

RESEARCH ARTICLE

Undecidability of the Spectral Gap

Toby Cubitt¹, David Perez-Garcia^{2,3} and Michael M. Wolf^{4,5}

¹Department of Computer Science, University College London, Gower Street, London WC1E 6EA, United Kingdom; E-mail: t.cubitt@ucl.ac.uk.

²Department of Applied Mathematics and Mathematical Analysis, Facultad de CC Matemáticas, Universidad Complutense de Madrid, Plaza de Ciencias 3, 28040 Madrid, Spain.

³ICMAT, C/Nicolás Cabrera, Campus de Cantoblanco, 28049 Madrid, Spain; E-mail: dperezga@ucm.es.

⁴Department of Mathematics, Technische Universität München, 85748 Garching, Germany.

⁵Munich Center for Quantum Science and Technology (MCQST), Germany; E-mail: m.wolf@tum.de.

Received: 11 June 2020; **Revised:** 29 July 2021; **Accepted:** 12 October 2021

Keywords: quantum many-body physics, quantum computation, undecidability, spectral theory

2020 Mathematics Subject Classification: *Primary* – 81Q10; *Secondary* – 03D35, 81P68, 03D10, 03D80, 05b45

Abstract

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. Moreover, this constant can be taken equal to the operator norm of the local operator that generates the Hamiltonian (the local interaction strength). The result still holds true if one restricts to arbitrarily small quantum perturbations of classical Hamiltonians. The proof combines a robustness analysis of Robinson’s aperiodic tiling, together with tools from quantum information theory: the quantum phase estimation algorithm and the history state technique mapping Quantum Turing Machines to Hamiltonians.

Contents

| | | |
|----------|----------------------------------------------------|-----------|
| 1 | Introduction | 2 |
| 1.1 | Definitions and notation | 2 |
| 1.2 | Main results | 3 |
| 1.3 | Relation to physics | 4 |
| 1.4 | Brief overview of the proof | 7 |
| 2 | Unconstrained local Hilbert space dimension | 8 |
| 2.1 | Undecidability of the spectral gap via tiling | 8 |
| 2.2 | Undecidability of low energy properties | 11 |
| 3 | Quantum Phase Estimation Turing Machine | 15 |
| 3.1 | Quantum Turing Machinery | 16 |
| 3.2 | Quantum phase estimation overview | 33 |
| 3.3 | Preparation stage | 33 |
| 3.4 | Control- U_φ stage | 34 |
| 3.5 | Locating the LSB | 42 |

| | | |
|----------|--------------------------------------------|-----------|
| 3.6 | QFT stage | 46 |
| 3.7 | Reset Stage | 47 |
| 3.8 | Analysis | 47 |
| 4 | Encoding QTMs in local Hamiltonians | 48 |
| 4.1 | Preliminaries | 52 |
| 4.2 | Clock Oscillator | 56 |
| 4.3 | Initialisation sweep | 57 |
| 4.4 | Clock Counter | 57 |
| 4.5 | Clock Hamiltonian | 69 |
| 4.6 | QTM Hamiltonian | 71 |
| 4.7 | Analysis | 77 |
| 5 | Quasi-periodic tilings | 83 |
| 5.1 | Robinson's tiling | 83 |
| 5.2 | Rigidity of the Robinson tiling | 86 |
| 6 | Putting it all together | 90 |
| 6.1 | Undecidability of the g.s. energy density | 90 |
| 6.2 | From g.s. energy density to spectral gap | 96 |
| 6.3 | Periodic boundary conditions | 98 |

1. Introduction

1.1. Definitions and notation

Let us start by defining the setup we will be considering, which is the standard framework in mathematical physics for describing (finite) *spin lattice models*. We denote by $\Lambda(L) := \{1, \dots, L\}^2$ the set of vertices (referred to as *sites*) of a square lattice of side $L \in \mathbb{N}$, with $L \geq 2$. $\mathcal{E} \subset \Lambda(L) \times \Lambda(L)$ will denote the set of directed edges of the square lattice, oriented so that $(i, j) \in \mathcal{E}$ implies that j lies north or east of i . We will consider both *periodic* and *open* boundary conditions. In the first case, the outer rows and columns are also connected along the same direction, so that $\Lambda(L)$ becomes a square lattice on a torus. In the second case, these connections are not present.

We associate to each site $i \in \Lambda(L)$ a Hilbert space $\mathcal{H}^{(i)} \simeq \mathbb{C}^d$, and to any subset $S \subseteq \Lambda(L)$ the tensor product $\bigotimes_{i \in S} \mathcal{H}^{(i)}$. The *interactions* between neighbouring pairs $(i, j) \in \mathcal{E}$ are given by Hermitian operators $h^{(i,j)} \in \mathcal{B}(\mathcal{H}^{(i)} \otimes \mathcal{H}^{(j)})$. In addition, we may assign an *on-site interaction* given by a Hermitian matrix $h_1^{(k)} \in \mathcal{B}(\mathcal{H}^{(k)})$ to each site $k \in \Lambda(L)$.

We will restrict in this paper to models that are built up from such nearest-neighbour and possibly on-site terms in a translationally invariant way. That is, when identifying the isomorphic Hilbert spaces on which they act, $h_1^{(k)} = h_1^{(l)}$ for all $k, l \in \Lambda(L)$, and $h^{(i',j')} = h^{(i,j)}$ whenever there is a $v \in \mathbb{Z}^2$ such that $(i', j') = (i + v, j + v)$. The total Hamiltonian for the system is then defined as

$$H^{\Lambda(L)} := \sum_{(i,j) \in \mathcal{E}} h^{(i,j)} + \sum_{k \in \Lambda(L)} h_1^{(k)}. \tag{1.1}$$

It is fully described for any system size L 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 define the interactions between neighbouring sites within any row and column respectively. Hence, it may alternatively be written as

$$H^{\Lambda(L)} = \sum_{\text{rows}} \sum_c h_{\text{row}}^{(c,c+1)} + \sum_{\text{columns}} \sum_r h_{\text{col}}^{(r,r+1)} + \sum_{i \in \Lambda(L)} h_1^{(i)}. \tag{1.2}$$

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

The set of eigenvalues, or *energy levels*, of the Hamiltonian $H^{\wedge(L)}$ will be denoted by $\text{spec } H^{\wedge(L)} := \{\lambda_0(H^{\wedge(L)}), \lambda_1(H^{\wedge(L)}), \dots\}$. When the Hamiltonian is clear from context, we will sometimes omit it and just write $\{\lambda_0, \lambda_1, \dots\}$. They are always assumed to be listed in non-decreasing order $\lambda_0 \leq \lambda_1 \leq \dots$. The smallest eigenvalue $\lambda_0(H^{\wedge(L)})$ is called the *ground state energy* and the corresponding eigenvectors *ground states*. The *ground state energy density* is defined as

$$E_\rho := \lim_{L \rightarrow \infty} E_\rho(L), \quad \text{where } E_\rho(L) := \frac{\lambda_0(H^{\wedge(L)})}{L^2}. \tag{1.3}$$

It is not difficult to show [22] that this limit is well defined.

$H^{\wedge(L)}$ is called *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. $H^{\wedge(L)}$ is called *classical* if its defining interactions $h^{(i,j)}, h_1^{(k)}$ are diagonal in a given product basis (e.g. the canonical one).

We can define now the main quantity under study: the *spectral gap*

$$\Delta(H^{\wedge(L)}) := \lambda_1(H^{\wedge(L)}) - \lambda_0(H^{\wedge(L)}). \tag{1.4}$$

In this paper we are considering the behaviour of $\Delta(H^{\wedge(L)})$ in the thermodynamic limit, that is, when $L \rightarrow \infty$. For that, we introduce 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). 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 $\lim_{L \rightarrow \infty} c(L) = \infty$.

The ideas and techniques for the proof will mainly come from quantum information theory. We therefore use notation standard to that field, such as Dirac’s notation for linear algebra operations. A summary of all the relevant standard notation can be found, for example, in Chapter 2 of the classic book of Nielsen and Chuang [62].

1.2. Main results

For each natural number n , we define $\varphi = \varphi(n)$ to be the rational number whose binary fraction expansion consists of the digits of n in reverse order after the decimal point. We also fix throughout a specific Universal Turing Machine, denoted by UTM.

Theorem 3 (Main theorem). *For any given universal Turing Machine UTM, we can construct explicitly a dimension d , $d^2 \times d^2$ matrices A, A', B, C, D, D' a $d \times d$ diagonal projector Π and a rational number β which can be as small as desired, with the following properties:*

- (i) A is diagonal with entries in \mathbb{Z} .
- (ii) A' is Hermitian with entries in $\mathbb{Z} + \frac{1}{\sqrt{2}}\mathbb{Z}$,
- (iii) B, C have integer entries,
- (iv) D is diagonal with entries in \mathbb{Z} ,
- (v) D' is hermitian with entries in \mathbb{Z} .

For each natural number n , define:

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

where $\alpha(n) \leq \beta$ is an algebraic number computable from n . Then:

- (i) The local interaction strength is bounded by 1, i.e. $\max(\|h_1(n)\|, \|h_{\text{row}}(n)\|, \|h_{\text{col}}(n)\|) \leq 1$.
- (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.

Some comments are in order:

- Using the classic result of Turing [72] that the Halting Problem for UTM on input n is undecidable, we can conclude that *the spectral gap problem is also undecidable*.
- The interaction terms in the Hamiltonians given in the Theorem are β -perturbations of the classical Hamiltonian given by $h_{\text{row}} = A$, $h_{\text{col}} = D$. Since β can be taken as small as desired, all interactions appearing in the Theorem are arbitrarily small quantum perturbations of a classical system.
- The only dependency on n in the interactions appears in some *prefactors* to a set of fixed interactions.
- In the gapped case, the size of the gap is larger than or equal to the local interaction strength, hence is *as large as possible*.
- It is straightforward (if tedious) to extract an explicit value for d in Theorem 3 and Corollary 4 from the construction described in this paper.
- By exploiting the well known connection between (algorithmic) undecidability and (axiomatic) independence (see e.g. [68]) one immediately obtains the following corollary:

Corollary 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.*

Moreover, as a key intermediate step in the proof of Theorem 3, we will prove that the ground state energy density problem is also undecidable, a result clearly interesting on its own.

Theorem 5 (Undecidability of g.s. energy density). *Let $d \in \mathbb{N}$ be sufficiently large but fixed. We can explicitly construct a one-parameter family of 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 for which determining whether $E_\rho = 0$ or $E_\rho > 0$ is an undecidable problem.*

A short version of this paper – including a statement of the main result, discussion of its implications, an outline of the main ideas behind the proof, together with a sketch of the argument – was published recently in Nature [22]. We encourage the reader to consult it in order to gain some high-level intuition about the full, rigorous proof given in this work.

1.3. Relation to physics

As already discussed in [22], this result has a number of implications for condensed matter and mathematical physics, that we briefly recall here.

Quantum spin lattice models are ubiquitous in mathematical physics. They play a central role in condensed matter physics, in quantum computation and even in high energy physics, where one route to formalising the continuum is to consider quantum many body systems on a lattice and later sending the lattice spacing to zero. In all these contexts, one often starts from a description of the microscopic interactions that govern the system. The question is then how to infer the relevant observable macroscopic properties that emerge as the lattice size (number of particles) tends to infinity.

The spectral gap is amongst the most important properties of a quantum many-body Hamiltonian. It is intimately related to the physics of a quantum many-body system. Quantum phase transitions can only occur at critical points where the gap vanishes. The intuitive reason is that the spectral gap “protects” the ground state properties of the system against small perturbations, since an energy of the order of the gap must be pumped into the system to transition to a different state. Formalising this intuition is a major open question. Recently Bravyi, Hastings and Michalakis [19] and Michalakis and Pytel [57] proved this for the particular case of frustration-free systems satisfying certain additional conditions, by proving that the spectral gap of these systems is stable under arbitrary local perturbations throughout the system.

The low-temperature behaviour of a system is also governed by the spectral gap: gapped systems exhibit “non-critical” behaviour, with low-energy excitations that behave as massive particles [35], preventing long-range correlations [42, 61]; gapless systems exhibit “critical” behaviour, with low-energy excitations that behave as massless particles, allowing long-range correlations. This implies that ground states of gapped systems are somehow less complex. That intuition has been formalised in the *area law conjecture* [26], which has been proven for 1D spin systems [39], and in 2D if certain additional hypothesis on the spectrum are satisfied [40]. This in turn translates into better algorithms for computing properties of such systems [65]. A paradigmatic example of this line of research is the recent proof [51, 7], based on an improved version of the 1D area law theorem [6], that the ground state energy density problem for gapped quantum spin systems in 1D is in the complexity class P (i.e. the run-time of the algorithm scales polynomially in the system size).

Because of its central importance to many-body quantum systems, many seminal results in mathematical physics concern the spectral gap of specific systems. Examples include the Lieb-Schultz-Mattis theorem showing that the Heisenberg chain for half-integer spins is gapless [53], extended to higher dimensional bipartite lattices by Hastings [38], and the proof of a spectral gap for the 1D AKLT model [1]. More recently, shortly after this work first appeared on preprint servers, Bravyi and Gosset gave a complete characterisation of the spectral gap for the simplest non-trivial case of 1D chains of spin- $\frac{1}{2}$ particles (qubits) with frustration-free, nearest-neighbour interactions [18]. (The frustrated case remains open.)

The same is true of some of the most important open questions in the field. For instance, the *Haldane conjecture* [36], first formulated in 1983, states that the integer-spin antiferromagnetic Heisenberg model in 1D has a non-vanishing spectral gap. Despite strong supporting evidence from numerical simulations, a proof remains elusive. The analogous question in 2D for non-bipartite lattices can be traced back to the work of Anderson in 1973 [5], where he suggested the existence of new types of materials, nowadays called topological quantum spin liquids. The existence of minerals in nature, such as herbertsmithite, whose interactions can be well approximated by the spin- $\frac{1}{2}$ antiferromagnetic Heisenberg model on the Kagome lattice – hence are compelling candidates for systems with a topological quantum spin liquid phase [37] – brought the problem of its spectral gap to the forefront of physics. The existence of a spectral gap in these systems is an active area of research, with techniques reducing the spectral gap problem to finite-volume criteria very recently being applied successfully (partially numerically) to certain cases such as the honeycomb lattice [52, 67], whilst other cases remain disputed even at the level of numerical evidence (see e.g. [76, 44]).

Undecidability of the spectral gap has a number of immediate corollaries relevant to physics. It implies that one can write down models whose phase diagrams are so complex they are in fact uncomputable. It implies the standard approach of trying to gain insight into physics models by solving numerically for larger and larger lattice sizes necessarily fails for some systems; the system can 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. (In Section 2 we also construct models exhibiting the opposite transition, from gapped to gapless.) Not only can the threshold size be arbitrarily large; the threshold size itself is an uncomputable number. In a recent paper [11], we constructed very simple models with small local dimensions which exhibit this type of “size-driven phase transition” (without, however, the undecidability properties proven here). Our findings also imply that a result showing robustness of the spectral gap under perturbations, as for the case of frustration-free Hamiltonians [57] and for free-fermion systems [24, 41], 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 [66, 43]. Since membership in many fractal sets is not decidable (when formulated in the framework of real computation; see [17]), 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 in quantum field theory, 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.

The idea that some of the most difficult open problems in physics could be mathematically proven to be impossible to solve is not new. Indeed, proving the impossibility of obtaining exact formulae for certain physical quantities in natural physical models 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 [3]. Apart from the pioneering work of Komar in this direction in 1964 [49], stating undecidable properties in Quantum Field Theories, and the influential paper of Anderson entitled “More is different” [4] in 1972, most of the initial connections between physical problems and undecidability arose in the 80s and 90s. In 1981, [69] studied non-computable solutions to the wave equation. One year later, [31] found undecidable questions in models of hard frictionless balls. In 1984, Wolfram wrote the paper “Undecidability and Intractability in Theoretical Physics” where, motivated by the emergent complexity of very simple cellular automata, he conjectured that many natural problems in physics, and in particular in classical statistical physics, should be undecidable. Many results in this direction – using cellular automata to prove undecidability results in physical problems – have been shown since then (e.g [25] or [64]). The work of Gu, Weedbrook, Perales and Nielsen [34] in 2009 found undecidable properties in the 2D classical Ising model. Besides cellular automata, another natural connection between undecidability and physics, that we will exploit in this paper, came from tiling problems, shown to be undecidable by the works of Wang and Berger in the 60s [74, 15]. This idea has been exploited for instance in 1990 by Kanter [45], finding undecidable properties in anisotropic 1D Potts Hamiltonians. That same year, in a completely different direction, Moore wrote one of the most influential papers on undecidability in physics [59], proving undecidability of the long-term behaviour of a particle-in-a-box problem. (See also [14] for a nice commentary on that paper.)

In recent years, mainly motivated by quantum information theory and the link it established between physics and computer science, there has been a revival of interest in undecidability in quantum physics. Results in this direction have appeared in several contexts, such as measurement and control [27, 75], tensor networks [60, 48, 21], measurement-based quantum computation [73], channel capacities [28], or Bell inequalities [12, 71].

A precursor of those results, due to Lloyd, appeared already in 1993 [54, 55] in the early days of quantum information theory. There, it is argued that quantum computing implies certain spectral properties of quantum mechanical operators are undecidable. More precisely, the unitary evolution U associated to the evolution of a computer (classical or quantum) capable of universal computation has invariant subspaces with discrete spectrum (roots of unity) and other invariant subspaces with continuous spectrum (the whole unit sphere), corresponding respectively to computations that halt and

do not halt. Then, since the halting problem is undecidable, Lloyd concludes that given a quantum state associated with a program in the infinite dimensional space in which U is defined, it is undecidable to know whether it has overlap with an invariant subspace having discrete spectrum or, conversely, it is supported on an invariant subspace in which the spectrum is the full unit circle. See [56, 23] for a recent discussion between the relation between Lloyd's result and the result of this paper.

In subsequent follow-up work, we have extended the undecidability results described here to 1D systems [8] and to uncomputability of phase diagrams [9], and have used insights from our undecidability results to construct simple, explicit examples of the new type of "size-driven" phase transition implied by undecidability of the spectral gap [11].

We refer to a forthcoming review paper in collaboration with Gu and Perales [20] for an extensive analysis and discussion of all results mentioned above, and many more, related to undecidability in physics.

1.4. Brief overview of the proof

An extended overview and discussion of the ideas behind our proof can be found in the Supplementary Material of [22]. We will simply sketch here the four main steps of the proof, whose rigorous statements and proofs are given, respectively, in Sections 3 to 6.

Step 1: write n on the tape.

In order to feed an arbitrary input $n \in \mathbb{N}$ to the UTM, and at the same time keep a uniform upper bound on the local dimension of the Hilbert space of each site in the final Hamiltonian, we need to construct a Quantum Turing Machine (QTM) with a fixed number of internal states such that, when fed with an input given by any sequence of 1's larger than the size of n , writes n deterministically on the tape and halts. We need also to keep strict control on the time and space required for this computation. The precise result we prove along these lines is given in Theorem 10, and Section 3 is devoted to proving it. The idea is to construct explicitly the QTM associated to the *quantum phase estimation algorithm*. To do this, we rely heavily on results and ideas from [16].

Step 2: construct, for each $n \in \mathbb{N}$, a 1D Hamiltonian whose ground state energy on a finite chain depends on the behaviour of the UTM on input n .

For this, we heavily rely on the remarkable paper of Gottesman and Irani [32] (see also [10]), which can be seen as a milestone in a long history of papers [47, 46, 63, 2], dating back to Feynman [30], relating the ground state energy of a Hamiltonian with the time-evolution of a quantum computation. In order to achieve the desired properties of our construction, we have to modify the original construction of Gottesman and Irani, but most of the essential ideas of this step appear already in [32].

After these two steps, we have for each $n \in \mathbb{N}$ a 1D translationally invariant Hamiltonian whose ground state energy for a chain of size L and open boundary conditions is 0 if the UTM does not halt on n in space less than $O(L)$ and time less than $O(c^L)$ (for any desired constant c). Otherwise, it has energy larger than c^{-L} . So there the difference in ground state energy depending on the halting behaviour of the UTM on n vanishes exponentially fast with chain length, and is zero in the thermodynamic limit. Theorem 32 states the precise result we obtain, and Section 4 is devoted to its proof.

Step 3: amplify the ground state energy difference.

For this, we turn to 2D lattices, and exploit the properties of an aperiodic tiling of the 2D plane due to Robinson. The idea is to construct a Hamiltonian whose ground state mimics the tiling pattern of Robinson's tiling shown in Figure 9 and, at the same time, places the ground state of the 1D Hamiltonian constructed in Step 2 on top of each of the 1D borders appearing in this pattern. Using the fact that there is a constant density of squares of size 4^r for all $r \in \mathbb{N}$, and making a shift in energies, we manage to show for the resulting Hamiltonian that, if the UTM halts on n , the ground state energy diverges to $+\infty$, whereas in the non-halting case it diverges to $-\infty$. From there, we conclude undecidability of the ground state energy density. Section 5 takes care of proving all the new results we require for Robinson's tiling. The rest of the proof, together with the final step, appears in Section 6.

Step 4: from ground state energy difference to spectral gap.

The final step combines the Hamiltonian H_u constructed in Step 3 with two others: a trivial Hamiltonian having ground state energy 0 and a constant spectral gap, and a critical Hamiltonian H_d having ground state energy 0 and a spectrum that becomes dense in $[0, \infty)$ as the system size goes to infinity. We couple these Hamiltonians in such a way that the spectrum of the resulting Hamiltonian on $\Lambda(L)$ is

$$\{0\} \cup \left[\text{spec}(H_u^{\wedge(L)}) + \text{spec}(H_d^{\wedge(L)}) \right] \cup S_L,$$

with $S_L \geq \min \left\{ 1, 1 + \lambda_0(H_u^{\wedge(L)}) \right\}$. This, together with the spectral properties of H_u shown in Step 3, completes the proof.

2. Unconstrained local Hilbert space dimension

Before starting in Section 3 on the proof of the main theorem, in this section we will describe two approaches that exploit known undecidability results for tiling and completion problems. Based on these, we can derive simpler but weaker undecidability results for the spectral gap and for other low energy properties of translationally invariant, nearest-neighbour Hamiltonians on a 2D square lattice, defined by their local interactions $\{(h_{\text{row}}(n), h_{\text{col}}(n))\}_{n \in \mathbb{N}}$. In contrast to the main theorem of this paper, however, the families of Hamiltonians that we construct here have local Hilbert space dimension d_n that differs for different elements n of the family, with no upper bound on $\{d_n\}_{n \in \mathbb{N}}$.

