completed.

TSC is supported by the Royal Society. DPG acknowledges support from MINECO (grant MTM2011-26912), Comunidad de Madrid (grant QUITEMAD+CM, ref. S2013/ICE-2801), and financial support from the John von Neumann Professorship program of the Technische Universität München. DPG and MMW acknowledge support from the European CHIST-ERA project CQC (funded partially by MINECO grant PRI-PIMCHI-2011-1071).

This work was made possible through the support of grant #48322 from the John Templeton Foundation. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the John Templeton Foundation.

References

- [Aff+88] I. Affleck et al. “Valence bond ground states in isotropic quantum antiferromagnets”. *CMP* 115 (3 1988), p. 477.
- [Aha+07] D. Aharonov et al. “Adiabatic quantum computation is equivalent to standard quantum computation”. *SIAM Journal of Computing* 37 (2007), pp. 166–194.
- [Aha+09] Dorit Aharonov et al. “The power of quantum systems on a line”. *Commun. Math. Phys.* 287 (1 2009). arXiv:0705.4077v2 [quant-ph], p. 41.
- [Aiz98] M. Aizenman, ed. *Open Problems in Mathematical Physics*. <https://web.math.princeton.edu/~aizenman/OpenProblems.iamp/>. 1998-1999.

- [And73] P. W. Anderson. “Resonating valence bonds: a new kind of insulator?": *Mat. Res. Bull.* 8 (1973), pp. 153–160.
- [Bal10] L. Balents. “Spin liquids in frustrated magnets”. *Nature* 464 (2010), pp. 199–208.
- [Bar77] Jon Barwise, ed. *Handbook of mathematical logic*. Elsevier, 1977.
- [Ben73] Charles H. Bennett. “Logical reversibility of computation”. *IBM J. Res. Develop.* 17 (1973), p. 525.
- [Ber66] R. Berger. “The undecidability of the domino problem”. *Mem. Amer. Math. Soc.* 66 (1966), pp. 1–72.
- [BS93] Lenore Blum and Steve Smale. “The Gödel Incompleteness Theorem and Decidability over a Ring”. English. *From Topology to Computation: Proceedings of the Smalefest*. Ed. by Morris W. Hirsch, Jerrold E. Marsden, and Michael Shub. Springer US, 1993, pp. 321–339. ISBN: 978-1-4612-7648-7. DOI: [10.1007/978-1-4612-2740-3_32](https://doi.org/10.1007/978-1-4612-2740-3_32). URL: http://dx.doi.org/10.1007/978-1-4612-2740-3_32.
- [BV97] Ethan Bernstein and Umesh Vazirani. “Quantum complexity theory”. *SIAM Journal on Computing* 26.5 (1997), p. 1411. URL: www.cs.berkeley.edu/~vazirani/pubs/bv.ps.
- [Far+00] E. Farhi et al. “Quantum computation by adiabatic evolution”. [arXiv: quant-ph/0001106](https://arxiv.org/abs/quant-ph/0001106). 2000.
- [Fey85] Richard Feynman. “Quantum mechanical computers”. *Optics News* 11 (1985), p. 11.
- [FG+12] Carlos Fernández-González et al. “Frustration free gapless Hamiltonians for Matrix Product States”. *arXiv:1210.6613* (2012).
- [GI09] Daniel Gottesman and Sandy Irani. *The quantum and classical complexity of translationally invariant tiling and Hamiltonian problems*. arXiv:0905.2419[quant-ph]. 2009.
- [Gu+09] Mile Gu et al. “More Really is Different”. *Physica D* 238 (2009), pp. 835–839.
- [Hal83] F.D.M. Haldane. “Nonlinear field theory of large-spin Heisenberg anti-ferromagnets: Semiclassically quantized solitons of the one-dimensional easy-axis Néel state”. *PRL* 50 (1983), p. 1153.
- [Han+12] T.-H. Han et al. “Fractionalized excitations in the spin-liquid state of a Kagome-lattice antiferromagnet”. *Nature* 492 (2012), pp. 406–410.
- [Has04] M.B. Hastings. “Lieb-Schultz-Mattis in higher dimensions”. *Phys. Rev. B* 69 (2004), p. 104431.