2.1. 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:

2.1.1. Ingredient 1: a tiling Hamiltonian

We start by recalling the notion of Wang tilings and tiling problems. A unit square whose edges are coloured with colours chosen from a finite set is called a *Wang tile*. A finite set \mathcal{K} of Wang tiles is said to *tile the plane* \mathbb{Z}^2 if there is an assignment $\mathbb{Z}^2 \rightarrow \mathcal{K}$ such that abutting edges of adjacent tiles have the same colour. (Rotations or reflections of the tiles are not allowed here.) It is a classic result of Berger [15] that, given any set of tiles as input, determining whether or not this set can tile the plane is undecidable.

There is an easy way to rewrite any tiling problem as a ground state energy problem for a classical Hamiltonian. If $\mathcal{K} = \{1, \dots, K\}$ we assign a Hilbert space $\mathcal{H}^{(i)} \simeq \mathbb{C}^K$ to each site i of a square lattice, and define the local interactions via

$$h_c^{(i,j)} := \sum_{(m,n) \in C^{(i,j)}} |m\rangle\langle m|_{(i)} \otimes |n\rangle\langle n|_{(j)}, \quad (2.1)$$

where 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 . The overall Hamiltonian on the lattice $\Lambda(L)$ is then simply

$$H_c^{\wedge(L)} := \sum_{(i,j) \in \mathcal{E}} h_c^{(i,j)}. \quad (2.2)$$

By construction, we have $h_c^{(i,j)} \geq 0$, $H_c^{\wedge(L)}$ is translationally 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

$0 \in \text{spec } H_c^{\Lambda(L)}$ 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

$$\forall L > L_0 : \text{spec } H_c^{\Lambda(L)} \geq 1. \tag{2.3}$$

This is due to Wang's extension theorem (see e.g. [33]).

2.1.2. Ingredient 2: a gapless frustration-free Hamiltonian

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

$$H_q^{\Lambda(L)} := \sum_{(i,j) \in \mathcal{E}} h_q^{(i,j)}, \tag{2.4}$$

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}_+ [29]. 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.

2.1.3. 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

$$H^{(i,j)} := |0\rangle\langle 0|^{(i)} \otimes \mathbb{1}_{cq}^{(j)} + \mathbb{1}_{cq}^{(i)} \otimes |0\rangle\langle 0|^{(j)} \tag{2.5}$$

$$+ \sum_{(m,n) \in \mathcal{C}^{(i,j)}} |m\rangle\langle m|^{(i)} \otimes \mathbb{1}_q^{(i)} \otimes |n\rangle\langle n|^{(j)} \otimes \mathbb{1}_q^{(j)} \tag{2.6}$$

$$+ \sum_{\alpha, \beta, \gamma, \delta=1}^D \mathbb{1}_c^{(i)} \otimes |\alpha\rangle\langle \beta|^{(i)} \otimes \mathbb{1}_c^{(j)} \otimes |\delta\rangle\langle \gamma|^{(j)} \langle \alpha, \delta | h_q^{(i,j)} | \beta, \gamma \rangle. \tag{2.7}$$

Here $\mathbb{1}_c, \mathbb{1}_q$ and $\mathbb{1}_{cq}$ 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 on the square lattice $\Lambda(L)$ as

$$H^{\Lambda(L)} := \sum_{(i,j) \in \mathcal{E}} H^{(i,j)}. \tag{2.8}$$

Theorem 6 (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 the excitation spectrum 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.*

$$\text{spec } H^{\Lambda(L)} \setminus \{0\} \geq 1. \tag{2.9}$$

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 (2.5), (2.6) and (2.7), respectively. Let us now assign a *signature* $\sigma \in \{0, \dots, K\}^{L^2}$ to every state of our computational product basis, in the following way: $|0\rangle^{(i)}$ is assigned the signature $\sigma_i = 0$, and $|k, \alpha\rangle^{(i)}$ is assigned the signature $\sigma_i = k$ irrespective of α . By collecting computational basis states with the same signature we can then decompose the Hilbert space as

$$\bigotimes_{i \in \Lambda} \mathcal{H}^{(i)} \simeq \bigoplus_{\sigma} \mathcal{H}_{\sigma}. \tag{2.10}$$

The Hamiltonian is block-diagonal w.r.t. this decomposition, i.e. it can be written as

$$H^{\Lambda(L)} = \bigoplus_{\sigma} H_{\sigma}, \quad \text{and thus} \quad \text{spec } H^{\Lambda(L)} = \bigcup_{\sigma} \text{spec } H_{\sigma}. \tag{2.11}$$

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

$$\text{spec } H^{\Lambda(L)} = \{0\} \cup \left\{ \text{spec } H_c^{\Lambda(L)} + \text{spec } H_q^{\Lambda(L)} \right\} \cup S, \tag{2.12}$$

The equivalence between the impossibility of tiling with \mathcal{K} and the existence of a spectral gap of size at least 1 now follows immediately from the properties of the ingredients h_c and h_q . 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. □

From the undecidability of the tiling problem we now obtain the following corollary. Recall from Section 1.1 that a pair of Hermitian matrices $(h_{\text{row}}, h_{\text{col}})$ defines a Hamiltonian $H^{\Lambda(L)}$ for each square $\Lambda(L)$ via (1.2), where we are now choosing open boundary conditions.¹

Corollary 7 (Undecidability of the spectral gap for unconstrained dimension). *Let $\epsilon > 0$. There is no algorithm that, on input $(h_{\text{row}}, h_{\text{col}})$ with rational entries and operator norm smaller than $1 + \epsilon$, decides whether the associated family of Hamiltonians $\{H^{\Lambda(L)}\}$ describe a gapped or a gapless system (according to the definitions given in Section 1.1). This holds even under the promise that one of these is true, that $H^{\Lambda(L)}$ is frustration free for all Λ , and that in the gapped case the spectral gap is at least 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 arbitrarily 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.

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

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 any given consistent formal system with a recursive set of axioms, one can construct a specific 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.² However, the local dimension, whilst finite, depends on the choice of axiomatic system and there is no upper-bound on the values it can take. The main aim of this paper is to prove the much stronger Theorem 3, which shows that the spectral gap problem remains undecidable even for a fixed value of the local Hilbert space dimension. This in turn implies that for any consistent, recursive formal system, axiomatic independence holds for Hamiltonians with this fixed local dimension; the local dimension is now independent of the choice of axioms.

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.

2.2. Undecidability of low energy properties

2.2.1. 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 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. A drawback is that we lose the frustration-free property. On the other hand, we do not have to rely on the undecidability of tiling (which is a rather non-trivial result [15, 70]), but can exploit the simpler undecidability of the *completion problem*, already proven by Wang [74]. In the tiling completion problem, one fixes a tile at the left-bottom corner and asks whether the first quadrant can be tiled.

Undecidability of the completion problem is relatively simple by reduction from the Halting Problem for Turing Machines (TM), and explained in full detail in Robinson [70, Section 4]. There are many different, computationally equivalent definitions of Turing Machines. (The Halting Problem is undecidable for any of these variants.) For the completion problem, the most convenient is to take a Turing Machine to be specified by a finite number of states $Q = \{q_0, q_1, \dots\}$ with q_0 the initial state, a finite alphabet of tape symbols $S = \{s_0, s_1, \dots\}$ with s_0 the blank symbol, together with a finite set of transition rules specified by quintuples of the form

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

The TM has a half-infinite tape of cells indexed by \mathbb{N} , and a single read/write tape head that moves along the tape. The machine starts in state q_0 with the head in tape cell 0. If the TM is currently in state q and reads the symbol s from the current tape cell, then it writes s' into the current tape cell, transitions to state q' , and moves in the direction D . For each q, s , there must be at most one valid quintuplet. If there is none, then the machine halts.³

Following Robinson [70], we use the Wang tiles shown in Figure 1. For clarity, edges here are marked with labelled arrows rather than colours, and adjacent edges match if arrow heads abut arrow tails with the same label; clearly these tile markings can be identified with different colours to recover a set of standard Wang tiles. We denote the set of tiles by \mathcal{K} , which contains at most $K := |\mathcal{K}| \leq 2 + (3|Q| + 1)|S|$ elements.

If the tile that appears in the bottom-left corner in Figure 1, 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

²An explicit example can, for instance, be constructed from the explicit Turing Machines of [77] for which the halting problem is independent of ZFC

³Note that this defines a slightly different variant of Turing Machines from those used later in Section 3.




Figure 1. 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.

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} .

Assign a weight $w_{(i,j)}(m,n) \in \mathbb{Z}$ to 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 weights is summarised in the following table. Here \square stands for any tile different from \square , the unbracketed numbers define the weights for non-matching adjacent tiles, and the numbers in brackets define the weights for cases where the tiles match (i.e. all abutting labels and arrows match).

| | | | | | | | | | | | | | | | | | | | | | |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|------------|------|-----------|-----------|-----------|-----|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|------------|---|---|-----------|-----------|-----------|---|---|-----|--------|
| | | <table style="border-collapse: collapse;"> <tr> <td></td> <td style="border-right: 1px solid black; padding: 2px;">Right tile</td> <td></td> </tr> <tr> <td></td> <td style="border-right: 1px solid black; padding: 2px;">\square</td> <td style="padding: 2px;">\square</td> </tr> </table> | | Right tile | | | \square | \square | | <table style="border-collapse: collapse;"> <tr> <td></td> <td style="border-right: 1px solid black; padding: 2px;">Tile above</td> <td></td> </tr> <tr> <td></td> <td style="border-right: 1px solid black; padding: 2px;">\square</td> <td style="padding: 2px;">\square</td> </tr> </table> | | Tile above | | | \square | \square | | | | | |
| | Right tile | | | | | | | | | | | | | | | | | | | | |
| | \square | \square | | | | | | | | | | | | | | | | | | | |
| | Tile above | | | | | | | | | | | | | | | | | | | | |
| | \square | \square | | | | | | | | | | | | | | | | | | | |
| Left tile | <table style="border-collapse: collapse;"> <tr> <td style="border-right: 1px solid black; padding: 2px;">\square</td> <td style="padding: 2px;">4</td> <td style="padding: 2px;">2</td> <td style="padding: 2px;">(-1)</td> </tr> <tr> <td style="border-right: 1px solid black; padding: 2px;">\square</td> <td style="padding: 2px;">4</td> <td style="padding: 2px;">2</td> <td style="padding: 2px;">(0)</td> </tr> </table> | \square | 4 | 2 | (-1) | \square | 4 | 2 | (0) | <table style="border-collapse: collapse;"> <tr> <td style="border-right: 1px solid black; padding: 2px;">Tile below</td> <td style="border-right: 1px solid black; padding: 2px;">\square</td> <td style="padding: 2px;">4</td> <td style="padding: 2px;">2</td> <td style="padding: 2px;">(-1)</td> </tr> <tr> <td style="border-right: 1px solid black; padding: 2px;">\square</td> <td style="border-right: 1px solid black; padding: 2px;">\square</td> <td style="padding: 2px;">4</td> <td style="padding: 2px;">2</td> <td style="padding: 2px;">(0)</td> </tr> </table> | Tile below | \square | 4 | 2 | (-1) | \square | \square | 4 | 2 | (0) | (2.13) |
| \square | 4 | 2 | (-1) | | | | | | | | | | | | | | | | | | |
| \square | 4 | 2 | (0) | | | | | | | | | | | | | | | | | | |
| Tile below | \square | 4 | 2 | (-1) | | | | | | | | | | | | | | | | | |
| \square | \square | 4 | 2 | (0) | | | | | | | | | | | | | | | | | |

Assign 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

$$h_u^{(i,j)} := \sum_{m,n \in \mathcal{K}} w_{(i,j)}(m,n) |m\rangle\langle m|_{(i)} \otimes |n\rangle\langle n|_{(j)}, \tag{2.14}$$

so that positive and negative weights become energy ‘penalties’ or ‘bonuses’ in the Hamiltonian. 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 \square is the only one that can lead to negative weights. Assume first that a tile \square 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, \square 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 the Turing machine does not halt within L time-steps. Hence $\lambda_0(H_u^{\Lambda(L)}) = -2$, the ground state will be unique, and will be 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 \square 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 tile types at the top-left of Figure 1), we obtain:

Lemma 8 (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 translationally 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 the Turing Machine does not halt within L time-steps, then $\lambda_0(H_u^{\Lambda(L)}) = -2$. In this case the corresponding ground state is a product state, and is the unique eigenstate with negative energy.*
- (ii) *If the Turing machine halts, then $\exists L_0$ (given by the halting time) such that $\forall L > L_0 : \lambda_0(H_u^{\Lambda(L)}) = 0$.*

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.

2.2.2. Reduction of the halting problem to arbitrary low energy properties

Consider translationally invariant Hamiltonians on square lattices with open boundary conditions. Let $H_q^{\Lambda(L)}$ 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 9 (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 translationally invariant Hamiltonians on square lattices with open boundary conditions. Let $H_q^{\Lambda(L)}$ with nearest-neighbour interaction $h_q^{(i,j)} \in \mathcal{B}(\mathcal{H}_q^{(i)} \otimes \mathcal{H}_q^{(j)})$, $\mathcal{H}_q^{(i)} \simeq \mathbb{C}^q$, and on-site term $h_{q_1}^{(k)} \in \mathcal{B}(\mathcal{H}_q^{(k)})$ describe such a Hamiltonian. Assume that there is a $c \in \mathbb{R}$ such that $\lambda_0(H_q^{\Lambda(L)}) - cL^2 \in [-\frac{1}{2}, \frac{1}{2}]$ for all lattice sizes L . Then we can construct a family of Hamiltonians $H^{\Lambda(L)}$ with the following properties:*

- (i) *$H^{\Lambda(L)}$ is translationally 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 $H^{\Lambda(L)}$ are nearest-neighbour (possibly with on-site terms), computable and independent of L .*
- (iii) *If the TM does not halt within L time-steps, then $H^{\Lambda(L)}$ has a non-degenerate product ground state and a spectral gap $\Delta(H^{\Lambda(L)}) \geq \frac{1}{2}$.*
- (iv) *If TM halts, then $\exists L_0$ (determined by the halting time) such that $\forall L > L_0$ the following identity between multisets holds in the interval $[0, \frac{1}{2}]$:*

$$\text{spec } H_q^{\Lambda(L)} - \lambda_0(H_q^{\Lambda(L)}) = \text{spec } H^{\Lambda(L)} - \lambda_0(H^{\Lambda(L)}). \tag{2.15}$$

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, (ii) and (iii) 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.

We first embed the Hamiltonian given in the Theorem in a larger system, such that its ground state energy is shifted by -1 . Define a translationally invariant, nearest-neighbour Hamiltonian on an auxiliary system with local Hilbert space $\mathcal{H}_a^{(i)} \simeq \mathbb{C}^2$, via

$$\eta := \sum_{(i,j) \in \mathcal{E}} \eta^{(i,j)}, \quad \eta^{(i,j)} := -\frac{1}{2} |0\rangle\langle 0|^{(i)} \otimes |1\rangle\langle 1|^{(j)} + 2\mathbb{1}^{(i)} \otimes |0\rangle\langle 0|^{(j)}. \tag{2.16}$$

Note that this is nothing but $\frac{1}{2} \times$ the Hamiltonian defined by (2.13) (with the number in brackets), where $|0\rangle$ equals \square and $|1\rangle$ is the blank tile. Hence $\lambda_0(\eta) = -1$ 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 $\mathcal{H}_Q^{(i)} \otimes \mathcal{H}_Q^{(j)}$. 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 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

$$h^{(i,j)} := h_Q^{(i,j)} + h_u^{(i,j)} + 6\mathbb{1}_Q^{(i)} \otimes \mathbb{1}_u^{(j)} + 6\mathbb{1}_Q^{(j)} \otimes \mathbb{1}_u^{(i)}, \tag{2.17}$$

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

$$H^{\Lambda(L)} = \sum_{(i,j) \in \mathcal{E}} h^{(i,j)} + \sum_{k \in \Lambda(L)} \mathbb{1}_a^{(k)} \otimes h_{q_1}^{(k)}. \tag{2.18}$$

To analyse the ground state of $H^{\Lambda(L)}$, notice that all the terms in the Hamiltonian $\eta, h_q, h_u, \mathbb{1}_Q \otimes \mathbb{1}_u$ commute with each other. Let $\tilde{H}_q = \sum_{(i,j) \in \mathcal{E}} h_q^{(i,j)} + h_{q_1}^{(i,j)}$, $\tilde{H}_\eta = \sum_{(i,j) \in \mathcal{E}} \eta^{(i,j)} + 2\mathbb{1}_Q^{(i)} \otimes \mathbb{1}_u^{(j)} + 2\mathbb{1}_Q^{(j)} \otimes \mathbb{1}_u^{(i)}$, and $\tilde{H}_u = \sum_{(i,j) \in \mathcal{E}} h_u^{(i,j)} + 4\mathbb{1}_Q^{(i)} \otimes \mathbb{1}_u^{(j)} + 4\mathbb{1}_Q^{(j)} \otimes \mathbb{1}_u^{(i)}$. We can then decompose $H^{\Lambda(L)} = \tilde{H}_q + \tilde{H}_u + \tilde{H}_\eta$ where the three Hamiltonians commute with each other.

We now assign a signature $\sigma \in \{0, 1, 2\}^{L^2}$ to every state of our computational product basis, so that $|i\rangle^{(j)}$ is identified with $\sigma_j \in \{0, 1\}$ for all $|i\rangle \in \mathcal{H}_Q^{(j)}$ of the form $|\sigma_j\rangle_a |\alpha_j\rangle_q$ and $\sigma_j = 2$ for all $|i\rangle \in \mathcal{H}_u^{(j)}$. By collecting computational basis states with the same signature, we can then decompose the Hilbert space as

$$\bigotimes_{i \in \Lambda} \mathcal{H}^{(i)} \simeq \bigoplus_{\sigma} \mathcal{H}_{\sigma}. \tag{2.19}$$

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

$$H^{\Lambda(L)} = \bigoplus_{\sigma} H_{\sigma}, \quad \text{and thus} \quad \text{spec } H^{\Lambda(L)} = \bigcup_{\sigma} \text{spec } H_{\sigma}. \tag{2.20}$$

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

Case 1:

$\sigma = (2, \dots, 2)$. Restricted to this sector (i.e. the \mathcal{H}_{σ} component of the direct sum), the Hamiltonian acts as $H_u^{\Lambda(L)} = \sum_{(i,j) \in \mathcal{E}} h_u^{(i,j)}$. But according to Lemma 8, if up to time L the TM did not halt, then $H_u^{\Lambda(L)}$ has one negative eigenvalue -2 whose associated eigenstate is a product. However, if $L \geq L_0$ where L_0 is related to the halting time, then $\text{spec } H_u^{\Lambda(L)} \geq 0$.

Case 2:

$\sigma_j = 1$ for all $j \in \Lambda(L)$, except in the lowest left corner, where $\sigma_j = 0$. In this sector we have the following identity between multisets:

$$\text{spec } H_{\sigma} = \text{spec } H_q^{\Lambda(L)} - 1. \tag{2.21}$$

Case 3:

Any other signature σ . It is easy to see that the energy given by \tilde{H}_{η} in the sector \mathcal{H}_{σ} is exactly the energy given to the configuration σ by the tiling Hamiltonian (2.14) with colours $\{0, 1, 2\}$ and weights given by the following tables from (2.13), hence the energy is ≥ 0 .

| | | | | |
|-----------|------------|---|------------|---|
| | Right tile | | Tile above | |
| | 0 | 2 | 0 | 2 |
| Left tile | 0 | 2 | 1 | 2 |
| | 1 | 2 | 1 | 2 |
| | 2 | 2 | 2 | 0 |
| | | | Tile below | |
| | | | 0 | 2 |
| | | | 1 | 2 |
| | | | 2 | 2 |

(2.22)

Similarly, one can show that $\tilde{H}_u \geq 0$. Therefore, in this Case 3, the energy of $H^{\Lambda(L)}$ in the sector \mathcal{H}_{σ} is $\geq \lambda_0(H_q^{\Lambda(L)})$. Theorem 9 now follows straightforwardly. □

3. Quantum Phase Estimation Turing Machine

We will see in Section 4 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 $n \in \mathbb{N}$, we can choose the transition rules such that the QTM, when started from a sequence of N ones (N any number larger than the number of bits $|n|$ in the binary representation of n) writes out the binary representation of n to its tape and then halts deterministically.

Given n , we recall from Section 1.2 that $\varphi = \varphi(n)$ is defined as the rational number whose binary fraction expansion contains the digits of n in reverse order after the decimal.⁴

The QTM that we construct will be based on the quantum phase estimation algorithm [78, 62], applied to the single qubit unitary $U_{\varphi} = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi\varphi} \end{pmatrix}$, to obtain the phase φ . The input N can be any upper

⁴The fact that the digits are in reverse order is purely for convenience in the exposition – it is trivial to modify the construction so that they are ordered the other way around.

bound on the number of binary digits in the output of the quantum phase estimation algorithm, such that φ can be represented exactly as a binary fraction, with no approximation error. Note that the transition rules of the QTM are *not* allowed to depend on N , though they can depend on φ (and hence on n).

The important role played by N can be sketched as follows. Firstly, N fixes the number of output qubits to be used in the quantum phase estimation algorithm, such that an exact binary fraction representation of φ is possible. But N has a more subtle use in the last step of the quantum phase estimation algorithm, where one needs to implement an inverse quantum Fourier transform (QFT). To implement such a QFT in the standard way on the full set of output qubits, one would require quantum gates implementing phase rotations by angles that depend on N . But this would require N -dependent QTM transition rule amplitudes, which is not allowed. The solution is to use N as a way to locate the least significant bit of φ , which will in turn allow the QFT to be implemented using gates that depend only on φ . (This is explained in more detail when we give the precise description and analysis of the QTM, later in this section.)

This way of carrying out the QFT on a QTM comes at a cost of exponential run-time, so would not be useful if using the QFT to solve some computational problem. But for our peculiar use of the phase-estimation algorithm to generate an input string for a universal TM, which may anyway run indefinitely, this exponential run-time is not an issue. (Though it will entail a new clock construction when we come to encode the evolution of this QTM in a Hamiltonian, see Section 4.)

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

Theorem 10 (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 any n , and for each choice of states p, q , alphabet symbols σ, τ and directions D , the transition amplitude $\delta(p, \sigma, \tau, q, D)$ is one of the elements of the set*

$$\left\{ 0, 1, \pm \frac{1}{\sqrt{2}}, e^{i\pi\varphi}, e^{i\pi 2^{-|\varphi|}} \right\} \tag{3.1}$$

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

We emphasise again 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 10 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 10.

3.1. Quantum Turing Machinery

For completeness we quote the basic definitions of Turing Machines and Quantum Turing Machines verbatim from [16].

Definition 11 (Turing Machine – Definition 3.2 in [16]). 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 is a transition function

$$\delta : Q \times \Sigma \rightarrow \Sigma \times Q \times \{L, R\}. \quad (3.2)$$

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.

As amplitudes in quantum mechanics can take arbitrary complex values, one has to be careful when defining quantum Turing Machines to ensure the transition amplitudes are themselves computable. Following [16, Definition 3.2], we restrict the transition function to computable numbers in the standard way. Let $\tilde{\mathbb{C}}$ be the set consisting of α such that there is a deterministic classical Turing Machine M_α that, given input $n \in \mathbb{N}$, outputs the real and imaginary parts of α to within 2^{-n} in time polynomial in n .⁵

Definition 12 (Quantum Turing Machine – Definition 3.2 in [16]). 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

$$\delta : Q \times \Sigma \rightarrow \tilde{\mathbb{C}}^{\Sigma \times Q \times \{L, R\}}. \quad (3.3)$$

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 a 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**.

⁵We have followed [16, Definition 3.2] by restricting to *efficiently* computable numbers in the definition of QTMs. This is not strictly necessary in our case, as we are concerned with computability rather than complexity theory here. [16] show that this restriction is without loss of generality. In fact, in our case, the transition amplitudes of all the QTMs we construct will belong to the even more restrictive set of algebraic numbers.

Definition 13 (Generalised TM and QTM). A generalised TM or generalised QTM is defined exactly as a standard TM (Definition 11) or standard QTM (Definition 12) except that the set of head movement directions is $\{L, R, N\}$ instead of just $\{L, R\}$.

Following Bernstein and Vazirani [16], we define various classes of deterministic and quantum Turing Machines:

Definition 14 (Well-formed – Definition 3.3 in [16]). 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 [16]) 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 15 (Normal form – Definition 3.13 in [16]). A well-formed QTM or reversible TM $M = (\Sigma, Q, \delta)$ is in *normal form* if

$$\forall \sigma \in \Sigma \quad \delta(q_f, \sigma) = |\sigma\rangle |q_0\rangle |N\rangle. \tag{3.4}$$

In [16] 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 22, below).

There is a specific class of QTMs, called *unidirectional*, that play an important role in the general theory developed by Bernstein and Vazirani [16].

Definition 16 (Unidirectional – Definition 3.14 in [16]). 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 17 (Local well-formedness – Thm. B.1 and Cor. B.3 in [16]). 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 [16] proved that one can w.l.o.g. restrict to *unidirectional* QTMs. We therefore restrict the following quantum analogue of Theorem 17 to the unidirectional case:

Theorem 18 (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

$$\forall p, \sigma \in Q \times \Sigma \quad \|\delta(p, \sigma)\| = 1. \tag{3.5a}$$

Orthogonality

$$\forall (p_1, \sigma_1) \neq (p_2, \sigma_2) \in Q \times \Sigma \quad \langle \delta_r(p_1, \sigma_1), \delta_r(p_2, \sigma_2) \rangle = 0. \quad (3.5b)$$

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

Bernstein and Vazirani [16] 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 [16].) We therefore extend the proof of Theorem 18 to generalised QTMs, which is not difficult – indeed, it is made particularly straightforward by our restriction to unidirectional machines.

Proof of Theorem 18. 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 19 (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 cell⁶, 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 movements apply in superposition.

Behaving properly is not a severe restriction on classical Turing Machines.⁷ 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

⁶This property is called *stationarity* in [16]

⁷Nor is it a severe restriction on QTMs, but we will not need or prove this here.

(see any standard text book on the theory of computation, e.g. Kozen [50]), 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 [16] 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.⁸ To avoid technical issues, we follow Bernstein and Vazirani [16] in defining QTMs on two-way infinite tapes, but we will ensure that none of the QTMs (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.

3.1.1. 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 $[\sigma_1, \sigma_2, \dots, \sigma_k] \in \Sigma_1 \times \Sigma_2 \times \dots \times \Sigma_k$ specifying the symbol written on each track. Similarly, the configuration of the tape will be specified by a tuple $[c_1, c_2, \dots, c_n] \in \Sigma_1^* \times \Sigma_2^* \times \dots \times \Sigma_k^*$, 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 17 or Theorem 18, 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 [16], 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 20 (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; when started on \mathcal{X}_2 , M_2 leaves the first track untouched.

⁸Effectively, we can only encode quantum bounded linear automata rather than QTMs, though there is no limit on the finite tape length we can encode.

- (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}_2 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 [16]. □

Lemma 21 (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 [16]. □

Lemma 22 (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 [16]. 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 $\tau \in \Sigma$,

$$\delta'(q, \tau) = \sum_{p,q} \delta(p, \sigma, \tau, q, d_q)^* |\sigma\rangle |p\rangle |\bar{d}_p\rangle. \tag{3.6}$$

- (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 18 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 (ii). 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 (ii) 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 (iv) 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 (ii) implies that M^\dagger executes the sequence of n steps $\pi(|\phi_{n-1}\rangle), \dots, \pi(|\phi_0\rangle)$. Finally, by (iv), 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 15 as opposed to that in [16]) 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 [16]).

In the following sections, we give explicit constructions of the reversible and quantum Turing Machines that, when combined appropriately using Lemmas 20 to 22 as described below, implement a phase-estimation QTM with the properties required to prove Theorem 10. The formal definitions of all these Turing Machines are given by their transition rule tables, which can directly be verified to satisfy the well-formedness conditions of Theorems 17 and 18.

Since understanding how a TM works just from a table of transition rules is not always straightforward, we additionally give an informal pseudocode description of how each machine operates on suitable input. These pseudocode descriptions also help in verifying that the machines specified in the transition rule tables behave properly and satisfy the claimed space and run-time bounds.

Each pseudocode listing describes how the Turing Machine moves its head and reads and writes tape symbols as the algorithm proceeds, in order to implement the claimed operation. The internal state that the TM transitions into *after* carrying out the corresponding line of pseudocode is indicated in the margin.⁹ The internal state that the TM is in at the beginning of an algorithm or subroutine is also indicated in the margin. Where no internal state is indicated in the margin, this means the internal state is left unchanged. Conditional branches in the pseudocode (i.e. **if**, **else if** and **else** statements) are structured to reflect how the different computational paths that can be followed by the TM during the computation branch off and are then re-merged to ensure reversibility.

However, it should be emphasised that these pseudocode listings are not rigorous specifications of the corresponding Turing Machines; the formal, mathematically rigorous specifications are the transition rule tables.

3.1.2. 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 10, we will be interested in keeping tight control of the space requirements of all our constructions.

Lemma 23 (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. In pseudocode:

⁹Generally, the new internal state is indicated on the line that carries out the head movement associated with the transition rule that transitions into that new state, except for transition rules where the TM head stays put.

Algorithm 1 COPY

```

1: input: string  $s \in (\Sigma \setminus \#)^*$  on track 1, blank track 2 ▷  $q_0$ 
2: if read blank symbol # on both tracks then
3:   halt ▷  $q_f$ 
4: else
5:   step Right ▷  $q_1$ 
6:   while read non-blank symbol  $\sigma$  on track 1, blank symbol # on track 2 do
7:     copy symbol  $\sigma$  from track 1 to track 2
8:     step Right
9:   end while
10:  step Left ▷  $q_2$ 
11:  while read same non-blank symbol  $\tau$  on both tracks do
12:    step Left
13:  end while
14:  copy symbol  $\sigma$  from track 1 to track 2
15:  halt ▷  $q_f$ 
16: end if

```

It is straightforward to verify that the following normal form transition function implements COPY:

| | $[\sigma, \#]$ | $[\#, \#]$ | $[\tau, \tau]$ | |
|-------|------------------------------------------------|----------------------|--------------------------|--|
| 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 | $([\sigma, \#], q_0, N)$ | $([\#, \#], q_0, N)$ | $([\tau, \tau], q_0, N)$ | |
| | $\forall \sigma, \tau \in \Sigma \setminus \#$ | | | |

(3.7)

This partial transition function verifies the two conditions of Theorem 17, so it is well-formed and can be completed to give a reversible TM. □

Lemma 24 (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. The Turing Machine has a separate internal state q^σ corresponding to each tape symbol σ . It reads the symbol into this internal state, overwrites it with the preceding symbol, and steps right. In pseudocode:

Algorithm 2 SHIFT

```

1: input: string  $s \in (\Sigma \setminus \#)^*$  ▷  $q_0$ 
2: if read blank symbol # then
3:   step Right ▷  $q_2$ 
4: else
5:   read non-blank symbol  $\sigma$ 
6:   write #
7:    $v \leftarrow \sigma$ 
8:   step Right ▷  $q_1^\sigma$ 
9:   while read non-blank symbol  $\sigma$  do
10:    write  $v$ 
11:     $v \leftarrow \sigma$ 
12:    step Right ▷  $q_1^\sigma$ 
13:   end while
14:   write  $v$ 
15:   step Right ▷  $q_2$ 
16: end if
17: step Left ▷  $q_3$ 
18: while read non-blank symbol  $\sigma$  do
19:   step Left
20: end while
21: halt ▷  $q_f$ 

```

The reason for including two different states q_2 and q_3 is to guarantee the unidirectionality requirement of Theorem 17. The state q_2 is entered from the right and the state q_3 from the left. The effect of this is that, if the input is the empty string, this Turing Machine program still steps right and then left before halting.

It is straightforward to verify that the following normal-form transition function implements SHIFT:

| | | | |
|---------|----------------|-----------------------|-------|
| | # | σ | |
| q_0 | $(\#, q_2, R)$ | $(\#, q_1^\sigma, R)$ | |
| q_1^v | (v, q_2, R) | (v, q_1^σ, R) | |
| q_2 | $(\#, q_3, L)$ | | |
| q_3 | $(\#, q_f, N)$ | (σ, q_3, L) | |
| q_f | $(\#, q_0, N)$ | (σ, q_0, N) | (3.8) |

$\forall \sigma, v \in \Sigma \setminus \#$

As this partial transition function verifies the two conditions of Theorem 17, it is well-formed and can be completed to give a reversible TM. □

Lemma 25 (Equality machine). *There is a three-track, normal form, reversible TM E_{QL} 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, E_{QL} runs for time $4|s| + 1$ and uses $|s| + 1$ space.*

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 reach the end of the input without encountering a non-matching pair, return to the starting cell and flip the output bit. If we encounter 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.

To accomplish this, we use a state q_4 that is the only state that can transition to q_f (except for the special case in which the first symbols in both input strings already differ, which is handled immediately). The state q_1 searches along the strings, until it either finds the end of both strings (first computational path), or it finds a point where the symbols in the two inputs differ (second computational path).

In the first path, state q_2 will return to the starting cell, switch the output bit and transition to q_3 . State q_3 will move right again doing the same as q_1 did and, at the end of both strings, will transition to q_4 . Note that we need two different states for this, due to unidirectionality: q_2 will be entered from the left and q_3 from the right.

In the second path, states q'_2 and q'_3 play the corresponding roles of q_2 and q_3 , respectively. In both paths, q_4 will return to the starting cell and halt. Within each path there is no issue with reversibility. The key to this is that, when both computational paths are merged into state q_4 , they do so on exactly the same input states as when the paths diverged (with q_1 replaced by q_3 or q'_3 depending on the path). Thus these transitions fulfil the properties required for a well-formed (partial) transition function by Theorem 17.

This implementation requires being able to identify the starting cell, so that we can return to it again. This is also important since we want to avoid ever moving the head before the starting cell, to ensure the TM is proper, see Definition 19.) 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 in this case, by copying over the symbol from the first track. (This is not the only way one could handle this.)

Proof of Lemma 25. In pseudocode, the Turing Machine program for EQL does the following:

Algorithm 3 EQL

```

1: input: string  $s \in (\Sigma \setminus \#)^*$  on track 1
2:   string  $t \in (\Sigma \setminus \#)^*$  on track 2
3:   bit  $b \in \{0, 1\}$  on track 3                                     ▷  $q_0$ 
4: if read symbols  $\sigma \neq \tau$  on tracks 1 and 2 then
5:   halt                                                         ▷  $q_f$ 
6: else
7:   write # on track 2
8:   step Right                                                  ▷  $q_1$ 
9:   while read non-blank symbols  $\sigma = \tau$  on tracks 1 and 2 do
10:    step Right
11:   end while
12:   if read blank symbols # on both tracks 1 and 2 then
13:    step Left                                                  ▷  $q_2$ 
14:    while read non-blank symbols  $\sigma = \tau$  on tracks 1 and 2 do
15:     step Left
16:    end while
17:    flip bit  $b$  on track 3
18:    step Right                                                  ▷  $q_3$ 
19:    while read symbols  $\sigma = \tau$  on tracks 1 and 2 do
20:     step Right
21:    end while
22:   else if read different symbols  $\sigma \neq \tau$  on tracks 1 and 2 then
23:    step Left                                                  ▷  $q'_2$ 
24:    while read non-blank symbols  $\sigma = \tau$  on tracks 1 and 2 do
25:     step Left
26:    end while
27:    step Right                                                  ▷  $q'_3$ 
28:    while read symbols  $\sigma = \tau$  on tracks 1 and 2 do
29:     step Right
30:    end while
31:   end if
32:   step Left                                                  ▷  $q_4$ 
33:   while read symbols  $\sigma = \tau$  on tracks 1 and 2 do
34:    step Left
35:   end while
36:   copy track 1 symbol  $\sigma$  to track 2
37:   halt                                                         ▷  $q_f$ 
38: end if

```

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 | $([\sigma, \sigma, \#], q_0, N)$ | $([\sigma, \tau, \#], q_0, N)$ | $([\sigma, \#, \#], q_0, N)$ | $([\#, \tau, \#], q_0, N)$ | |

(3.9)

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

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

One can verify that this partial transition function satisfies the two conditions of Theorem 17, so it is well-formed and can be completed to give a reversible TM, as required. □

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 26 (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 use 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 the general procedure for reversible

simulation and uncomputation of a non-reversible TM due to Bennett [13], 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 26. We first use the `SHIFT` machine from Lemma 24 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`[†] (the reversal of the `SHIFT` machine, constructed using Lemma 22) to shift the output back one cell to the left.

To implement reversibly the incrementing part of this procedure, we assign states as follows. State q_1 will search along the input for the first 0 to change it to a 1, and along the way will change all 1s to 0s. In this way, the carry from incrementing the first (least significant) bit is propagated along the binary string to the appropriate bit. If no such 0 exists, a 1 will be appended to the end of the input string by changing the blank symbol `#` there to a 1.

These two possible computation paths have to be merged back into a state q_4 , which will then return the TM head to the starting cell and halt. To merge these two paths, which end up in states q_2 and q'_2 respectively, we add an extra state q_3 together with transition rules from q_2 and q'_2 into q_3 , such that both computational paths arrive at a configuration in which the TM is in state q_3 and reading a 1. Thereupon, the TM can transition into q_4 and complete the computation.

In pseudocode:

Algorithm 4 `INC`

| | |
|-------------------------------------------------------------------------------------------------------------|----------|
| 1: input: blank <code>#</code> followed by little-endian binary representation of $n \in \mathbb{N}$ | ▷ q_0 |
| 2: step Right | ▷ q_1 |
| 3: while read 1 do | |
| 4: write a 0 | |
| 5: step Right | |
| 6: end while | |
| 7: if read 0 then | |
| 8: write a 1 | |
| 9: step Right | ▷ q_2 |
| 10: step Left | ▷ q_3 |
| 11: else if read blank symbol <code>#</code> then | |
| 12: write a 1 | |
| 13: step Right | ▷ q'_2 |
| 14: step Left | ▷ q_3 |
| 15: end if | |
| 16: step Left | ▷ q_4 |
| 17: while read a 0 do | |
| 18: step Left | |
| 19: end while | |
| 20: halt | ▷ q_f |

The following normal-form transition function implements this:

| | # | 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 | $(\#, q_3, L)$ | | | |
| q_3 | | | $(1, q_4, L)$ | |
| q_4 | $(\#, q_f, N)$ | $(0, q_4, L)$ | | |
| q_f | $(\#, q_0, N)$ | $(0, q_0, N)$ | $(1, q_0, N)$ | (3.10) |

This partial transition function verifies the two conditions of Theorem 17, 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 22. □

We can use the INC construction to construct a looping primitive. (This gives a slightly more general version of the Looping Lemma from [16, Lemma 4.13], which also has tighter control on the space requirements.)

Lemma 27 (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 starting 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 [16, 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 the 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 $b' = -b$ if $x = n$ or $x = m - 1$ but not both, otherwise $b' = b$. M_1 can be constructed by dovetailing together the EQL machine from Lemma 25 (with the first and third tracks as its input tracks and the fourth track as its output), then the INC machine from Lemma 26 (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 21 and the fact that all the constituent machines are proper, normal form and reversible, M_1 is proper, normal form and reversible. From Lemmas 25 and 26, M_1 runs for time $O(\log m)$ and takes at most $|m| + 2$ space.

M_2 has new initial and final states q_α, q_ω , with q_0 and q_f as its only other normal (i.e. not 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 following normal form, partial transition function for M_2 is essentially the same as the corresponding construction in Lemma 4.2.6 of Bernstein and Vazirani [16], 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 requirements (i) to (iii), above:

| | | | | |
|------------|------------------------|------------------------|------------------------|--------|
| | # | 0 | 1 | |
| q_α | (1, q_0 , N) | | | (3.11) |
| q_f | (0, q_0 , N) | | (0, q_ω , N) | |
| q_ω | (#, q_α , N) | (0, q_α , N) | (1, q_α , N) | |

M_2 is clearly normal form, and satisfies the two conditions of Theorem 17.

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:

$$\begin{aligned}
 n; m; n; 0 &\xrightarrow{(i)} n; m; n; 1 \xrightarrow{M_1} n; m; n+1; 0 \xrightarrow{(ii)} n; m; n+1; 0 \xrightarrow{M_1} n; m; n+2; 0 \xrightarrow{(ii)} \dots \\
 &\dots \xrightarrow{(ii)} n; m; m-1; 0 \xrightarrow{M_1} n; m; m; 1 \xrightarrow{(iii)} n; m; m; 0
 \end{aligned} \tag{3.12}$$

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 23 (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[†] of the copying machine from Lemma 23 (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 28, 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 show that combining M_1 and M_2 as described above to form M' gives a proper, normal form, reversible TM. 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 combining it with M_2 to construct 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 17. 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 additional 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. □

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 28 (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 27 on the second track, inserting for its special state the INC machine of Lemma 26 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 22. □

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

Lemma 29 (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 $1^n; \#$ (n written in unary on the first track), $UTOB$ behaves properly and outputs $1^n; n$ (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 26 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 | $(\#, q_0, N)$ | $(1, q_0, N)$ |

(3.13)

This verifies the two conditions of Theorem 17, 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 M 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.

In pseudocode:

Algorithm 5 M

```

1: input: unary representation of  $n \in \mathbb{N}$ , i.e. string of  $n$  1's ▷  $q_0$ 
2: if read blank symbol # then
3:   halt ▷  $q_f$ 
4: else
5:   call  $q$  ▷  $q$ 
6:   write a #
7:   step Right ▷  $q_1$ 
8:   call PLACE_MARKER_AND_RESET
9:   loop
10:    call  $q$  ▷  $q$ 
11:    step Right ▷  $q'_2$ 
12:    while read a 1 do
13:      step Right
14:    end while
15:    write a 1
16:    step Right ▷  $q_1$ 
17:    call PLACE_MARKER_AND_RESET
18:  end loop
19: end if
20:
21: procedure PLACE_MARKER_AND_RESET ▷  $q_1$ 
22:   if read a # then
23:     step Left ▷  $q_4$ 
24:     call RESET_AND_HALT
25:   else if read a 1 then
26:     write a #
27:     Step Left ▷  $q_2$ 
28:     while read a 1 do
29:       step Left
30:     end while
31:   end if
32: end procedure
33:
34: procedure RESET_AND_HALT ▷  $q_4$ 
35:   while read a 1 do
36:     step Left
37:   end while
38:   write a 1
39:   halt ▷  $q_f$ 
40: end procedure

```

The purpose of this apparently pointless procedure is that, by Lemma 20, we can substitute the INC machine from Lemma 26 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

3.2. 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 fraction 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:

- (i) 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 $|1\rangle$);
- (ii) The controlled-unitary stage, during which control- U_φ operations are applied between the output qubits and the ancilla qubit;
- (iii) A stage in which we locate the least significant bit of the output;
- (iv) The quantum Fourier transform stage, in which the inverse quantum Fourier transform is applied to the output qubits;
- (v) 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 21 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 1^N (the unitary representation of N) 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.

3.3. 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 $1^N; \#$ of the quantum and input tracks into $1^N; N$. Running the unary-to-binary converter U_{ToB} from Lemma 29 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 to $|\#\rangle$ 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.)

In pseudocode:

Algorithm 6 Preparation QTM

| | |
|--------------------------------------------------------------------------------------|----------------------|
| 1: input: unary representation of $N \in \mathbb{N}$, i.e. string of N 1's | $\triangleright q_0$ |
| 2: apply unitary gate that rotates from $ 1\rangle$ to $\#$ | |
| 3: step Right | $\triangleright q_1$ |
| 4: while read $ 1\rangle$ do | |
| 5: apply Hadamard gate to rotate to $ +\rangle$ | |
| 6: step Right | |
| 7: end while | |
| 8: apply unitary gate that rotates from $\#$ to $ +\rangle$ | |
| 9: step Right | $\triangleright q_2$ |
| 10: step Left | $\triangleright q_3$ |
| 11: while not read $\#$ do | |
| 12: Step Left | |
| 13: end while | |
| 14: apply unitary gate that rotates from $\#$ to $ 1\rangle$ | |
| 15: halt | $\triangleright q_f$ |

It is straightforward to verify that the following partial transition function satisfies the conditions of Theorem 18, 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 | $\frac{1}{\sqrt{2}}(0\rangle + 1\rangle) q_2\rangle R\rangle$ | | $\frac{1}{\sqrt{2}}(0\rangle + 1\rangle) q_1\rangle R\rangle$ | (3.14) |
| q_2 | $\# \rangle q_3\rangle L\rangle$ | | | |
| q_3 | $ 1\rangle q_f\rangle N\rangle$ | $ 0\rangle q_3\rangle L\rangle$ | $ 1\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$ | |

3.4. 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 2). Constructing a QTM that implements this is more complex. The basic idea is to supply the QTM with an internal state q that 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 2. The preparation and control- U_φ stages of the quantum phase estimation circuit for φ (cf. Fig. 5.2 in [62]).

We will use three auxiliary tracks. A “mark track”, 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 27.

3.4.1. 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 30 (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 24 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^\dagger machine at the very end (where SHIFT^\dagger is the reversal of M_1 constructed using Lemma 22) 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 effectively applies a CNOT gate between the qubits on the quantum track marked by c and t on the mark track, with c as the control and t as the target. However, as these qubits are not necessarily adjacent on the tape, M_1 must temporarily make use of an auxiliary qubit stored in the internal states of the QTM, which gets “carried” along with the QTM head. As we cannot leave a qubit in the internal state at the end of the computation, however, M_1 must reversibly erase this internal qubit after it has served its purpose, by applying a second CNOT between it and the control qubit. (See below for further details and quantum circuit diagrams.)