- [Has07] M.B. Hastings. “An area law for one-dimensional quantum systems”. *J. Stat. Mech.: Theory and Experiment* 2007 (08 2007).
- [Hof76] Douglas R. Hofstadter. “Energy levels and wave functions of Bloch electrons in rational and irrational magnetic fields”. *Phys. Rev. B* 14 (6 Sept. 1976), pp. 2239–2249. doi: [10.1103/PhysRevB.14.2239](https://doi.org/10.1103/PhysRevB.14.2239). URL: <http://link.aps.org/doi/10.1103/PhysRevB.14.2239>.
- [Hol73] Alexander S. Holevo. “Bound for the quantity of information transmitted by a quantum communication channel”. *Problems of Information Transmission* 9 (1973), pp. 177–183.
- [JW00] Arthur Jaffe and Edward Witten. *Quantum Yang-Mills Theory*. <http://www.claymath.org/sites/default/files/yangmills.pdf>. 2000.
- [KKR06] J. Kempe, A. Kitaev, and O. Regev. “The complexity of the Local Hamiltonian problem”. *SIAM Journal of Computing* 35 (5 2006), p. 1070.
- [Koz97] Dexter Kozen. *Automata and computability*. Springer Science & Business Media, 1997.
- [KSV02] A.Y. Kitaev, A.H. Shen, and M.N. Vyalyi. *Classical and Quantum Computation*. Vol. 47. AMS, 2002.
- [Llo93] S. Lloyd. “Quantum-mechanical computers and uncomputability”. *PRL* 71 (1993), pp. 943–946.
- [LSM61] E.H. Lieb, T.D. Schultz, and D.C. Mattis. “Two-Dimensional Ising Model as a Soluble Problem of Many Fermions”. *Ann. Phys.* 16 (1961), p. 407.
- [LVV13] Z. Landau, U. Vazirani, and T. Vidick. “A polynomial-time algorithm for the ground state of 1D gapped local Hamiltonians”. (2013). arXiv:1307.5143v1.
- [MH06] T. Komatsu, M.B. Hastings. “Spectral Gap and Exponential Decay of Correlations”. *CMP* 265 (2006), pp. 781–804.
- [Mię97] J. Miękisz. “Stable quasiperiodic ground states”. *J. Stat. Phys.* 88 (1997), pp. 691–711.
- [MP13] S. Michalakis and J. Pytel. “Stability of Frustration-Free Hamiltonians”. *CMP* 322 (2013), pp. 277–302.
- [NC00] Michael A. Nielsen and Isaac L. Chuang. *Quantum computation and quantum information*. Cambridge University Press, 2000.

- [NO02] Harumichi Nishimura and Masanao Ozawa. “Computational Complexity of Uniform Quantum Circuit Families and Quantum Turing Machines”. *Theor. Comput. Sci.* 276 (2002), p. 147. eprint: [quant-ph/9906095](https://arxiv.org/abs/quant-ph/9906095).
- [OA01] D. Osadchy and J. E. Avron. “Hofstadter butterfly as quantum phase diagram”. *Journal of Mathematical Physics* 42.12 (2001), pp. 5665–5671. DOI: [http://dx.doi.org/10.1063/1.1412464](https://doi.org/10.1063/1.1412464). URL: <http://scitation.aip.org/content/aip/journal/jmp/42/12/10.1063/1.1412464>.
- [OT08] R. Oliveira and B. Terhal. “The complexity of quantum spin systems on a two-dimensional square lattice”. *QIC* 8 (2008). arXiv:0504050[quant-ph], p. 0900.
- [PER89] Marian B. Pour-El and Jonathan I. Richards. *Computability in Analysis and Physics*. Springer, 1989.
- [Poo12] Bjorn Poonen. *Undecidable Problems: A Sampler*. arXiv:1204.0299. 2012.
- [Rad62] Tibor Rado. “On non-computable functions”. *Bell System Technical Journal* 41 (1962), pp. 877–884.
- [Rob71] R. Robinson. “Undecidability and Nonperiodicity for Tilings of the Plane”. *Invent. Math.* 12 (1971), p. 177.
- [Tur36] Alan Turing. “On computable numbers, with an application to the Entscheidungsproblem”. *Proc. London Math. Soc.* 2nd ser. 42 (1936–7).
- [Wan61] H. Wang. “Proving theorem by pattern recognition”. *Bell System Tech. J.* 40 (1961), pp. 1–41.
- [YHW11] S. Yan, D. A. Huse, and S. R. White. “Spin-liquid ground state of the S=1/2 Kagome Heisenberg antiferromagnet”. *Science* 332 (2011), pp. 1173–1176.