M_1 implements this in stages: first, it scans right until it encounters a c on the mark track, at which point it applies a CNOT between the quantum track and a qubit stored in its internal state. It then returns the head to the starting cell, and applies a CNOT between this internal qubit and an auxiliary qubit m on the mark track (defined by the two orthogonal states $|m_0\rangle$ and $|m_1\rangle$, which are identified with $|0\rangle$ and $|1\rangle$ respectively). Finally, it returns to the location of the c on the mark track, applies a second CNOT between the quantum track and the internal qubit, returns to the starting cell, and halts.

In pseudocode:

Algorithm 7 M_1

```

1: input: mark track containing one  $c$ , one  $t$ , and # everywhere else;
2:     quantum track containing an arbitrary quantum state.                                ▷  $q_0$ 
3: write  $m_0$  to mark track
4: step Right                                                                           ▷  $q_1$ 
5: while not read  $c$  on mark track do
6:     step Right
7: end while
8: apply CNOT gate with quantum track as control, internal qubit  $j$  as target
9: step Left                                                                           ▷  $q_2^j$ 
10: while read # or  $t$  on mark track do
11:     step Left
12: end while
13: apply CNOT with internal qubit  $j$  as control, mark track qubit  $m$  as target          ▷  $q_3^j$ 
14: while not read  $c$  on mark track do
15:     step Right
16: end while
17: apply CNOT gate with quantum track as control, internal qubit  $j$  as target
18: step Left                                                                           ▷  $q_4$ 
19: while read # or  $t$  on mark track do
20:     step Left
21: end while
22: halt                                                                                ▷  $q_f$ 

```

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 two of the internal states of M_1 comes in two varieties, indicated by a superscript $j = \{0, 1\}$ and identified with $|0\rangle, |1\rangle$, respectively, which are used to temporarily store a qubit in the internal state of the machine:

| | [#, #] | [#, i] | [t , ·] | [c , j] | [m_k , ·] | |
|---------|--------------------------------------------|--------------------------------------------|-----------------------------------------------|-------------------------------------------|-----------------------------------------------------|--------|
| 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_3^j \rangle R \rangle$ | (3.15) |
| q_3^j | | $ \#, i \rangle q_3^j \rangle R \rangle$ | $ t, \cdot \rangle q_3^j \rangle R \rangle$ | $ c, j \rangle q_4 \rangle L \rangle$ | | |
| q_4 | | $ \#, i \rangle q_4 \rangle L \rangle$ | $ t, \cdot \rangle q_4 \rangle L \rangle$ | | $ m_k \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$ | |

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

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 27, 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 starting 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.

In pseudocode:

Algorithm 8 M'

- 1: **input:** mark track containing auxiliary qubit m in the first cell,
- 2: one c , one t and $\#$ everywhere else. ▷ q_0
- 3: apply CNOT gate with mark track qubit m as control, internal qubit j as target
- 4: step Right ▷ q_1^j
- 5: **while** not read t on mark track **do**
- 6: step Right
- 7: **end while**
- 8: apply cU gate with internal qubit j as control, quantum track as target
- 9: step Left ▷ q_2^j
- 10: **while** read $\#$ or c on mark track **do**
- 11: step Left
- 12: **end while**
- 13: apply CNOT gate with mark track qubit m as control, internal qubit j as target
- 14: halt ▷ q_f


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 | | | $[\#, \cdot] q_1^0\rangle R\rangle$ | $[t, j] q_2^0\rangle L\rangle$ | $[c, \cdot] q_1^0\rangle R\rangle$ | |
| q_1^1 | | | $[\#, \cdot] q_1^1\rangle R\rangle$ | $\sum_i u_{ij} [t, i] q_2^1\rangle L\rangle$ | $[c, \cdot] q_1^1\rangle R\rangle$ | |
| q_2^0 | $ m_0, \cdot\rangle q_f\rangle N\rangle$ | | $[\#, \cdot] q_2^0\rangle L\rangle$ | | $[c, \cdot] q_2^0\rangle L\rangle$ | (3.16) |
| 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$ | $[\#, \cdot] q_0\rangle N\rangle$ | $[t, j] q_0\rangle N\rangle$ | $[c, \cdot] q_0\rangle N\rangle$ | |

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

M' is proper, never alters the mark track and, for given mark track configuration, always runs for the same number of time steps. So substituting M' in the looping machine Loop of Lemma 27 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 (which we label m), and the two internal qubits stored in the internal states q_2^j of M_1 and q_1^j of M_2 (let's label these $j_{1,2}$, respectively). Qubits m, j_1 and j_2 are all 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 j_1 qubits, followed by a CNOT between the j_1 and m qubits, and finally another CNOT between the c and j_1 qubits:



Similarly, M' applies a CNOT between the k and j_2 qubits, a control- U operation between the j_2 and t qubits, and a final CNOT between the k and j_2 qubits. Letting cU denote the control- U gate, and recalling that the j_2 qubit starts off in the $|0\rangle$ state, the overall effect of this is:

$$CNOT_{m,j_2} cU_{j_2,t} CNOT_{m,j_2} |\psi\rangle_{m,t} |0\rangle_{j_2} = cU |\psi\rangle_{m,t} |0\rangle_{j_2}, \tag{3.17}$$


or, as a quantum circuit diagram:



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:

$$(CNOT_{m,j_2} cU_{j_2,t} CNOT_{m,j_2})^k |\psi\rangle_{m,t} |0\rangle_{j_2} = cU^k |\psi\rangle_{m,t} |0\rangle_{j_2}, \tag{3.18}$$


or, in quantum circuit diagrams:



Since the CNOT gate is self-inverse, the time-reversal M_1^\dagger in fact applies the same sequence of CNOTs as M_1 . So, on the control, target, mark and internal qubits, dovetailing M_1, M_2, M_1^\dagger carries out the operation:

$$\begin{aligned} &(CNOT_{c,j_1} CNOT_{j_1,m} CNOT_{c,j_1} cU_{m,t}^k \cdot \\ &CNOT_{j_1,m} CNOT_{c,j_1} CNOT_{c,j_1}) |\psi\rangle_{c,t} |0\rangle_{j_1} |0\rangle_m \\ &= cU_{c,t}^k |\psi\rangle_{c,t} |0\rangle_{j_1} |0\rangle_m, \end{aligned} \tag{3.19}$$

or, in quantum circuit diagrams:



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

3.4.2. 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 output qubits are stored in reverse order on the quantum track starting from the 2nd cell, 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; recall that the 1st cell holds the target qubit.) 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 30), 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.

In pseudocode, $STEP_c$ does the following:

Algorithm 9 $STEP_c$

- 1: **input:** string $s \in \{t, c, \#\}^*$ starting with the symbol t , followed by zero or
 - 2: more #'s, ending with a single c symbol. ▷ q_0
 - 3: step Right ▷ q_1
 - 4: **while** read # **do**
 - 5: step Right
 - 6: **end while**
 - 7: write a #
 - 8: step Left ▷ q_2
 - 9: write a c
 - 10: step Left ▷ q_3
 - 11: **while** not read a t **do**
 - 12: step Left
 - 13: **end while**
 - 14: halt ▷ q_f
-

The following well-formed, normal form, partial transition function implements the reversible TM $STEP_c$ (which acts only on the mark track):

| | | | | |
|-------|----------------|---------------|----------------|--|
| | # | t | c | |
| q_0 | | (t, q_1, R) | | |
| q_1 | $(\#, q_1, R)$ | | $(\#, q_2, L)$ | |
| q_2 | (c, q_3, L) | | | |
| q_3 | $(\#, q_3, L)$ | (t, q_f, N) | | |
| q_f | $(\#, q_0, N)$ | (t, q_0, N) | (c, q_0, N) | |

(3.20)

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

Doubling a number in binary simply appends a 0 onto the binary representation of the number. The Turing Machine implementation of this is particularly convenient in the little-endian convention we are using, as it simply means finding the first blank symbol and appending the 0 there. We can simplify the implementation slightly and avoid the Turing Machine head ever moving before the starting cell, by taking advantage of the fact that the input in our case is always a power of 2, so, in little-endian order, always consists of a string of zero or more 0’s followed by a single 1; in particular, the input is never blank. Thus we can overwrite the initial 0 or 1 with a blank symbol to mark the start of the tape, and restore the initial tape cell to 0 at the end of the algorithm. In pseudocode:

Algorithm 10 DBL

```

1: input: binary representation of  $k = 2^n$  for  $n \in \mathbb{N}$ . ▷  $q_0$ 
2: if read a 1 then
3:   write a #
4:   step Right ▷  $q_2$ 
5: else
6:   write a #
7:   step Right ▷  $q_1$ 
8:   while read a 0 do
9:     step Right
10:  end while
11:  write a 0
12:  step Right ▷  $q_2$ 
13: end if
14: write a 1
15: step Left ▷  $q_3$ 
16: while read 0 do
17:   step Left
18: end while
19: write a 0
20: halt ▷  $q_f$ 

```

The following well-formed, normal form, partial transition function for DBL (which acts only on the cU looping track) does precisely this:

$$\begin{array}{c|ccc}
 & \# & 0 & 1 \\
 \hline
 q_0 & & (\#, q_1, R) & (\#, q_2, R) \\
 q_1 & & (0, q_1, R) & (0, q_2, R) \\
 q_2 & (1, q_3, L) & & \\
 q_3 & (0, q_f, N) & (0, q_3, L) & \\
 q_f & (\#, q_0, N) & (0, q_0, N) & (1, q_0, N)
 \end{array} \tag{3.21}$$

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 $[t, 1]$, then returns to the starting cell and halts. (Constructing a reversible TM for this is trivial.) We dovetail this with a proper, normal form reversible TM that scans to the end of the output qubits, initialises the mark track in that cell to a c , then returns to the starting cell and halts. In pseudocode:

Algorithm 11 controlled- U_φ initialisation

- 1: **input:** mark track containing t in the first cell;
- 2: quantum track with ancilla and output qubits followed by #s ▷ q_0
- 3: step Right ▷ q_1
- 4: **while** not read # on quantum track **do**
- 5: step Right
- 6: **end while**
- 7: step Left ▷ q_2
- 8: write c on the mark track
- 9: step Left ▷ q_3
- 10: **while** not read t on mark track **do**
- 11: step Left
- 12: **end while**
- 13: halt ▷ q_f

The following normal form, partial transition function (which acts only on the mark and quantum tracks) accomplishes this:

$$\begin{array}{c|ccc}
 & [\#, \#] & [\#, i] & [t, j] \\
 \hline
 q_0 & & & ([t, j], q_1, R) \\
 q_1 & ([\#, \#], q_2, L) & ([\#, i], q_1, R) & \\
 q_2 & & ([c, i], q_3, L) & \\
 q_3 & & ([\#, i], q_3, L) & ([t, j], q_f, N) \\
 q_f & ([\#, \#], q_0, N) & ([\#, i], q_0, N) & ([t, j], q_0, N)
 \end{array} \tag{3.22}$$

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

On the mark track, this leaves a t and c over the ancilla qubit and the first output qubit,¹⁰ 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 23. We then run a reversible looping TM from Lemma 27 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 30 with the STEP_c and DBL machines constructed above.

On appropriate input, the cU^k machine from Lemma 30 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 30) 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 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

¹⁰Recall that the first output qubit is the last one on the quantum track.




Figure 3. The inverse QFT stage of the quantum phase estimation circuit (cf. Fig. 5.1 in [62]).

track, a t and c in the first and second cell 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 27). In the n 'th iteration it runs the cU^k machine from Lemma 30 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 $N; N; tc; N + 1$ 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 23 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 tc 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 3.3, the overall QTM we have constructed implements the control- U circuit of Figure 2. If φ_k denotes the k 'th digit in the binary fraction expansion of φ , then this prepares the state [62]

$$\frac{1}{2^N} \left(|0\rangle + e^{2\pi i 0.\varphi_N} |1\rangle \right) \left(|0\rangle + e^{2\pi i 0.\varphi_{N-1}\varphi_N} |1\rangle \right) \dots \left(|0\rangle + e^{2\pi i 0.\varphi_2 \dots \varphi_{N-1}\varphi_N} |1\rangle \right) \left(|0\rangle + e^{2\pi i 0.\varphi_1 \varphi_2 \dots \varphi_N} |1\rangle \right) \quad (3.23)$$

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.

3.5. 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 3).

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 3.4 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.¹¹

¹¹The universal QTM construction of Bernstein and Vazirani [16] can approximate the QFT on N qubits to arbitrary precision, but it does not implement it exactly.

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 10 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 fraction 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 [62]

$$\frac{1}{2^{\frac{N}{2}}} \left(|0\rangle + e^{2\pi i 0.\varphi_N} |1\rangle \right) \left(|0\rangle + e^{2\pi i 0.\varphi_{N-1}\varphi_N} |1\rangle \right) \dots \left(|0\rangle + e^{2\pi i 0.\varphi_2 \dots \varphi_{N-1}\varphi_N} |1\rangle \right) \left(|0\rangle + e^{2\pi i 0.\varphi_1 \varphi_2 \dots \varphi_N} |1\rangle \right) \tag{3.24}$$

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

The least significant bit of φ is the $|\varphi|$ 'th bit, and by assumption (cf. Theorem 10) $|\varphi| \leq N$. So the $|\varphi|$ 'th bit of φ is 1 and the $|\varphi| + 1 \dots N$ 'th bits are all 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.

Algorithm 12 LSB QTM

- 1: **input:** quantum tape containing $|1\rangle$ followed by N -qubit state from (3.24). ▷ q_0
 - 2: write $|\#\rangle$
 - 3: **repeat**
 - 4: step Right ▷ q_1
 - 5: apply a Hadamard gate ▷ q_2
 - 6: **until** read a $|1\rangle$
 - 7: step Left ▷ q_3
 - 8: **while** read $|0\rangle$ **do**
 - 9: step Left
 - 10: **end while**
 - 11: halt ▷ q_f
-

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 it obeys the three conditions of Theorem 18, 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$ | $\frac{1}{\sqrt{2}}(0\rangle + 1\rangle) q_2\rangle N\rangle$ | $\frac{1}{\sqrt{2}}(0\rangle - 1\rangle) q_2\rangle N\rangle$ | |
| 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$ |

(3.25)

(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:

$$\begin{aligned}
 &|#\rangle |0\rangle^{N-|\varphi|} |0\rangle |1\rangle \frac{1}{2^{\frac{|\varphi|-1}{2}}} \left(|0\rangle + e^{2\pi i 0 \cdot \varphi_{|\varphi|-1} \varphi_{|\varphi|}} |1\rangle \right) \left(|0\rangle + e^{2\pi i 0 \cdot \varphi_{|\varphi|-2} \varphi_{|\varphi|-1} \varphi_{|\varphi|}} |1\rangle \right) \dots \\
 &\dots \left(|0\rangle + e^{2\pi i 0 \cdot \varphi_1 \varphi_2 \dots \varphi_{|\varphi|-1} \varphi_{|\varphi|}} |1\rangle \right). \tag{3.26}
 \end{aligned}$$

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 29. 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 17 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 | $(\#, q, N)$ | $(0, q_2, L)$ | | |
| 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 | $(\#, q_0, N)$ | $(0, q_0, N)$ | $(1, q_0, N)$ | |

(3.27)

This machine effectively steps right over 0’s until it encounters a 1, 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 a special 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. In pseudocode:

Algorithm 13 LSB-counting TM

| | |
|-----------------------------------------------------------------------------------------|---------|
| 1: input: string $s \in \{0, 1, \#\}^*$ consisting of a single # followed by any | |
| 2: binary string | ▷ q_0 |
| 3: step Right | ▷ q_1 |
| 4: while read a 0 do | |
| 5: write a # | |
| 6: step Left | ▷ q_2 |
| 7: while read a 0 do | |
| 8: step Left | |
| 9: end while | |
| 10: enter special state q | ▷ q |
| 11: step Right | ▷ q_3 |
| 12: while not read a # do | |
| 13: step right | |
| 14: end while | |
| 15: write a 0 | |
| 16: step Right | ▷ q_1 |
| 17: end while | |
| 18: step Left | ▷ q_4 |
| 19: while read a 0 do | |
| 20: step Left | |
| 21: end while | |
| 22: halt | ▷ q_f |

By Lemma 20, we can substitute for the state q the INC machine from Lemma 26 (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 (3.26) 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 27, which uses the original count track as its input and work track, to run a SHIFT machine from Lemma 24 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 (3.26), 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.

3.6. 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 3](#)). 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.

In each cascade, we want to apply a $cU_{2^{-|\varphi|}}^{2^{|\varphi|+m-n}}$ 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 30](#). 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 26](#) to increment the number $N - |\varphi|$ stored on the count track to $N - |\varphi| + 1$. We then run the SUB machine from [Lemma 28](#), 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.

3.4.2. 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 27](#), 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^\dagger of the machine constructed in [Section 3.4.2](#).) We dovetail this with the $cU_{2^{-|\varphi|}}^k$ machine from [Lemma 30](#). 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^\dagger of the machine from [Section 3.4.2](#) and replacing c with t in its transition rules to give a STEP_t^\dagger 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 $cU_{2^{-|\varphi|}}^{2^{|\varphi|+m-n}}$ 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 3](#).

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.)

3.4.2. Outer loop

For the outer looping machine, we use a LOOP2 machine from [Lemma 27](#) running on the input and count tracks (which hold the numbers N and $N - |\varphi| + 1$, respectively). By [Lemma 27](#), 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.

3.7. 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 3.5). 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 3.5). We can then run the reversal SHIFT^\dagger (constructed using Lemma 22) of the machine used in Section 3.5, 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 3.5. (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 24 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 3.5. We can then run the reversal UTOB^\dagger of the unary-to-binary machine from Lemma 29, 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.

3.8. Analysis

By dovetailing together the preparation stage (Section 3.3), control- U_φ stage (Section 3.4), inverse-QFT stage (Section 3.6) and reset stage (Section 3.7), 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 10.

Now, we know [62] that the quantum phase estimation algorithm outputs the exact binary fraction 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 10.

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 10.

The last thing we must check in order to finish the proof of Theorem 10 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 18 the problem is equivalent to completing an orthonormal set of vectors with coefficients in

$$\mathcal{S} = \left\{ 0, 1, \pm \frac{1}{\sqrt{2}}, e^{i\pi\varphi}, e^{i\pi 2^{-|\varphi|}} \right\} \quad (3.28)$$

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 10.

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 fraction 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 6 to bound ground state energies of a Hamiltonian constructed out of this QTM.)

4. Encoding QTMs in local Hamiltonians

In Section 3, we obtained a QTM that implements exact quantum phase estimation deterministically, using a number of time-steps exponential in the bit-length of the phase being estimated. This gives us a way to generate any desired input to feed into a universal (reversible) Turing Machine (UTM). Dovetailing the phase estimation QTM and the UTM, we therefore obtain a family of QTMs of constant size (fixed alphabet size and number of internal states) for which the halting problem is undecidable on blank input tape.

In this section we will show how for any QTM, one can construct an associated Hamiltonian whose ground state encodes in a concrete way the evolution of the QTM. In particular, for the family of QTMs just described, it encodes the associated halting problem. We start by distilling the main ideas behind the construction from a historical perspective, to serve as a roadmap for this section.

The idea of encoding quantum computations into ground states of Hamiltonians goes back to Feynman [30], and was developed into its modern form by Kitaev [47]. The key object is the computational history state, which encodes the entire history of the computation in superposition. As these history states will play an important role in this section, we define them rigorously here:

Definition 31 (Computational history state). A *computational history state* $|\psi\rangle_{CQ} \in \mathcal{H}_C \otimes \mathcal{H}_Q$ is a state of the form

$$|\psi\rangle_{CQ} = \frac{1}{\sqrt{\dim \mathcal{H}_C}} \sum_{t=0}^{d_C} |t\rangle |\psi_t\rangle, \tag{4.1}$$

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 this quantum computation.

As Feynman realised, given a quantum circuit defined by the sequence of unitaries $\{U_t\}_{1 \leq t \leq T}$, it is straightforward to write down a Hamiltonian for which the ground space is precisely spanned by the set of computational history states (Definition 31), where one allows any initial state $|\psi_0\rangle_Q$.

In order to obtain such a Hamiltonian, one first focuses on the clock register \mathcal{H}_C . Let $d_C = T$, and look for a Hamiltonian that has as ground state the superposition of all clock time steps:

$$|\psi\rangle_C = \frac{1}{\sqrt{T}} \sum_{t=1}^T |t\rangle. \tag{4.2}$$

This can be enforced by a standard hopping Hamiltonian:




$$H_C = \sum_{t=1}^T (|t\rangle - |t-1\rangle) (\langle t| - \langle t-1|). \tag{4.3}$$


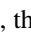

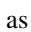
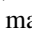
To obtain the history state of Definition 31 from (4.2), one applies the controlled unitary $\mathcal{U}_{CQ} = \sum_t |t\rangle\langle t| \otimes U_{[1\dots t]}$ on $|\psi\rangle_C |\psi_0\rangle_Q$, where $U_{[1\dots t]}$ is a shorthand notation for $U_t U_{t-1} \dots U_1$. At the level of the ground state, this boils down to considering the rotated Hamiltonian $H_{CQ} = \mathcal{U}_{CQ} (H_C \otimes \mathbb{1}_Q) \mathcal{U}_{CQ}^\dagger$, which is nothing more than

$$H_{CQ} = \sum_{t=1}^T \left(|t\rangle\langle t| \otimes \mathbb{1} + |t-1\rangle\langle t-1| \otimes \mathbb{1} - |t\rangle\langle t-1| \otimes U_t - |t-1\rangle\langle t| \otimes U_t^\dagger \right). \tag{4.4}$$

The difficulty arises when one wants to implement this construction with a Hamiltonian that is (a) local, (b) one-dimensional and (c) translationally invariant. Issues (a)–(c) were addressed consecutively by Kitaev et al. [47], Aharonov et al. [2] and Gottesman and Irani [32].

As we already illustrated, the key ingredient to construct the *history state Hamiltonian* (4.4) is the clock. An important insight from [2] was to define a 1D local Hamiltonian whose ground state encodes in superposition a series of sweeps of particular product states, which act as an oscillator of a clock. (Cf. Figure 4, which represents the particular clock oscillation we will use later).

Now, apart from hopping terms in (4.3) (called *evolution* or *transition rule* terms) enforcing the evolution of the clock, one now needs to also include *penalty terms* that restrict the type of configurations that can appear in the ground state, by giving a positive energy to configurations that do not appear in the clock oscillation cycle. In the example of Figure 4, for instance, giving a positive energy to configurations of the form  guarantees that no  can appear to the right of a  in any component of the ground state.

The type of constraints that can be enforced by penalty terms is characterised by the notion of a regular expressions, which for example (informally) allows one to express constraints of the form “write one , then as many ’s as you want, then write a  or a , then as many  as you want, and then




Figure 4. Example of a clock oscillation cycle.

a $\textcircled{>}$ ". (Note that this regular expression captures precisely the states present in the oscillation cycle of Figure 4.) The formal definition of regular expression is recalled in Definition 33, and the key result from [32] connecting regular expressions with penalty terms is stated in Lemma 36. An important point here is that we will be able to guarantee by other means that the sites at the two ends are always in states $\textcircled{<}$ and $\textcircled{>}$, respectively. (We will enforce this using tiling constructions, described in detail Section 5; the same condition is enforced in a different way in [32].) We can therefore restrict our analysis just to these “bracketed” configurations.

The idea of a local clock was developed further by Gottesman and Irani [32] to obtain a translationally invariant clock where, analogous to a real clock, after each oscillation cycle (representing one “tick” of the clock), a counter gets incremented by one. In [32] the clock counter counts in unary, and is represented on a separate track of the TM. For a finite chain of length L , this means the clock can tick for $O(L)$ steps before the counter cannot be incremented further and the clock stops ticking. In our case, we need the clock to tick for an exponentially larger number of time steps. Hence we need to count in binary rather than unary. In fact, it will be convenient to have it count in an even larger (but fixed) base ζ . To achieve this, we will have to generalise the binary counter machine constructed in the previous section. This is done in Section 4.4.1.

There is an important subtlety. When one takes into account the full clock (oscillation and counter), penalty terms and transition rule terms are not enough to characterise the desired clock evolution. An example of this will be given below, immediately after the statement of Lemma 37. The key idea used to solve this problem in [2], also used crucially later in [32], was a more general way to guarantee for those particular constructions that all other product state configurations (except for those giving the desired clock evolution) have larger energy. Note that if a state does not follow the correct evolution, it picks up a positive energy contribution from the transition rule terms. Some of the undesired configurations pick up a positive energy contribution directly from the penalty terms, as before. There are other undesirable configurations that are not detected directly by the penalty terms. But these remaining undesirable configurations all evolve under the transition rule terms into configurations that *do* pick up such an energy penalty. Aharonov et al. [2] call this the *Clairvoyance Lemma*. This result may or may not hold for any given construction; the clock and Hamiltonian construction must be designed such that this property holds and the Clairvoyance Lemma proven for the specific construction at hand.

Aharonov et al. [2] proved their Clairvoyance Lemma by dividing the Hilbert space of the clock into subsets of product states that are invariant under the clock evolution, i.e. into orbits of the action of the clock Hamiltonian. If the above property holds – that undesired states get penalties or evolve to states with penalties – all orbits pick up energy penalties except the one corresponding to a valid clock evolution. This in turn ensures that the ground states of the full Hamiltonian are precisely the set of history states (Definition 31). Moreover, the analysis gives explicit lower bounds for the energies of the excited states.

We will prove a version of the Clairvoyance Lemma for our construction following the same approach. We first introduce penalty terms and transition rule terms for the clock Hamiltonian. We prove (building on the key result Lemma 37) that all product states that do not appear in the desired clock evolution either pick up an energy penalty from the penalty terms, or evolve to a configuration that picks up such a penalty. This allows us to prove that the clock Hamiltonian consisting of the sum over these penalty and transition rule terms has as its unique ground state the uniform superposition over the product states appearing in the desired clock evolution. This is the key ingredient in our version of the Clairvoyance Lemma (Section 4.7), which shows that for the complete Hamiltonian that includes both the clock and computational registers C and Q , the unique ground state is the computational history state of Definition 31, and all other eigenstates have energies lower bounded by $\Omega\left(\frac{1}{T^3}\right)$.

The final insight to solve the issue (c) related to the translationally invariance property comes from Gottesman and Irani [32]. There, it was shown how all the ideas of [2] to address the local 1D case, could be implemented keeping the Hamiltonian translationally invariant. Indeed, this section can be seen as a self-contained analysis of a local Hamiltonian construction that generalises that of Gottesman and Irani [32] for the particular purposes of this paper.

Apart from adapting the clock construction to a translationally invariant setting, one of the main changes with respect to [2] (or to the original Feynman-Kitaev construction outlined above) in order to make the construction translationally invariant was to switch from quantum circuits to QTMs. That is, after each clock tick, a step of a QTM evolution is executed on a different track. In this case, the computational register Q in Definition 31 corresponds to the whole QTM configuration (internal state, head location, and tape configuration), encoded across a number of separate tracks. However, in order to apply our version of the Clairvoyance Lemma, one needs to guarantee that the U_t terms appearing in (4.4) are indeed unitaries, and not partial isometries, even in the case in which the evolution corresponds to a QTM rather than a quantum circuit. In our case, this is guaranteed by restricting to a special class of QTMs, first identified in Bernstein and Vazirani [16], that has especially nice properties (well-formed Definition 38, normal-form Definition 15 and unidirectional Definition 16, to which we also add proper Definition 19) but that, at the same time, is general enough for our purposes.¹² We give a more detailed overview of how these properties, together with a modification to the encoded computation, allow us to prove our version of the key Clairvoyance Lemma, in Section 4.7.

The following is the main result of this section:

Theorem 32 (Local Hamiltonian QTM encoding). *Let $\mathbb{C}^d = \mathbb{C}^C \otimes \mathbb{C}^Q$ be the local Hilbert space of a 1-dimensional chain of length L , with special marker states $|\langle\rangle\rangle, |\otimes\rangle\rangle$. Denote the orthogonal complement of $\text{span}(|\langle\rangle\rangle, |\otimes\rangle\rangle)$ in \mathbb{C}^d by \mathbb{C}^{d-2} .*

For any well-formed unidirectional Quantum Turing Machine $M = (\Sigma, Q, \delta)$ and any constant $K > 0$ ¹³, 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 .

¹²Indeed, it is general enough to admit universal QTMs [16].

¹³ K will be needed to deal with the extra space required by the QTM defined in Theorem 10. It will be set to $K = 3$ in Section 6 where all components of the proof are combined.

(iii) Denote $\mathcal{H}(L - 2) := (\mathbb{C}^{d-2})^{\otimes L-2}$ and define $S_{br} = \text{span}(|\langle \rangle\rangle) \otimes \mathcal{H}(L - 2) \otimes \text{span}(|\rangle \rangle) \subset \mathcal{H}$. Then 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 $\Omega(\zeta^L)$ time-steps, where $\zeta = |\Sigma \times Q|$ ¹⁴. 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 10, then h has the following form

$$h = A + (e^{i\pi\varphi} B + e^{i\pi 2^{-|\varphi|}} C + \text{h.c.}), \tag{4.5}$$

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 $H(L)|_{S_{br}}$ has energy $\Omega(1/T^3)$, where T is the total number of time-steps encoded in the computational history state.

4.1. Preliminaries

As our construction draws heavily on [32], we will follow their notation and terminology, which we summarise here. We divide the chain into multiple tracks:

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

Tracks 1-3 correspond to the clock register C and tracks 4-6 to the computational register Q in Definition 31 (see Figure 5 for an illustration).

The local Hilbert space at each site is the tensor product of the local Hilbert space of each of the six tracks $\mathcal{H} = \otimes_{i=1}^6 \mathcal{H}_i$, where

$$\begin{aligned} \mathcal{H}_1 &:= \text{span}\{|s\rangle\} \oplus \{|\langle \rangle\rangle, |\rangle \rangle\} \\ s &\in \{\overset{\rightarrow}{0}, \overset{\leftarrow}{0}, \overset{\rightarrow}{0}, \overset{\leftarrow}{0}, \overset{\rightarrow}{1}, \overset{\leftarrow}{1}, \bigcirc, \odot\} \text{ for } i \in \{1, \dots, K\}, \\ \mathcal{H}_2 &:= \text{span}\{|p\rangle\} \oplus \{|\langle \rangle\rangle, |\rangle \rangle\} \\ p &\in P' \cup \{\bigcirc, \odot\} \text{ where } P := P_L \cup P_N \cup P_R \text{ and } P' := P \cup P'_R, \end{aligned}$$

¹⁴This choice of ζ guarantees that the QTM M has enough time to halt (if it is going to halt) within this number of time-steps in the finite tape segment available.

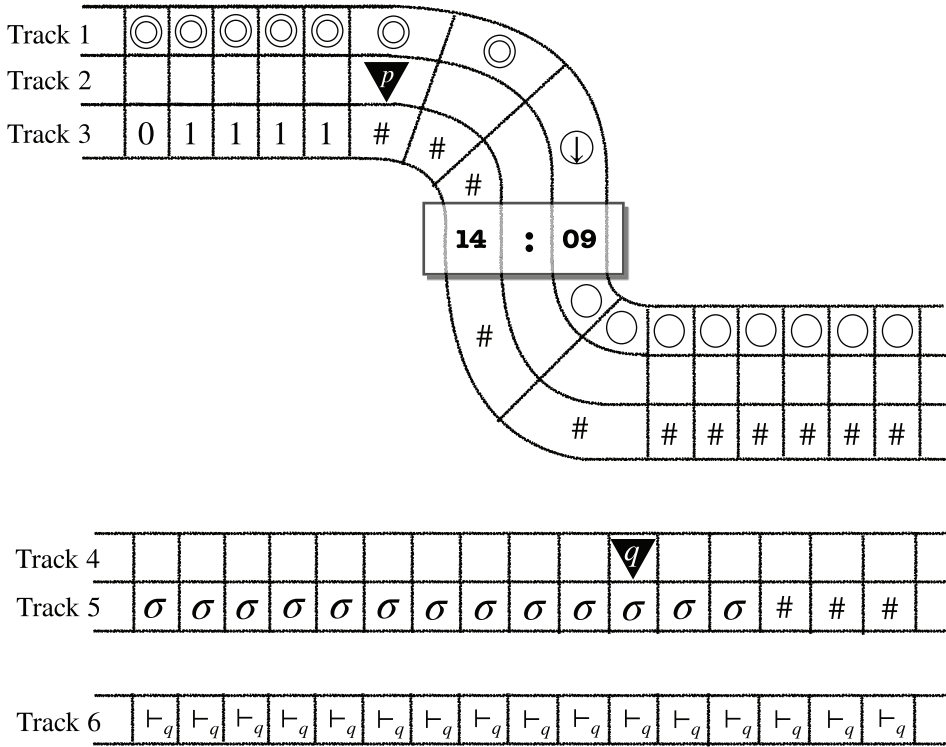


Figure 5. Cartoon illustration of the role of the different tracks. Track 1 acts as the clock’s “second hand”, repeating one full cycle for every clock increment. Track 3 acts as the “minute hand”, getting incremented by one for every cycle of Track 1. Track 2 stores the counter Turing Machine head and internal state machinery needed to implement this incrementing. These clock tracks drive the Quantum Turing Machine in Tracks 4 and 5, Track 4 storing the QTM head and internal state, Track 5 storing the QTM tape qudits. If the QTM ever halts or runs out of space on Track 5, it switches to writing garbage on the “time-wasting” trap stored on Track 6 to use up the remaining time on the clock.

$$\begin{aligned}
 \mathcal{H}_3 &:= \text{span}\{|\tau\rangle\} \oplus \{|\langle\rangle, |\rhd\rangle\} \\
 &\quad \tau \in \Xi \text{ where } \Xi := \{\dagger, \#, 0, \dots, \zeta - 1\} \text{ with } \zeta = |\Sigma \times Q|, \\
 \mathcal{H}_4 &:= \text{span}\{|q\rangle\} \oplus \{|\langle\rangle, |\rhd\rangle\} \\
 &\quad q \in Q' \cup P \cup P'_L \cup \{r_x\} \cup \{\circ, \odot\} \text{ where } Q := Q_L \cup Q_N \cup Q_R, Q' := Q \cup Q'_L \\
 &\quad \text{and } x \in Q, \\
 \mathcal{H}_5 &:= \text{span}\{|\sigma\rangle\} \oplus \{|\langle\rangle, |\rhd\rangle\} \\
 &\quad \sigma \in \Sigma, \\
 \mathcal{H}_6 &:= \text{span}\{|\gamma\rangle\} \oplus \{|\langle\rangle, |\rhd\rangle\} \\
 &\quad \gamma \in \Xi \cup \{\dagger_q\} \text{ where } q \in Q'.
 \end{aligned} \tag{4.6}$$

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 17 and 18), these sets are disjoint.

The \odot and \ominus Track 2 and 4 symbols are used for cells that do not currently hold the head.¹⁵ The role of the Track 1 states is described in Section 4.2. 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 4.6.2.

The marker states $|\langle\rangle\rangle, |\rangle\rangle$ appearing in Theorem 32 will just be the states $|\langle\rangle\rangle = \bigotimes_{i=1}^6 |\langle\rangle\rangle_{\text{track } i}$ and $|\rangle\rangle = \bigotimes_{i=1}^6 |\rangle\rangle_{\text{track } i}$. When a subset of tracks T is clear from the context, we will also write $|\langle\rangle\rangle, |\rangle\rangle$ to denote $\bigotimes_{i \in T} |\langle\rangle\rangle_{\text{track } i}$ and $\bigotimes_{i \in T} |\rangle\rangle_{\text{track } i}$, respectively.

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*.

We will sometimes use regular expressions to specify subsets of standard basis states. A regular expression denotes a (possibly infinite) subset of finite-length strings over a finite alphabet. Equivalently, it can be thought of as a pattern that matches all the strings in the subset and no others. Regular expressions, and the regular expression notation we will use, can be defined inductively in the standard way:

Definition 33 (Regular expression). Given a finite alphabet Σ , let Σ^* denote the set of all finite-length strings of symbols from that alphabet. The following are regular expressions:

- The empty regular expression ϵ denotes the set containing only the empty string.
- Any symbol $x \in \Sigma$ denotes the singleton set $\{x\}$ containing only the string x .


Given two regular expressions R_1, R_2 , the following are also regular expressions:

- $R_1 R_2$ denotes the set $\{xy : x \in R_1, y \in R_2\}$ of all concatenations of strings matching R and S .
- The alternation $R_1 | R_2$ denotes the set $R_1 \cup R_2$ of strings matching either R_1 or R_2 (or both).
- The Kleene-star R^* denotes the Kleene closure of R , i.e. the smallest set of strings S such that $\epsilon, R \in S$ and S is closed under concatenation.
- For $x_1, \dots, x_n \in \Sigma$, square brackets $[x_1, x_2, \dots, x_n] := x_1 | x_2 | \dots | x_n$ are a shorthand notation for alternations of symbols.

Parentheses (\dots) are used to group subexpressions, and have the highest precedence. Square brackets have higher precedence than the Kleene-star operator, which has higher precedence than the alternation operator. Concatenation has the lowest precedence.

Definition 34 (Bracketed state). We call standard basis states that match the regular expression $\langle \odot \cdot^* \ominus \rangle$ *bracketed states* (where \odot stands for any single-site state other than \langle, \rangle). We denote by S_{br} the subspace spanned by the bracketed states.

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




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 \langle and \rangle 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


¹⁵There 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.

Table 1. Single-site illegal states enforcing that end-marker states must appear across all tracks simultaneously.


Track i - j single-site illegal states, for all pairs $i \neq j$



involving $\langle \leftarrow \rangle$ and $\langle \rightarrow \rangle$ with symbols stretching vertically across all tracks, e.g.



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



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, even if we don't allow ourselves to use single-site penalty terms, 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.

By using single-site illegal states one can enforce that, on all legal states, the marker $|\langle \leftarrow \rangle|$ (resp. $|\langle \rightarrow \rangle|$) can only ever appear simultaneously on all tracks. In this case, if one restricts to global bracketed states, one also gets bracketed states in each of the tracks individually. The single-site illegal states enforcing this are summarised in [Table 1](#).

We will make repeated use of Lemma 5.2 from [32], 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 $\langle \leftarrow \rangle$ and end with a right-bracket $\langle \rightarrow \rangle$. 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 [32], 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 \mathcal{S}_{br} of bracketed states, with \langle at the left end of the chain and \rangle at the right. (We show how to enforce this later, in Section 6.)

4.2. 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 4.4. This section describes the Track 1 clock oscillator.

The legal configurations of Track 1 are states matching the regular expression $\langle \circ \circ^* [\overset{\rightarrow}{0}, \overset{\leftarrow}{0}, \overset{\leftarrow}{(i)}, \vec{i}, \overset{\leftarrow}{i}] \circ \circ^* \rangle$. By Lemma 36, we can enforce this using penalty terms. Independent of the configuration of any other track, a \rightarrow arrow (either $\overset{\rightarrow}{0}$ or \vec{i}) sweeps right along the chain until it reaches the end, whereupon it turns around and becomes a \leftarrow . (When a $\overset{\rightarrow}{0}$ 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 32 ensures that there is enough space for such sequences of transitions.) This \leftarrow arrow sweeps left until it reaches the beginning of the chain, at which point it turns around and becomes a \rightarrow again. We call this entire sequence an *oscillator cycle*. (One complete cycle is illustrated in Figure 4 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:

$$\begin{array}{ll}
 \boxed{\overset{\rightarrow}{0}} \boxed{\circ} \rightarrow \boxed{\circ} \boxed{\overset{\rightarrow}{0}}, & \boxed{\overset{\leftarrow}{0}} \boxed{\rangle} \rightarrow \boxed{\overset{\leftarrow}{(0)}} \boxed{\rangle}, \\
 \boxed{\circ} \boxed{\overset{\leftarrow}{(i)}} \rightarrow \boxed{\overset{\leftarrow}{(i+1)}} \boxed{\circ}, & \boxed{\circ} \boxed{\overset{\leftarrow}{(K)}} \rightarrow \boxed{\overset{\leftarrow}{0}} \boxed{\circ}, \\
 \boxed{\circ} \boxed{\overset{\leftarrow}{0}} \rightarrow \boxed{\overset{\leftarrow}{i}} \boxed{\circ}, & \boxed{\vec{i}} \boxed{\circ} \rightarrow \boxed{\circ} \boxed{\vec{i}}, \\
 \boxed{\overset{\leftarrow}{i}} \boxed{\rangle} \rightarrow \boxed{\overset{\leftarrow}{i}} \boxed{\rangle}, & \boxed{\circ} \boxed{\overset{\leftarrow}{i}} \rightarrow \boxed{\overset{\leftarrow}{i}} \boxed{\circ}.
 \end{array} \tag{4.7}$$

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

$$\begin{array}{ll}
 \boxed{\langle} \boxed{\overset{\leftarrow}{0}} \rightarrow \boxed{\langle} \boxed{\vec{i}}, & \boxed{\langle} \boxed{\overset{\leftarrow}{i}} \rightarrow \boxed{\langle} \boxed{\vec{i}}, \\
 & \boxed{\langle} \boxed{\neg p_\alpha} \rightarrow \boxed{\langle} \boxed{\neg p_\alpha},
 \end{array} \tag{4.8}$$

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

4.3. 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

$$\langle p_0 \circledast \rangle \text{ and } \langle \#^* \rangle, \tag{4.9}$$

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

$$\begin{array}{|c|c|} \hline \neg \langle & \overset{\rightarrow}{0} \\ \hline & \circledast \\ \hline \end{array}, \tag{4.10}$$

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):

$$\begin{array}{|c|c|} \hline \neg \langle & \overset{\rightarrow}{0} \\ \hline & \neg \# \\ \hline \end{array}, \tag{4.11}$$

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

At the beginning of the chain, we use illegal pairs to forbid Track 2 from being anything other than p_α when $\overset{\rightarrow}{0}$ is over it. We also forbid it from being p_α when a $\overset{\leftarrow}{1}$ is at the beginning of the chain:

$$\begin{array}{|c|c|} \hline \langle & \overset{\rightarrow}{0} \\ \hline & \neg p_\alpha \\ \hline \end{array}, \quad \begin{array}{|c|c|} \hline \langle & \overset{\leftarrow}{1} \\ \hline & p_\alpha \\ \hline \end{array}. \tag{4.12}$$

Thus any standard basis state containing a $\overset{\rightarrow}{0}$, $\overset{\leftarrow}{0}$ or $\overset{\leftarrow}{1}$ on Track 1 that matches the regular expressions (4.9), 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 $\overset{\leftarrow}{1}$ 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:

$$\begin{array}{|c|c|c|c|c|} \hline & \overset{\rightarrow}{0} & \circledast & \dots & \circledast \\ \hline \langle & p_\alpha & \circledast & \dots & \circledast \\ \hline & \vdash & \# & \dots & \# \\ \hline \end{array}. \tag{4.13}$$

4.4. Clock Counter

4.4.1. Counter TM construction

We construct the counter TM using a generalisation of the binary incrementing machine INC from Lemma 26, which will increment integers written in base- ζ instead of in binary. Recall that $\zeta = |\Sigma \times Q|$. Constructing a base- ζ version of INC is largely a straightforward extension of Lemma 26.

Ensuring reversibility is similar to the binary case. The TM state q_1 scans along the digits of the little-endian input, changing all $\zeta - 1$'s to 0's as it goes, until it finds the first digit different from $\zeta - 1$. It then increments this by one, to add the carry from all the previous $\zeta - 1$ s digits. If no such digit exists, a 1 will be appended to the end of the input. The difference with the binary case is that now there are three cases (or computational paths) that we must distinguish: the case in which a digit other than 0 is incremented by one, and two different cases in which a 1 can be written to the tape: incrementing 0, or writing a new 1 in place of the first # symbol after the input. These three computational paths correspond

respectively to TM states p_2, p'_2 and p''_2 . The latter two paths are first merged into state p'_3 . The former transitions instead into state p_3 . Then, in a second stage, the two remaining paths are merged into state p_4 which, as in the binary case, is the state that then returns the head to the starting cell and halts.

In pseudocode:

Algorithm 14 INC_ζ

```

1: input: string  $s \in \{0, \dots, \zeta - 1, \vdash, \#\}^*$  consisting of the symbol  $\vdash$  followed by
2:     the little-endian base- $\zeta$  representation of  $n \in \mathbb{N}$ . ▷  $q_0$ 
3: step Right ▷  $p_1$ 
4: while read a  $\zeta - 1$  do
5:     write a 0
6:     step Right
7: end while
8: if read  $0 < i < \zeta - 2$  then
9:     write  $i + 1$ 
10:    step Right ▷  $p_2$ 
11:    step Left ▷  $p_3$ 
12: else
13:    if read 0 then
14:        write 1
15:        step right ▷  $p'_2$ 
16:    else if read # then
17:        write a 1
18:        step Right ▷  $p''_2$ 
19:    end if
20:    step Left ▷  $p'_3$ 
21: end if
22: step Left ▷  $p_4$ 
23: while not read a  $\vdash$  do
24:    step Left
25: end while
26: halt ▷  $p_f$ 

```

The following well-formed, normal-form, partial transition function implements a reversible TM INC_ζ 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 | (\vdash, p_0, R) | $(\#, p_0, R)$ | $(0, p_0, R)$ | $(1, p_0, R)$ | (i, p_0, R) | $(\zeta - 1, p_0, R)$ | (4.14) |

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

We now modify this machine slightly in order to make it loop forever. However, we must ensure that the loop reenters the initial state p_0 of INC_ζ reversibly. To accomplish 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.

In pseudocode:

Algorithm 15 BASE- ζ COUNTER

- | | | |
|----|----------------------------------------------------------|---------------------------|
| 1: | input: \vdash in first cell, followed by a blank tape. | $\triangleright p_\alpha$ |
| 2: | step Right | $\triangleright p_{-1}$ |
| 3: | step Left | $\triangleright p_0$ |
| 4: | loop | |
| 5: | call INC_ζ | $\triangleright p_f$ |
| 6: | step Right | $\triangleright p'_f$ |
| 7: | step Left | $\triangleright p_0$ |
| 8: | end loop | |

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) | $(\zeta - 1, p_4, L)$ | (4.15) |
| 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) | $(\zeta - 1, p_\alpha, R)$ | |

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

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:

$$\begin{array}{|c|c|} \hline p & \cdot \\ \hline a & b \\ \hline \end{array} \quad \text{or} \quad \begin{array}{|c|c|} \hline \cdot & p \\ \hline a & b \\ \hline \end{array}. \tag{4.16}$$

A single row always denotes a section of tape:

$$\begin{array}{|c|c|} \hline a & b \\ \hline \end{array}. \tag{4.17}$$

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 introduce the following notation. We denote the starting cell by placing a \langle symbol on the left. We denote tape cells *other than* the starting cell with $\neg\langle$ on the left. For illustration, two examples of this would be:

$$\begin{array}{|c|c|} \hline \langle & p \\ \hline a & \\ \hline \end{array}, \quad \begin{array}{|c|c|} \hline \neg\langle & p \\ \hline a & \\ \hline \end{array}. \tag{4.18}$$

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:

$$\begin{array}{|c|c|} \hline \neg\langle & p_\alpha \\ \hline & \cdot \\ \hline \end{array}. \tag{4.19}$$

We also declare any configuration in state p_ω to be illegal:¹⁶

$$\begin{array}{|c|c|} \hline p_\omega & \cdot \\ \hline \cdot & \cdot \\ \hline \end{array}. \tag{4.20}$$

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

$$\begin{array}{|c|c|} \hline p_\alpha & \cdot \\ \hline \cdot & \neg\# \\ \hline \end{array}. \tag{4.21}$$

The transition table in (4.15) is partial; some transitions are never used, so were not defined. For each (p, τ) for which no transition rule is defined in (4.15), 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

$$\begin{array}{|c|} \hline p \\ \hline \tau \\ \hline \end{array}. \tag{4.22}$$

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 (4.15), we declare the corresponding TM head and tape configurations to be illegal, i.e. any configurations containing one of the following illegal pairs:

$$\begin{array}{|c|c|}, \quad \begin{array}{|c|c|}, \quad \begin{array}{|c|c|} \hline \cdot & p_R \\ \hline \tau & \cdot \\ \hline \end{array}, \quad \begin{array}{|c|c|} \hline p_L & \cdot \\ \hline \cdot & \tau \\ \hline \end{array}, \quad \begin{array}{|c|c|} \hline p_N & \cdot \\ \hline \tau & \cdot \\ \hline \end{array}. \tag{4.23}$$

¹⁶We could alternatively remove state p_ω from the construction entirely. We have elected to retain it, since inclusion of a halting state is necessary to conform to the definition of a TM.

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

$$\begin{array}{|c|c|} \hline \cdot & p \\ \hline \# & \# \\ \hline \end{array} \cdot \tag{4.24}$$

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

$$\vdash [0, \dots, \zeta - 1]^* \#^* \tag{4.25}$$

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

$$\begin{array}{|c|c|} \hline 0 & \# \\ \hline \end{array} \cdot \tag{4.26}$$

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:

$$\begin{array}{|c|c|} \hline 0 & \textcircled{>} \\ \hline \end{array} \tag{4.27}$$

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 $[0, L]$.*

Before delving into the technical proof, let us briefly explain why the set of configurations reached by the counter TM during its computation (or rather, the complement of this set) cannot be characterised by illegal pairs. And why, therefore, one needs to exclude many of the invalid configurations by showing that they necessarily go on to evolve into an illegal configuration.

As an illustrative example, consider the valid configurations associated with state p_1 located in position 5. These are of the form:

$$\begin{array}{|c|c|c|c|c|c|} \hline \vdash & 0 & 0 & 0 & \tau & \dots \\ \hline \end{array} \tag{4.28}$$

with τ any element other than \vdash . Any attempt to guarantee a 0 in position 3 in this configuration by declaring extra illegal pairs (which we recall can only be defined between nearest neighbours) would

contradict the fact that, for state p_0 , all configurations of the form:

$$\begin{array}{c}
 p_0 \\
 \boxed{\vdash \quad i_1 \quad i_2 \quad i_3 \quad i_4 \quad \cdots}
 \end{array}$$

with $i_j \in \{0, \dots, \zeta - 1\}$ are valid configurations reached by the counter TM during its evolution.

Proof. The first part of the Lemma is true by construction, since the partial transition rules of (4.15) 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

$$\vdash (\neg\#)^* \neg[0, \#]\#^* \tag{4.28}$$

Configurations in which the head is more than one cell away from a non-blank symbol are illegal by (4.24), 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 (4.19), except those with the head in the starting cell. The latter are illegal by (4.21) 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 (4.15), so by (4.22) the only legal p_{-1} configurations have their head over a #. Moreover, there is only one transition into p_{-1} in (4.15). Hence, evolving any such configuration one step *backwards* according to (4.15) either enters a configuration that is illegal due to (4.22), 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

$$\begin{array}{c}
 p_1 \\
 \boxed{\vdash \quad 0 \quad \cdots \quad 0 \quad \tau \quad \cdots} \\
 \underbrace{\hspace{10em}}_n \\
 \tau \in \Xi / \{\vdash\}, \quad n \geq 0
 \end{array}$$

with the rest of the tape compatible with (4.28) 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 :

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

Evolving any such configuration *backwards* according to (4.15) either enters into an illegal configuration due to (4.22), 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 :

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

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 (4.15), so this configuration is illegal by (4.23).

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

$$\begin{array}{c}
 \begin{array}{|c|c|c|c|c|c|}
 \hline
 \vdash & 0 & \cdots & 0 & x & \tau & \cdots \\
 \hline
 \end{array} \\
 \underbrace{\hspace{6em}} \\
 n \\
 x \in \Xi/\{\vdash, \#, 0, 1\}, \quad \tau \in \Xi/\{\vdash\}, \quad n \geq 0.
 \end{array}$$

with the rest of the tape compatible with (4.28) 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

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

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 (4.15), so these configurations are illegal by (4.23).

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

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

Evolving any such configuration *backwards* one time step according to (4.15) either produces an illegal configuration due to (4.22), or produces the configuration:

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

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

$$\begin{array}{c}
 \begin{array}{|c|c|c|c|c|}
 \hline
 \vdash & 0 & \dots & 0 & \tau & \dots \\
 \hline
 \end{array} \\
 \underbrace{\hspace{4em}} \\
 n \\
 \tau \in \Xi/\{\vdash, \#, 0, 1\}, \quad n \geq 0.
 \end{array}$$

with the rest of the tape compatible with (4.28) 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

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

There is no transition out of (p_3, x) where $\tau \in \{\vdash, \#, 0, 1\}$ in (4.15), so all such configurations are illegal by (4.22). The remaining legal configurations have at least one tape symbol x to the left of the head that is neither 0, #, nor \vdash :

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

Evolving any such configuration forwards one time step according to (4.15) either produces an illegal configuration (for example if $n = 0$), or produces the configuration:

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

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




with the rest of the tape compatible with (4.28) 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 :



Evolving any such configuration according to (4.15) steps the head left over the 0's until the head is over the x :



But there is no transition out of (p_4, x) in (4.15), so this configuration is illegal by (4.22).

We now analyse the case in which p_4 is over \vdash . Evolving one step backwards according to (4.15) steps the head right and either enters an illegal configuration, transitions to state p_4 with the head 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 (4.15), 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 (4.15) for one time step transitions to the state p_4 , and we have already shown than 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 (4.15), 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 (4.15) for one time step transitions to the state p_f , and we have already shown than 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 (4.15), 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 (4.15) into p_0 are from p_1 or p'_f , evolving any such configuration one step backwards according to (4.15) 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. □

4.4.2. 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 a regular expression $\langle \circ \circ^* p \circ \circ^* \rangle$ 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 $\overset{\rightarrow}{\uparrow}$ 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 $\overset{\rightarrow}{\uparrow}$ 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 $\overset{\rightarrow}{\uparrow}$ 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

$$\begin{array}{|c|} \hline \overset{\rightarrow}{\uparrow} \\ \hline p \\ \hline \sigma \\ \hline \end{array} \tag{4.29}$$

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 $\overset{\rightarrow}{\uparrow}$ 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 $\overset{\rightarrow}{\uparrow}$, 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 [32]). Thus, when the $\overset{\rightarrow}{\uparrow}$ is on top of the head and the head will move right, we will update the tape, move the head, and step the $\overset{\rightarrow}{\uparrow}$ 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 $\overset{\rightarrow}{\uparrow}$ now on top of the p'_R state, we transition to the correct state p_R , and step the $\overset{\rightarrow}{\uparrow}$ 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 $\overset{\rightarrow}{\uparrow}$, so we add a penalty term to make all other p'_R configurations illegal:

$$\begin{array}{|c|} \hline \neg \overset{\rightarrow}{\uparrow} \\ \hline p'_R \\ \hline \end{array} \tag{4.30}$$

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 $p \in P \setminus \{p_\alpha\}$, $p_{L,N,R} \in P_{L,N,R}$ 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):

$$\begin{array}{|c|c|} \hline \overrightarrow{\tau} & \odot \\ \hline p & \odot \\ \hline \sigma & \cdot \\ \hline \end{array} \rightarrow \begin{array}{|c|c|} \hline \circ & \overrightarrow{\tau} \\ \hline p_N & \odot \\ \hline \tau & \cdot \\ \hline \end{array}, \quad \begin{array}{|c|} \hline \overrightarrow{\tau} \\ \hline p \\ \hline \sigma \\ \hline \end{array} \rightarrow \begin{array}{|c|} \hline \overleftarrow{\tau} \\ \hline p_N \\ \hline \tau \\ \hline \end{array}, \quad (4.31a)$$

$$\begin{array}{|c|c|} \hline \overrightarrow{\tau} & \odot \\ \hline \circ & p \\ \hline \cdot & \sigma \\ \hline \end{array} \rightarrow \begin{array}{|c|c|} \hline \circ & \overrightarrow{\tau} \\ \hline p_L & \odot \\ \hline \cdot & \tau \\ \hline \end{array}, \quad \begin{array}{|c|c|} \hline \overrightarrow{\tau} & \odot \\ \hline p & \odot \\ \hline \sigma & \cdot \\ \hline \end{array} \rightarrow \begin{array}{|c|c|} \hline \circ & \overrightarrow{\tau} \\ \hline \circ & p'_R \\ \hline \tau & \cdot \\ \hline \end{array}, \quad (4.31b)$$

$$\begin{array}{|c|c|} \hline \overrightarrow{\tau} & \odot \\ \hline p'_R & \odot \\ \hline \cdot & \cdot \\ \hline \end{array} \rightarrow \begin{array}{|c|c|} \hline \circ & \overrightarrow{\tau} \\ \hline p_R & \odot \\ \hline \cdot & \cdot \\ \hline \end{array}, \quad \begin{array}{|c|} \hline \overrightarrow{\tau} \\ \hline p'_R \\ \hline \cdot \\ \hline \end{array} \rightarrow \begin{array}{|c|} \hline \overleftarrow{\tau} \\ \hline p_R \\ \hline \cdot \\ \hline \end{array}. \quad (4.31c)$$

Transitions in which the head moves left from the end of the chain are already covered by the left-moving transition in (4.31b). 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 4.4 will never attempt such a transition when initialised properly. Moreover, we just declared such configurations illegal in (4.29).)

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 incrementer TM is implemented in (4.14)). The incrementer 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 $O(L\zeta^L)$) 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 (4.31a) and (4.31b) will apply during each $\overrightarrow{\tau}$ sweep. Similarly, for any configuration containing a $p'_R \in P'_R$, a (4.31c) rule will apply. It is easy to verify that if any of the transition rules from (4.31a)–(4.31c) is applied to any legal configuration, none of the (4.31a) and (4.31b) rules can apply again to the resulting configuration. If the final (4.31b) rule applies, the (4.31c) 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 (4.31a)–(4.31c) therefore implement a single step of the TM during each left-to-right sweep of the $\overrightarrow{\tau}$ state on Track 1, as required. We call this phase of the evolution, in which Track 1 contains a $\overrightarrow{\tau}$ or $\overleftarrow{\tau}$, the *computation phase*.

We also modify the clock oscillator rules from Section 4.2 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:

$$\begin{array}{|c|c|} \hline \overrightarrow{\tau} & \odot \\ \hline \circ & \circ \\ \hline \circ & \circ \\ \hline \end{array} \rightarrow \begin{array}{|c|c|} \hline \circ & \overrightarrow{\tau} \\ \hline \circ & \circ \\ \hline \circ & \circ \\ \hline \end{array}, \quad \begin{array}{|c|c|} \hline \overrightarrow{\tau} & \odot \\ \hline \circ & \odot \\ \hline \odot & \odot \\ \hline \end{array} \rightarrow \begin{array}{|c|c|} \hline \circ & \overrightarrow{\tau} \\ \hline \circ & \odot \\ \hline \odot & \odot \\ \hline \end{array}, \quad \begin{array}{|c|} \hline \overrightarrow{\tau} \\ \hline \odot \\ \hline \odot \\ \hline \end{array} \rightarrow \begin{array}{|c|} \hline \overleftarrow{\tau} \\ \hline \odot \\ \hline \odot \\ \hline \end{array}. \quad (4.32)$$

We must also enforce illegality of all the configurations declared to be illegal in Section 4.4.1. To do this, we simply introduce local penalty terms corresponding to the combinations declared illegal in (4.19)–(4.24), (4.26) and (4.27), and enforce the Track 3 regular expression of (4.25) using penalty terms and Lemma 36.

The complete set of Track 1 to 3 transition rules is summarised in Table 2. 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 3.

Table 2. All transition rules for Tracks 1 to 3.

| $\vec{0}, \overleftarrow{0}$ rules | Track 1 rules | $\overleftarrow{1}$ rules |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------|
| $\begin{array}{c} \boxed{\vec{0} \quad \vec{0}} \rightarrow \boxed{\vec{0} \quad \vec{0}} \\ \boxed{\vec{0} \quad \triangleright} \rightarrow \boxed{\overleftarrow{0} \quad \triangleright} \\ \boxed{\circ \quad \overleftarrow{0}} \rightarrow \boxed{\overleftarrow{0} \quad \circ} \\ \boxed{\circ \quad \overleftarrow{0}} \rightarrow \boxed{\overleftarrow{0} \quad \circ} \end{array}$ | | $\boxed{\circ \quad \overleftarrow{1}} \rightarrow \boxed{\overleftarrow{1} \quad \circ} \quad (*)^1$ |
| $\vec{0}, \overleftarrow{0}$ rules | Track 1 and 2 rules | $\overleftarrow{1}$ rules |
| $\boxed{\leftarrow \quad \overleftarrow{0}} \rightarrow \boxed{\leftarrow \quad \overleftarrow{1}}$ | $\begin{array}{c} \boxed{\overleftarrow{1} \quad \circ} \rightarrow \boxed{\circ \quad \overleftarrow{1}} \\ \boxed{\overleftarrow{1} \quad \circ} \rightarrow \boxed{\circ \quad \overleftarrow{1}} \\ \boxed{\overleftarrow{1} \quad \triangleright} \rightarrow \boxed{\overleftarrow{1} \quad \triangleright} \quad (*) \end{array}$ | $\boxed{\leftarrow \quad \overleftarrow{1}} \rightarrow \boxed{\leftarrow \quad \overleftarrow{1}} \quad (*)$ |
| $\vec{0}, \overleftarrow{0}$ rules | Track 1, 2 and 3 rules | $\overleftarrow{1}$ rules |
| | $\begin{array}{c} \boxed{\overleftarrow{1} \quad \circ} \rightarrow \boxed{\circ \quad \overleftarrow{1}} \\ \boxed{\circ \quad p} \rightarrow \boxed{pL \quad \circ} \\ \boxed{\cdot \quad \sigma} \rightarrow \boxed{\cdot \quad \tau} \\ \boxed{\overleftarrow{1} \quad \circ} \rightarrow \boxed{\circ \quad \overleftarrow{1}} \\ \boxed{p \quad \circ} \rightarrow \boxed{pN \quad \circ} \\ \boxed{\sigma \quad \cdot} \rightarrow \boxed{\tau \quad \cdot} \\ \boxed{\overleftarrow{1} \quad \triangleright} \rightarrow \boxed{pN \quad \triangleright} \quad (*) \\ \boxed{\overleftarrow{1} \quad \circ} \rightarrow \boxed{\circ \quad \overleftarrow{1}} \\ \boxed{p \quad \circ} \rightarrow \boxed{\tau \quad p'_R} \\ \boxed{\sigma \quad \cdot} \rightarrow \boxed{\cdot \quad \cdot} \\ \boxed{\overleftarrow{1} \quad \circ} \rightarrow \boxed{\circ \quad \overleftarrow{1}} \\ \boxed{p'_R \quad \circ} \rightarrow \boxed{pR \quad \circ} \\ \boxed{\cdot \quad \cdot} \rightarrow \boxed{\cdot \quad \cdot} \\ \boxed{\overleftarrow{1} \quad \triangleright} \rightarrow \boxed{pR \quad \triangleright} \quad (*) \end{array}$ | |

¹Transition rules marked with an (*) will be replaced in Section 4.6 by a set of rules that also act on Tracks 4, 5 and 6.

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

| Regular expressions | | | |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|----------------|---------|
| Track 1 | Track 2 | Track 3 | |
| $\langle \circ \circ^* [\overset{\circ}{\leftarrow}, \overset{\circ}{\leftarrow}, \overset{\circ}{\leftarrow}, \overset{\circ}{\rightarrow}, \overset{\circ}{\rightarrow}] \circ \circ^* \rangle$ | | | |
| $\langle \circ \circ^* p \circ \circ^* \rangle$ | | | |
| $\langle \circ \circ^* i^* \# \circ \circ^* \rangle$ | | | |
| 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 (4.15) | | | |
| | | | |
| Illegal pairs | | | |
| Tracks 1, 2 and 3 if $\delta(p, \sigma) = (p_L, \tau, L)$ (4.29) | | | |
| | | | |

4.5. 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 $\langle \circ \circ^* [\overset{\circ}{\leftarrow}, \overset{\circ}{\leftarrow}, \overset{\circ}{\leftarrow}, \overset{\circ}{\rightarrow}, \overset{\circ}{\rightarrow}] \circ \circ^* \rangle$, 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 $\overset{\rightarrow}{0}$, $\overset{\leftarrow}{0}$, $\overset{\leftarrow}{1}$, $\overset{\rightarrow}{1}$ or $\overset{\leftarrow}{1}$ on Track 1. Thus clearly only transition rules from one of the sets in Table 2 (the set of $\overset{\leftarrow}{0}$, $\overset{\leftarrow}{0}$ rules, $\overset{\rightarrow}{1}$ rules, or $\overset{\leftarrow}{1}$ rules) can apply in the forward direction. The left hand sides of the rules within the $\overset{\leftarrow}{0}$, $\overset{\leftarrow}{0}$ and $\overset{\leftarrow}{1}$ 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 $\overset{\rightarrow}{1}$ 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 $\overset{\leftarrow}{0}$ turning rule that changes a $\overset{\rightarrow}{1}$ into a $\overset{\leftarrow}{0}$ (in the backwards direction). But this rule clearly cannot apply at the same time as the right hand side of any rule from the $\overset{\rightarrow}{1}$ set, which concludes the proof of the first part of the lemma.

It is straightforward to verify that all the transition rules in Table 2 preserve well-formedness, implying the second part. □

Recall that the state $|\phi_0\rangle$ is the standard basis state with $\langle \overset{\leftarrow}{0} \overset{\rightarrow}{0} \overset{\leftarrow}{0} \rangle$ on Track 1, $\langle p_\alpha \overset{\leftarrow}{0} \rangle$ on Track 2, and $\langle \vdash \#^* \rangle$ on Track 3. This corresponds to having the clock oscillator in the first state of the Track 1 sequence (see Figure 4), 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 $\overset{\rightarrow}{0}$ on Track 1. The only transition rules that apply when Track 1 contains a $\overset{\rightarrow}{0}$, $\overset{\leftarrow}{0}$ or $\overset{\leftarrow}{1}$ (see Table 2) 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 4.2 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 $\overset{\rightarrow}{0}$, $\overset{\leftarrow}{0}$ or $\overset{\leftarrow}{1}$ on Track 1, and Tracks 2 and 3 remain in their initial configurations.

At $t = 2(L - 2) + 1$, the $\overset{\leftarrow}{0}$ turns around and (because Track 2 contains a p_α at the beginning of the chain) becomes a $\overset{\rightarrow}{1}$. This $\overset{\rightarrow}{1}$ 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 4.4 that the transition rules implement one step of the counter TM during each left-to-right sweep of the $\overset{\rightarrow}{1}$ on Track 1, and then move the $\overset{\leftarrow}{1}$ back to the beginning of the chain in the second half of the oscillator cycle. Note that the counter TM construction of Section 4.4 never transitions back to its initial state p_α . So when the $\overset{\leftarrow}{1}$ reaches the beginning of the chain, it will never find a p_α on Track 3. Thus it will always transition into $\overset{\rightarrow}{1}$, never into $\overset{\leftarrow}{0}$ (see Table 2). Therefore, for all $t > 2(L - 2)$ the state $|\phi_t\rangle$ contains a $\overset{\rightarrow}{1}$ or $\overset{\leftarrow}{1}$ on Track 1. At $t = 2(L - 2)(n + 1)$ the $\overset{\rightarrow}{1}$ 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 \rangle 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 \rangle instead of a blank Tape 2 cell. When the $\overset{\rightarrow}{1}$ 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 2). 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 $O(L\zeta^L)$ steps of the TM (see Section 4.4) each of which requires $2(L - 2)$

transitions, we have that $T = O(L^2 \zeta^L)$. (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 $|\phi_t\rangle$ 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 $\overset{\circ}{\rightarrow}$, $\overset{\circ}{\leftarrow}$ or $\overset{\circ}{\downarrow}$. 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 $\overset{\circ}{\rightarrow}$ or $\overset{\circ}{\leftarrow}$, $\overset{\circ}{\downarrow}$ 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 3.

Case 2:

Track 1 contains $\overset{\circ}{\uparrow}$ or $\overset{\circ}{\downarrow}$. 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 (4.31b) to a configuration containing a $p \notin P'_R$ on Track 2, and evolve by a single forwards transition from (4.31c) 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 $\overset{\circ}{\uparrow}$, 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 \leftarrow and \rightarrow . 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 $\overset{\circ}{\uparrow}$ is in the starting cell then we have the state $|\phi_t\rangle$ for $t = 2(L - 2) + 1$. If not, either we have $\overset{\circ}{\uparrow}$ to the right of the head or $\overset{\circ}{\downarrow}$ 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 $\overset{\circ}{\downarrow}$ over the p_α within $O(L)$ transitions, which is illegal by the second Track 1 and 2 illegal pair from Table 3.

If the Track 2–3 configuration is not the initial one, then all positions of $\overset{\circ}{\uparrow}$ and $\overset{\circ}{\downarrow}$ give rise to elements in the set $\{|\phi_t\rangle\}_t$ except exactly those declared illegal in (4.29) and (4.30). This completes the proof of the Lemma. □

4.6. QTM Hamiltonian

With the clock in place, encoding an arbitrary quantum Turing Machine in the Hamiltonian is now possible using ideas from [32] (and is similar to the construction used in Section 4.4 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 $\{|\psi^i_{abcd}\rangle\}_i$ 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 $|ab\rangle |\psi^i_{abcd}\rangle \rightarrow |cd\rangle U_{abcd} |\psi^i_{abcd}\rangle$ for all i .
- (ii) $h_{\text{pen}} \in \mathcal{B}((\mathbb{C}^C \otimes \mathbb{C}^Q)^{\otimes 2})$ 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 2. 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 2, 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, firstly, that the new transition involving the rest of the tracks, when restricted to Tracks 1–3, recover exactly the transitions removed from Table 2; and, secondly, that orthogonality is preserved in the construction.

By the standard form,

$$h_{\text{trans}} = \sum_{ab \rightarrow cd} (|ab\rangle - |cd\rangle)(\langle ab| - \langle cd|) \otimes (2\mathbb{1} - U_{abcd} - U_{abcd}^\dagger). \tag{4.33}$$

Now, if we start with a family of QTM P_n that satisfies part (iii) of Theorem 10, it will be immediate from our construction that the partial definition of all U_{abcd} will only involve elements in the set

$$\mathcal{S} = \left\{ 0, 1, \pm \frac{1}{\sqrt{2}}, e^{i\pi\varphi}, e^{i\pi 2^{-|\varphi|}} \right\}. \tag{4.34}$$

As argued above, the same will hold for the full U_{abcd} , which gives part (vi) in Theorem 32.

4.6.1. 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 \leftarrow state sweeps from right to left. This is done by including the following transition rule terms (cf. (4.31a)–(4.31c)), where q is any state in $Q \setminus \{q_f\}$ (i.e. we exclude all the transitions out of the QTM’s final state):

$$\left| \begin{array}{|c|c|} \hline \bigcirc & \leftarrow \\ \hline q & \bigcirc \\ \hline \sigma & \cdot \\ \hline \end{array} \right\rangle \rightarrow \sum_{\tau, q_R} \delta(q, \sigma, \tau, q_R, R) \left| \begin{array}{|c|c|} \hline \leftarrow & \bigcirc \\ \hline \bigcirc & q_R \\ \hline \tau & \cdot \\ \hline \end{array} \right\rangle + \sum_{\tau, q_N} \delta(q, \sigma, \tau, q_N, N) \left| \begin{array}{|c|c|} \hline \leftarrow & \bigcirc \\ \hline q_N & \bigcirc \\ \hline \tau & \cdot \\ \hline \end{array} \right\rangle \tag{4.35a}$$

$$+ \sum_{\tau, q_L} \delta(q, \sigma, \tau, q_L, L) \left| \begin{array}{|c|c|} \hline \leftarrow & \bigcirc \\ \hline q'_L & \bigcirc \\ \hline \tau & \cdot \\ \hline \end{array} \right\rangle,$$

$$\left| \begin{array}{cc} \bigcirc & \leftarrow \\ \bigcirc & q'_L \\ \cdot & \cdot \end{array} \right\rangle \rightarrow \left| \begin{array}{cc} \leftarrow & \bigcirc \\ q_L & \bigcirc \\ \cdot & \cdot \end{array} \right\rangle. \tag{4.35b}$$

As in the implementation of the reversible counter TM (Section 4.4), these transition rules implement the **R**ight-moving, **N**on-moving and **L**eft-moving transitions separately. This time, since the \leftarrow 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 (4.35a) can never apply, because they require the \leftarrow 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 than¹⁷ $L + 1$. If the Track 4 QTM head ends up next to the \rightarrow 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 \leftarrow arrow to cover the case where there is no QTM head present:

$$\begin{array}{ccc} \begin{array}{|c|c|} \triple \bigcirc & \leftarrow \\ \triple \bigcirc & \bigcirc \end{array} & \rightarrow & \begin{array}{|c|c|} \triple \leftarrow & \bigcirc \\ \triple \bigcirc & \bigcirc \end{array}, & \begin{array}{|c|c|} \triple \bigcirc & \leftarrow \\ \triple \bigcirc & \bigcirc \end{array} & \rightarrow & \begin{array}{|c|c|} \triple \leftarrow & \bigcirc \\ \triple \bigcirc & \bigcirc \end{array}, & \begin{array}{|c|c|} \triple \bigcirc & \leftarrow \\ \triple \bigcirc & q \end{array} & \rightarrow & \begin{array}{|c|c|} \triple \leftarrow & \bigcirc \\ \triple \bigcirc & q \end{array}, \end{array} \tag{4.36}$$

4.6.2. QTM Unitarity

With the set of transition rule terms defined so far in (4.35a), 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 configurations trying to move the head left when it is located at the beginning of the chain.¹⁸

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 4.7.

If we complete the partial U (associated with the Track 1–3 transitions $\begin{array}{|c|c|} \triple \bigcirc & \leftarrow \\ \triple \bigcirc & \bigcirc \end{array} \rightarrow \begin{array}{|c|c|} \triple \leftarrow & \bigcirc \end{array}$) as defined so far by (4.35a)–(4.37) 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.

It will be important in what follows (Section 4.7) to ensure that the local update rules are extended to a unitary. We would therefore like 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

¹⁷Two cells are always used by the \leftarrow , \rightarrow markers.

¹⁸Although the QTMs we will encode will in fact never do this.

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 (4.35a) that encode all the transition rules of the counter TM from (4.15).¹⁹ Note that Track 6 contains additional 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 (4.15) 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 (4.35a) 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:

$$\begin{array}{|c|c|} \hline \bigcirc & \leftarrow \vdash \\ \hline q_f & \bigcirc \\ \hline \vdash & \cdot \\ \hline \end{array} \longrightarrow \begin{array}{|c|c|} \hline \leftarrow \vdash & \bigcirc \\ \hline p_\alpha & \bigcirc \\ \hline \vdash_{q_f} & \cdot \\ \hline \end{array} . \tag{4.37}$$

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 32, 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:

$$\begin{array}{|c|c|} \hline \leftarrow & \leftarrow \vdash \\ \hline \neg p_\alpha & \\ \hline q'_L & \\ \hline \vdash & \\ \hline \end{array} \longrightarrow \begin{array}{|c|c|} \hline \leftarrow & \leftarrow \vdash \\ \hline \neg p_\alpha & \\ \hline p_\alpha & \\ \hline \vdash_{q'_L} & \\ \hline \end{array} . \tag{4.38}$$

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 steps 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 4.6, 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_q (depending on the state q it was in immediately before). However, we have to make it compatible with the transitions already defined for Tracks 1–3. Hence we include the following

¹⁹Since these are the transition rules of a classical reversible TM, they have a particularly simple form, the right hand side of each (4.35) rule containing a single term.

transitions acting on Tracks 1,2,4 and 1,2,3,4:

$$\begin{array}{ccc}
 \begin{array}{|c|c|} \hline \textcircled{\leftarrow} \\ \hline \textcircled{\leftarrow} \\ \hline q \\ \hline \end{array} & \rightarrow & \begin{array}{|c|c|} \hline \textcircled{\leftarrow} \\ \hline \textcircled{\leftarrow} \\ \hline r_q \\ \hline \end{array} & , &
 \begin{array}{|c|c|} \hline \textcircled{\leftarrow} \\ \hline p \\ \hline \sigma \\ \hline q \\ \hline \end{array} & \rightarrow & \begin{array}{|c|c|} \hline \textcircled{\leftarrow} \\ \hline p_N \\ \hline \tau \\ \hline r_q \\ \hline \end{array} & , &
 \begin{array}{|c|c|} \hline \textcircled{\leftarrow} \\ \hline p'_R \\ \hline \cdot \\ \hline q \\ \hline \end{array} & \rightarrow & \begin{array}{|c|c|} \hline \textcircled{\leftarrow} \\ \hline p_R \\ \hline \cdot \\ \hline r_q \\ \hline \end{array}
 \end{array} \tag{4.39}$$

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

$$\begin{array}{|c|c|} \hline \textcircled{\leftarrow} \\ \hline r_q \\ \hline \end{array} \rightarrow \begin{array}{|c|c|} \hline \textcircled{\leftarrow} \\ \hline r_q \\ \hline \end{array} \tag{4.40}$$

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:

$$\begin{array}{|c|c|} \hline \textcircled{\leftarrow} \\ \hline \neg p_\alpha \\ \hline r_q \\ \hline \vdash \\ \hline \end{array} \rightarrow \begin{array}{|c|c|} \hline \textcircled{\leftarrow} \\ \hline \neg p_\alpha \\ \hline p_\alpha \\ \hline \vdash_q \\ \hline \end{array} \tag{4.41}$$

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

The transition rules in Table 4 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 (4.35a)–(4.41) are manifestly orthogonal. The right hand sides of all (4.36)–(4.41) and (4.35b) rules are clearly mutually orthogonal and normalised, and are also orthogonal to the right hand sides of all (4.35) rules.

It remains to show that (4.35) defines an isometry. Recall that there are two sets of (4.35) terms, one encoding the QTM transition rules, and the other the time-wasting counter TM. 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 of internal states Q' and P' (respectively). 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 (4.35). Then

$$\begin{aligned}
 \|\psi(p, \sigma)\|^2 &= \sum_{\tau, q_R} |\delta(p, \sigma, \tau, q_R, R)|^2 + \sum_{\tau, q_N} |\delta(p, \sigma, \tau, q_N, N)|^2 \\
 &\quad + \sum_{\tau, q_L} |\delta(p, \sigma, \tau, q_L, L)|^2 \\
 &= \|\delta(p, \sigma)\|^2 = 1
 \end{aligned} \tag{4.42}$$

by the normalisation condition of Theorem 18, so the right hand sides of (4.35) are normalised.

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

| Track 1 and 4 rules | |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| $\begin{array}{ c c } \hline \bigcirc & \uparrow \\ \hline \bigcirc & \bigcirc \\ \hline \end{array} \rightarrow \begin{array}{ c c } \hline \uparrow & \bigcirc \\ \hline \bigcirc & \bigcirc \\ \hline \end{array}, \quad \begin{array}{ c c } \hline \bigcirc & \uparrow \\ \hline \bigcirc & \bigcirc \\ \hline \end{array} \rightarrow \begin{array}{ c c } \hline \uparrow & \bigcirc \\ \hline \bigcirc & \bigcirc \\ \hline \end{array}$ | $\begin{array}{ c c } \hline \bigcirc & \uparrow \\ \hline \bigcirc & r_q \\ \hline \end{array} \rightarrow \begin{array}{ c c } \hline \uparrow & \bigcirc \\ \hline r_q & \bigcirc \\ \hline \end{array}, \quad \begin{array}{ c c } \hline \bigcirc & \uparrow \\ \hline \bigcirc & p \\ \hline \end{array} \rightarrow \begin{array}{ c c } \hline \uparrow & \bigcirc \\ \hline \bigcirc & p \\ \hline \end{array}$ |
| Track 1, 4 and 5 rules | |
| $\left \begin{array}{ c c } \hline \bigcirc & \uparrow \\ \hline p & \bigcirc \\ \hline \sigma & \cdot \\ \hline \end{array} \right\rangle \rightarrow \sum_{\tau, q_R} \delta(p, \sigma, \tau, q_R, R) \left \begin{array}{ c c } \hline \uparrow & \bigcirc \\ \hline \bigcirc & q_R \\ \hline \tau & \cdot \\ \hline \end{array} \right\rangle$ $+ \sum_{\tau, q_N} \delta(p, \sigma, \tau, q_N, N) \left \begin{array}{ c c } \hline \uparrow & \bigcirc \\ \hline q_N & \bigcirc \\ \hline \tau & \cdot \\ \hline \end{array} \right\rangle$ $+ \sum_{\tau, q_L} \delta(p, \sigma, \tau, q_L, L) \left \begin{array}{ c c } \hline \uparrow & \bigcirc \\ \hline q'_L & \bigcirc \\ \hline \tau & \cdot \\ \hline \end{array} \right\rangle$ $\left \begin{array}{ c c } \hline \bigcirc & \uparrow \\ \hline \bigcirc & q'_L \\ \hline \cdot & \cdot \\ \hline \end{array} \right\rangle \rightarrow \left \begin{array}{ c c } \hline \uparrow & \bigcirc \\ \hline q_L & \bigcirc \\ \hline \cdot & \cdot \\ \hline \end{array} \right\rangle$ | |
| Track 1, 4 and 6 rules | |
| $\begin{array}{ c c } \hline \bigcirc & \uparrow \\ \hline q_f & \bigcirc \\ \hline \vdash & \cdot \\ \hline \end{array} \rightarrow \begin{array}{ c c } \hline \uparrow & \bigcirc \\ \hline p_\alpha & \bigcirc \\ \hline \vdash_{q_f} & \cdot \\ \hline \end{array}, \quad \begin{array}{ c c } \hline \bigcirc & \uparrow \\ \hline p & \bigcirc \\ \hline \sigma & \cdot \\ \hline \end{array} \rightarrow \begin{array}{ c c } \hline \uparrow & \bigcirc \\ \hline \bigcirc & p_R \\ \hline \tau & \cdot \\ \hline \end{array}, \quad \begin{array}{ c c } \hline \bigcirc & \uparrow \\ \hline p & \bigcirc \\ \hline \sigma & \cdot \\ \hline \end{array} \rightarrow \begin{array}{ c c } \hline \uparrow & \bigcirc \\ \hline p_N & \bigcirc \\ \hline \tau & \cdot \\ \hline \end{array}$ | $\begin{array}{ c c } \hline \bigcirc & \uparrow \\ \hline p & \bigcirc \\ \hline \cdot & \sigma \\ \hline \end{array} \rightarrow \begin{array}{ c c } \hline \uparrow & \bigcirc \\ \hline p'_L & \bigcirc \\ \hline \cdot & \tau \\ \hline \end{array}, \quad \begin{array}{ c c } \hline \bigcirc & \uparrow \\ \hline \bigcirc & p'_L \\ \hline \cdot & \cdot \\ \hline \end{array} \rightarrow \begin{array}{ c c } \hline \uparrow & \bigcirc \\ \hline p_L & \bigcirc \\ \hline \cdot & \cdot \\ \hline \end{array}$ |
| Track 1, 2 and 4 rules | |
| $\begin{array}{ c c } \hline \uparrow & \\ \hline \bigcirc & \triangleright \\ \hline q & \\ \hline \end{array} \rightarrow \begin{array}{ c c } \hline \uparrow & \\ \hline \bigcirc & \triangleright \\ \hline r_q & \\ \hline \end{array}$ | |
| Track 1, 2, 4 and 6 rules | |
| $\langle \leftarrow \begin{array}{ c c } \hline \uparrow & \\ \hline \neg p_\alpha & \\ \hline q'_L & \\ \hline \vdash & \\ \hline \end{array} \right\rangle \rightarrow \langle \leftarrow \begin{array}{ c c } \hline \uparrow & \\ \hline \neg p_\alpha & \\ \hline p_\alpha & \\ \hline \vdash_{q'_L} & \\ \hline \end{array} \right\rangle, \quad \langle \leftarrow \begin{array}{ c c } \hline \uparrow & \\ \hline \neg p_\alpha & \\ \hline r_q & \\ \hline \vdash & \\ \hline \end{array} \right\rangle \rightarrow \langle \leftarrow \begin{array}{ c c } \hline \uparrow & \\ \hline \neg p_\alpha & \\ \hline p_\alpha & \\ \hline \vdash_{r_q} & \\ \hline \end{array} \right\rangle$ | |
| Track 1, 2, 3 and 4 rules | |
| $\begin{array}{ c c } \hline \uparrow & \\ \hline p & \triangleright \\ \hline \sigma & \\ \hline q & \\ \hline \end{array} \rightarrow \begin{array}{ c c } \hline \uparrow & \\ \hline p_N & \triangleright \\ \hline \tau & \\ \hline r_q & \\ \hline \end{array}, \quad \begin{array}{ c c } \hline \uparrow & \\ \hline p'_R & \triangleright \\ \hline \cdot & \\ \hline q & \\ \hline \end{array} \rightarrow \begin{array}{ c c } \hline \uparrow & \\ \hline p_R & \triangleright \\ \hline \cdot & \\ \hline r_q & \\ \hline \end{array}$ | |

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

$$\begin{aligned} \langle \psi(p_1, \sigma_1) | \psi(p_2, \sigma_2) \rangle &= \sum_{\tau, q_R} \delta(p_1, \sigma_1, \tau, q_R, R)^* \delta(p_2, \sigma_2, \tau, q_R, R) \\ &\quad + \sum_{\tau, q_N} \delta(p_1, \sigma_1, \tau, q_N, N)^* \delta(p_2, \sigma_2, \tau, q_N, N) \\ &\quad + \sum_{\tau, q_L} \delta(p_1, \sigma_1, \tau, q_L, L)^* \delta(p_2, \sigma_2, \tau, q_L, L) \\ &= \langle \delta(p_1, \tau_1), \delta(p_2, \tau_2) \rangle = 0 \end{aligned} \tag{4.43}$$

by the orthogonality condition of Theorem 18, so the right hand sides of (4.35) are mutually orthogonal. Thus the QTM and time-wasting counter TM transition rules from (4.35) preserve orthonormality, as required.

4.6.3. 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:

$$\begin{array}{|c|c|} \hline \neg \leftarrow & \begin{array}{c} \rightarrow \\ 0 \end{array} \\ \hline & \neg \odot \\ \hline \end{array}, \quad \begin{array}{|c|c|} \hline \leftarrow & \begin{array}{c} \rightarrow \\ 0 \end{array} \\ \hline & \neg q_0 \\ \hline \end{array}. \tag{4.44}$$

The $\begin{array}{c} \rightarrow \\ 0 \end{array}$ 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 32.²⁰

$$\begin{array}{|c|} \hline \begin{array}{c} \leftarrow \\ 0 \end{array} \\ \hline \neg 1 \\ \hline \end{array}, \quad \begin{array}{|c|} \hline \begin{array}{c} \leftarrow \\ 0 \end{array} \\ \hline \neg \# \\ \hline \end{array}. \tag{4.45}$$

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

$$\begin{array}{|c|c|} \hline \neg \leftarrow & \begin{array}{c} \rightarrow \\ 0 \end{array} \\ \hline & \neg \# \\ \hline \end{array}, \quad \begin{array}{|c|c|} \hline \leftarrow & \begin{array}{c} \rightarrow \\ 0 \end{array} \\ \hline & \neg \vdash \\ \hline \end{array}. \tag{4.46}$$

4.7. 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’’ from Aharonov et al. [2, Lemma 4.2] (see also [32, Lemma 5.6]).²¹ 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.

First, we give some intuition behind the argument, before giving the rigorous proof. The starting point for the proof is that the structure of standard-form Hamiltonians (Definition 41) implies that the invariant subspaces of H decompose as a tensor product across the clock and quantum tracks, with the subspaces spanned by states connected by transition rules.

²⁰These additional blank symbols will be needed later to provide for the small space overhead of the QTM constructed in Theorem 10.

²¹Note that our Definition 35 of ‘‘legal’’ states follows that of [32], which differs from the definition in [2].

Furthermore, these invariant subspaces come in three types:

- (1) subspaces in which all states are illegal (Definition 35), i.e. are in the support of one or more penalty terms;
- (2) subspaces that contain legal and illegal states; and
- (3) subspaces that only contain legal states, of which there can be at most one, spanned by the states corresponding to a valid evolution of the QTM.

We can then analyse the spectrum in each type of subspace separately, to show that the ground state energy (minimum eigenvalue) is either > 0 or $= 0$, respectively, depending on whether the QTM halts or not.

The type 1 invariant subspaces trivially have positive energy from the penalty terms. Lemma 40 guarantees that any legal clock state in a type 2 subspace evolves to an illegal state, implying that all eigenstates in type 2 subspaces also have positive energy.

The analysis of the type 3 subspaces is reminiscent of Kitaev's original QMA-hardness proof for local Hamiltonians [47]. We first construct a unitary transformation W that brings the Hamiltonian H into the form $E_C \otimes \mathbb{1}_Q$, where E_C is a tridiagonal almost-Toeplitz matrix acting on the Hilbert space of the classical tracks.

This unitary is constructed by taking a product over the local unitaries corresponding to the local transition rules at each time step of the QTM. It is crucial that we have a local *unitary* at each time step, and not a partial isometry, otherwise this product would not necessarily be unitary. As seen in Section 3.1, the fact that the QTM is well-formed, normal form and unidirectional guarantees via Theorem 18 that the quantum transition rules define an isometry, thus can always be extended to a local unitary. Note it is *not* required (nor is it the case) that the *clock* transitions be extended to a full unitary; only the QTM transitions.

However, as discussed in Section 4.6.2, extending the QTM transition rules to a local unitary requires including the transitions out of the Halting state required of normal-form QTM's (Definition 15), as well as transitions out of configurations where the QTM has run out of tape space. The time-wasting construction of Section 4.6.2 ensures we still have a handle on the behaviour of the QTM after such a transition. In particular, it will never enter the Halting state if it has not already done so before running out of tape.

Having brought the Hamiltonian into the almost-diagonal form $H \simeq E \otimes \mathbb{1}$, we can analyse the spectrum of the type 3 invariant subspaces. These subspaces do not contain any illegal states, so are in the kernel of all penalty terms except possibly the Halting penalty. In the absence of the Halting penalty, the minimum eigenvalue in this subspace is therefore given by that of the tridiagonal almost-Toeplitz matrix E , which is 0 by standard results. The associated eigenvector corresponds to the uniform superposition over all legal clock states. Inverting the unitary transformation W , this is exactly the history state of Definition 31.

Therefore, considering all invariant subspaces, if the QTM never halts, the ground state is the 0-energy history state, which lives in a type 3 subspace. If the QTM *does* halt, so that the type 3 subspace is no longer in the kernel of the Halting penalty term, the minimum eigenvalue in type 3 subspaces can also be bounded away from 0 by standard arguments, thus the ground state energy is necessarily positive.

We now give the rigorous version of this argument.

Lemma 42 (Invariant subspaces). *Let H_{trans} and H_{pen} define a standard-form Hamiltonian as 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 $H = H_{pen} + H_{trans}$ 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 $\mathcal{K}_S \otimes \mathcal{H}_Q$ 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 specified 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 $\lambda_0(\mathcal{K}_S) = \Omega(1/|S|^3)$.*
- (3) *S contains only legal states, and $\lambda_0(\mathcal{K}_S) = 0$. The corresponding eigenspace is*

$$\ker(H_{\text{trans}} + H_{\text{pen}}) = \text{span} \left\{ \frac{1}{\sqrt{|S|}} \sum_{t=0}^{|S|-1} |t\rangle_C |\psi_t\rangle_Q \right\} \tag{4.47}$$

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 [47]). *Let $A, B \geq 0$ be positive semidefinite 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*

$$A + B \geq 2\mu \sin^2 \left(\frac{\theta}{2} \right) \tag{4.48}$$

where

$$\cos \theta := \max_{\substack{|\psi\rangle \in \ker A \\ |\varphi\rangle \in \ker B}} |\langle \psi | \varphi \rangle|. \tag{4.49}$$

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

$$H_{\text{trans}}|_{\mathcal{K}_S \otimes \mathcal{H}_Q} = \sum_{t=0}^T \frac{1}{2} \left(|t\rangle\langle t| \otimes \mathbb{1} + |t+1\rangle\langle t+1| \otimes \mathbb{1} - |t+1\rangle\langle t| \otimes U_t - |t\rangle\langle t+1| \otimes U_t^\dagger \right), \tag{4.50}$$

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 $H_t := |t\rangle\langle t| \otimes \mathbb{1} + |t+1\rangle\langle t+1| \otimes \mathbb{1} + |t\rangle\langle t| \otimes U_t + |t+1\rangle\langle t+1| \otimes U_t^\dagger$ from (4.50). 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

$$W_t H_t W_t^\dagger = (|t\rangle - |t+1\rangle) \langle t - \langle t+1| \otimes \mathbb{1} \geq 0. \tag{4.51}$$

Thus, whether the path in S forms a loop or not, we have²²

$$\begin{aligned} H_{\text{trans}}|_{\mathcal{K}_S \otimes \mathcal{H}_Q} &\geq \sum_{t=0}^{|S|-2} \frac{1}{2} \left(|t\rangle\langle t| \otimes \mathbb{1} + |t+1\rangle\langle t+1| \otimes \mathbb{1} - |t+1\rangle\langle t| \otimes U_t - |t\rangle\langle t+1| \otimes U_t^\dagger \right) \\ &=: H_{\text{path}}, \end{aligned} \tag{4.52}$$

with equality if the path is not a loop.

Defining the unitary

$$W = \sum_{t=0}^{|S|-2} |t\rangle\langle t| \otimes \prod_{i=0}^t U_i^\dagger + (|S| - 1) \langle (|S| - 1) \rangle \otimes \mathbb{1}, \tag{4.53}$$

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

$$E = \begin{pmatrix} \frac{1}{2} & -\frac{1}{2} & 0 & \dots & \dots & 0 & 0 \\ -\frac{1}{2} & 1 & -\frac{1}{2} & \ddots & & & 0 \\ 0 & -\frac{1}{2} & 1 & -\frac{1}{2} & & & \vdots \\ \vdots & \ddots & -\frac{1}{2} & 1 & \ddots & & \\ & & & \ddots & \ddots & \ddots & \vdots \\ \vdots & & & & \ddots & \ddots & 0 \\ 0 & & & & \ddots & \ddots & 1 & -\frac{1}{2} \\ 0 & 0 & \dots & \dots & 0 & -\frac{1}{2} & \frac{1}{2} \end{pmatrix}. \tag{4.54}$$

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 $|\phi_k\rangle \propto \sum_{j=0}^{|S|-1} \cos\left(q_k\left(j + \frac{1}{2}\right)\right) |j\rangle$ [47]. Thus

$$\ker H_{\text{path}} = \text{span} \left\{ W^\dagger \frac{1}{\sqrt{|S|}} \sum_{t=0}^{|S|-1} |t\rangle_C |\psi_0\rangle_Q \right\} = \text{span} \left\{ \frac{1}{\sqrt{|S|}} \sum_{t=0}^{|S|-1} |t\rangle_C |\psi_t\rangle_Q \right\} \tag{4.55}$$

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

$$\cos^2 \theta = \max_{|\psi\rangle} \langle \phi_0|_C \langle \psi|_Q W^\dagger \Pi_{\ker B} W |\phi_0\rangle_C |\psi\rangle_Q \leq 1 - 1/|S|, \tag{4.56}$$

²²This operator has the same form as Kitaev’s Hamiltonian, and the analysis from here on is similar to that in Kitaev et al. [47].

Table 5. Initialisation illegal pairs for Tracks 4–6.

| Illegal pairs | | |
|-----------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------|------------------------------------------------------------------------------------------------|
| Tracks 1 and 4 | Tracks 1 and 5 | Tracks 1 and 6 |
| $\begin{array}{ c c } \hline \neg \langle & \bar{0} \\ \hline & \neg \bar{0} \\ \hline \end{array}$ | $\begin{array}{ c } \hline \bar{0} \\ \hline \neg 1 \\ \hline \end{array}$ | $\begin{array}{ c c } \hline \neg \langle & \bar{0} \\ \hline & \neg \# \\ \hline \end{array}$ |
| $\begin{array}{ c c } \hline \langle & \bar{0} \\ \hline & \neg q_0 \\ \hline \end{array}$ | $\begin{array}{ c } \hline \bar{0} \\ \hline \neg \# \\ \hline \end{array}$ | $\begin{array}{ c c } \hline \langle & \bar{0} \\ \hline & \neg \vdash \\ \hline \end{array}$ |

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

$$(H_{\text{trans}} + H_{\text{pen}})|_{\mathcal{K} \otimes \mathcal{H}_Q} \geq H_{\text{path}} + H_{\text{pen}}|_{\mathcal{K} \otimes \mathcal{H}_Q} = \Omega(1/|S|^3), \tag{4.57}$$

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 (4.53) and (4.54). 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 4.6.2 (or similar).

We are finally in a position to prove Theorem 32, 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 4–6 Hilbert spaces, respectively, for a chain of length L , as specified in (4.6). 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 2 and 4 (omitting those marked in Table 2) 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 Tables 1 and 3, 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 5.

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 $(H_{\text{trans}}(L) + H_{\text{pen}}(L) + H_{\text{init}}(L))|_{\mathcal{S}_{br}}$ is the computational history state*

$$\frac{1}{\sqrt{T}} \sum_{t=0}^T |\phi_t\rangle_C |\psi_t\rangle_Q, \tag{4.58}$$

where $T = \Omega(L \zeta^L)$, $|\psi_0\rangle_Q \in \mathcal{H}_Q$ 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 $(H_{\text{trans}}(L) + H_{\text{pen}}(L))|_{\mathcal{S}_{br}}$. 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

$$\ker(H_{\text{trans}} + H_{\text{pen}})|_{S_{\text{br}}} = \text{span} \left\{ \frac{1}{\sqrt{|S|}} \sum_{t=0}^{|S|-1} |\phi_t\rangle_C |\psi_t\rangle_Q \right\}, \tag{4.59}$$

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 $\overset{\circlearrowright}{0}$ 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 $\overset{\circlearrowleft}{(0)}, \dots, \overset{\circlearrowleft}{(K)}$, 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 5 give an additional energy penalty to any Track 4–5 states that are not in the desired initial QTM configuration when the $\overset{\circlearrowleft}{0}$ 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 (4.58).

The initial state $|\psi_0\rangle_Q$ on Track 4 matches the regular expression $\langle \langle \circ \circ^* q \circ \circ^* \rangle \rangle$. 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 $\overset{\circlearrowleft}{1}$ sweep (all other unitaries applied to Tracks 4 and 5 during each sweep being identities).

Finally, the clock counter TM counts up to $\Omega(\zeta^L)$ 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 $T = \Omega(L \zeta^L)$ as claimed.²³ We are done. \square

The local Hilbert space dimension in (4.6) 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 ($\zeta = |\Sigma \times Q|$). Thus the local Hilbert space dimension d depends only on the alphabet size and number of internal states of the QTM, as claimed in part (i) of Theorem 32.

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, satisfying part (ii) of Theorem 32.

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 32. Moreover, since the history state encodes $\Omega(|\Sigma \times Q|^L)$ complete

²³This lower-bound is certainly not tight, as incrementing the base- ζ number on the tape takes more than one step of the TM, but it suffices for our purposes.

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 $\Omega(|\Sigma \times Q|^L)$ steps of the computation. If the QTM halts before this, it transitions by the q_f rule of (4.37) 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 32 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 32, this 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, at no point during the evolution is its head ever in a superposition of locations. $|\psi_t\rangle$ is therefore of the form:

$$|\psi_t\rangle = |\langle\otimes\rangle \left(\sum_q |q \textcircled{\circ} \dots \textcircled{\circ}\rangle_{\text{Track 4}} \otimes |\varphi_q\rangle_{\text{Track 5}} \otimes |_{\vdash} q \# \dots \#\rangle_{\text{Track 6}} \right) |\langle\otimes\rangle, \tag{4.60}$$

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 (4.37) therefore applies to all states in the superposition, and the next state must have the form:

$$|\psi_{t+1}\rangle = |\langle\otimes\rangle \left(\sum_q |p_\alpha \textcircled{\circ} \dots \textcircled{\circ}\rangle_{\text{Track 4}} \otimes |\varphi_q\rangle_{\text{Track 5}} \otimes |_{\vdash} q \# \dots \#\rangle_{\text{Track 6}} \right) |\langle\otimes\rangle. \tag{4.61}$$

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 32 follows.

This concludes the proof of Theorem 32, and this section.

5. 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.

5.1. Robinson’s tiling

We now describe Robinson’s quasi-periodic tiling. It was discovered by Robinson in 1971 [70] as a tool to simplify Berger’s proof of the undecidability of the tiling problem [15].

Robinson’s tiling is based on the five basic tiles showed in Figure 6 and all rotations and reflections thereof. 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. With this orientation convention, two consecutive crosses (meaning with no other cross in between them) in the same row or column are said to be facing each other, or to be back-to-back, as shown in Figure 7. An arm is said to point in the direction of its unique complete central arrow. Arrows that do not start or end in the mid-point of a square’s side are called *side arrows*.

For each one of the five basic tiles of Figure 6, 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:




Figure 6. The five basic tiles of Robinson's tiling (top), and a simplified schematic representation of these used in Figure 9 (bottom).




Figure 7. (a) Facing crosses, and (b) back-to-back crosses in Robinson's tiling.




Figure 8. Parity markings to force the position of the crosses.

colours must match between adjacent side arrows; 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 potentially in other positions too).

The latter can be achieved by adding parity markings to the above tiles: the additional arrows showed in Figure 8, 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. Note that the parity markings Figure 8 already form a closed set under rotation and reflection—(c) is the rotation of (b)—in contrast to the five basic tiles of Figure 6 which are extended to the full set by including all rotations and reflections.

If we now draw only the red coloured lines, one possible tiling of the plane looks like Figure 9, which has the crucial (for our purposes) quasi-periodic structure consisting of interlocking squares of increasing size.

Let us try to explain briefly (full details can be found in [70]) how Figure 9 emerges, and the remaining freedom it allows in constructing valid tilings of the plane. The fact that crosses must appear in alternate rows in alternate positions, together with the form of the five basic tiles, means that any given cross 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




Figure 9. A possible Robinson tiling of the plane.




Figure 10. 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.)

having the cross we started with as the top/right corner, as shown in Figure 10. Furthermore, the central tile must be a red cross, with the only freedom being the choice of direction it faces.

Once the direction of the central cross is fixed, the 7×7 square obtained by extending the 3×3 square in the direction this red cross faces is fixed as shown in Figure 10. The green cross in the central position is forced, with the only freedom again being the choice of direction it faces. Continuing this procedure gives a tiling similar to that shown in Figure 9. It consists of a quasi-periodic structure of squares, formed by a continuous line of a single colour going through all the tiles on the perimeter of the




Figure 11. Left: an example of a fault line, in which the patterns are shifted. Right: a schematic way, with the fault line in grey, to complete the figure on the right to a 7×7 square (compare with the example of the 7×7 square in Figure 10).




Figure 12. The basic tiles of the modified Robinson tiling.

square. These are called *borders* and have sizes 4^n (which—following Robinson’s terminology—means that the distance between two of the facing crosses delimiting the border is $4^n + 1$) repeated with period 2^{2n+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 the 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, the line separating the two half-planes is called a *fault* (see Figure 11). However, if a half-plane is further divided into two quarter-planes, no further shift is possible between these quarter-planes.

5.2. 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 5.1 and [70, § 8] and Figure 11).

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 12). 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 5.1. We refer to these as the *modified Robinson tiles*. (Similar modifications to the Robinson tile set appear elsewhere in the literature, e.g. in [58].)

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 (see Figure 11). 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.²⁴

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 8, 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 [70], 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 [70, §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 [70, §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 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

²⁴Our “rigidity” results for the Robinson tiling could certainly be strengthened, but they suffice for our purposes.

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 the 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 [70, §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 [70, § 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. [70, § 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. □

We refer to the top edge of a 2^n -border as a 2^n -segment.²⁵

Lemma 48 (Segment bound). *The number of 2^n -segments in a tiling of an $L \times H$ rectangle (width L , height H) using modified Robinson tiles is $\geq \lfloor H/2^{n+1} \rfloor (\lfloor L/2^{n+1} \rfloor - 1)$ and $\leq (\lfloor H/2^{n+1} \rfloor + 1) \lfloor L/2^{n+1} \rfloor$ for all n .*

Proof. Lemma 47 implies that the tiling must contain the usual periodic pattern of 2^n -borders, up to translation. The result then is a trivial consequence of the fact that borders repeat vertically and horizontally with period 2^{n+1} . □

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

²⁵It is on these 2^n -segments that the QTMs will ultimately “run” in the final construction.

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.²⁶ 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 the 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. □

6. Putting it all together

6.1. Undecidability of the g.s. energy density

To prove our first main result, Theorem 5, 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 $|\langle \rangle\rangle, |\rangle \rangle\rangle$. A *Gottesman-Irani Hamiltonian* is a 1D, translationally invariant, nearest-neighbour Hamiltonian $H_q(r)$ on a chain of length $r \geq 2$ described by Hilbert space $(\mathbb{C}^Q)^{\otimes r}$ 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) $[h_q, |\langle \rangle\rangle\langle \langle \rangle| \otimes |\langle \rangle\rangle\langle \langle \rangle|] = [h_q, |\langle \rangle\rangle\langle \langle \rangle| \otimes |\rangle \rangle\rangle\langle \rangle \rangle|] = [h_q, |\rangle \rangle\rangle\langle \rangle \rangle| \otimes |\langle \rangle\rangle\langle \langle \rangle|] = [h_q, |\rangle \rangle\rangle\langle \rangle \rangle| \otimes |\rangle \rangle\rangle\langle \rangle \rangle|] = 0$.
- (iii) $\lambda_0(r) := \lambda_0(H_q(r)|_{S_{br}}) < 1$, where S_{br} is the subspace of states with fixed boundary conditions $|\langle \rangle\rangle, |\rangle \rangle\rangle$ at the left and right ends of the chain, respectively.
- (iv) $\forall n \in \mathbb{N} : \lambda_0(4^n) \geq 0$ and $\sum_{n=1}^{\infty} \lambda_0(4^n) < 1/2$.
- (v) $\lambda_0(H_q(r)|_{S_{<}}) = \lambda_0(H_q(r)|_{S_{>}}) = 0$, where $S_{<}$ and $S_{>}$ are the subspaces of states with, respectively, a $|\langle \rangle\rangle$ at the left end of the chain or a $|\rangle \rangle\rangle$ 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 $h^{\text{row}}, h^{\text{col}} \in \mathcal{B}(\mathbb{C}^C \otimes \mathbb{C}^{Q+1}) \otimes \mathcal{B}(\mathbb{C}^C \otimes \mathbb{C}^{Q+1})$ 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.*

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

²⁶For $d_i > n/2$, this substantially over-counts the number of strips, but it is sufficient for our purposes.

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.

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:

$$h_{j,j+1}^{\text{col}} = h_c^{\text{col}} \otimes \mathbb{1}_{eq}^{(j)} \otimes \mathbb{1}_{eq}^{(j+1)} \tag{6.1a}$$

$$h_{i,i+1}^{\text{row}} = h_c^{\text{row}} \otimes \mathbb{1}_{eq}^{(i)} \otimes \mathbb{1}_{eq}^{(i+1)} \tag{6.1b}$$

$$+ \mathbb{1}_c^{(i)} \otimes \mathbb{1}_c^{(i+1)} \otimes h_q \tag{6.1c}$$

$$+ |L\rangle\langle L|_c^{(i)} \otimes (\mathbb{1}_{eq} - |\otimes\rangle\langle\otimes|)^{(i)} \otimes \mathbb{1}_{ceq}^{(i+1)} \tag{6.1d}$$

$$+ (\mathbb{1}_c - |L\rangle\langle L|_c)^{(i)} \otimes |\otimes\rangle\langle\otimes|^{(i)} \otimes \mathbb{1}_{ceq}^{(i+1)} \tag{6.1e}$$

$$+ \mathbb{1}_{ceq}^{(i)} \otimes |R\rangle\langle R|_c^{(i+1)} \otimes (\mathbb{1}_{eq} - |\oslash\rangle\langle\oslash|)^{(i+1)} \tag{6.1f}$$

$$+ \mathbb{1}_{ceq}^{(i)} \otimes (\mathbb{1}_c - |R\rangle\langle R|_c)^{(i+1)} \otimes |\oslash\rangle\langle\oslash|^{(i+1)} \tag{6.1g}$$

$$+ \mathbb{1}_c^{(i)} \otimes |0\rangle\langle 0|_e^{(i)} \otimes |R\rangle\langle R|_c^{(i+1)} \otimes \mathbb{1}_{eq}^{(i+1)} \tag{6.1h}$$

$$+ |L\rangle\langle L|_c^{(i)} \otimes \mathbb{1}_{eq}^{(i)} \otimes \mathbb{1}_c^{(i+1)} \otimes |0\rangle\langle 0|_e^{(i+1)} \tag{6.1i}$$

$$+ \mathbb{1}_c^{(i)} \otimes |0\rangle\langle 0|_e^{(i)} \otimes (\mathbb{1}_c - |L\rangle\langle L|_c)^{(i+1)} \otimes (\mathbb{1}_{ceq} - |0\rangle\langle 0|_e)^{(i+1)} \tag{6.1j}$$

$$+ (\mathbb{1}_c - |R\rangle\langle R|_c)^{(i)} \otimes (\mathbb{1}_{eq} - |0\rangle\langle 0|_e)^{(i)} \otimes \mathbb{1}_c^{(i+1)} \otimes |0\rangle\langle 0|_e^{(i+1)}, \tag{6.1k}$$

where $\mathbb{1}_c$, $\mathbb{1}_{eq}$ and $\mathbb{1}_{ceq}$ are the identity operators on the corresponding Hilbert spaces. Operators living on \mathcal{H}_q , or $\mathcal{H}_q \otimes \mathcal{H}_q$, such as h_q in (6.1c), are assumed to be extended in the trivial way (they are defined as zero in the extra sectors) to the larger space $\mathcal{H}_q \oplus \mathcal{H}_e$, or $(\mathcal{H}_q \oplus \mathcal{H}_e) \otimes (\mathcal{H}_q \oplus \mathcal{H}_e)$.

The Hamiltonian can be understood as follows. (6.1e) and (6.1d) force a $|\otimes\rangle$ in the q -layer whenever there is an $|L\rangle$ in the c -layer. (6.1f) and (6.1g) do the same with $|\oslash\rangle$ and $|R\rangle$. (6.1h) and (6.1i) force a non-blank to the left and right of an $|R\rangle$ or $|L\rangle$, respectively. Finally, (6.1j) and (6.1k) 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, (6.1b) 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, (6.1c) 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. ((6.1c) 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 $\ell \in \mathcal{L}_L \cup \mathcal{L}_R$, 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 (6.1a) and (6.1b), and $\sum_{\ell \in \mathcal{L}} \lambda_0(|\ell|)$ the contribution of (6.1c), the rest of the terms being 0 as $|T\rangle_c |\psi_0\rangle_q$ satisfies all constraints imposed by (6.1d) to (6.1k).

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,

$$\bigotimes_{i \in \Lambda} \mathcal{H}_q^{(i)} \simeq \bigoplus_{\sigma} \mathcal{H}_{\sigma}. \tag{6.2}$$

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 H_{σ} . Since all Hamiltonian terms (6.1d) to (6.1k) are ≥ 0 , and for all $\ell \in \mathcal{L}$ the (6.1c) contribution is $\geq \lambda_0(|\ell|)$, we easily get that $\langle T|_c \langle \psi_{\sigma}|_q H |T\rangle_c |\psi_{\sigma}\rangle_q \geq \langle T|H_c|T\rangle + \sum_{\ell \in \mathcal{L}} \lambda_0(|\ell|)$, 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 (6.1d) or (6.1f). If it is the position next to the right or left end, it picks up a contribution of 1 from from (6.1h) or (6.1i). Finally, if it is not one of the above cases, it acquires a contribution of 1 from (6.1k). In all cases, the contribution is $\geq \lambda_0(|\ell|)$ since $\lambda_0(|\ell|) < 1$ by assumption for a Gottesman-Irani Hamiltonian (see Definition 50). As all terms in the Hamiltonian apart from (6.1a) and (6.1b) are positive-semidefinite, and the contribution from (6.1a) and (6.1b) is $\langle T|H_c|T\rangle$, we have that $\langle T|_c \langle \psi_{\sigma}|_q H |T\rangle_c |\psi_{\sigma}\rangle_q \geq \langle T|H_c|T\rangle + \sum_{\ell \in \mathcal{L}} \lambda_0(|\ell|)$, which completes the proof of the Lemma. \square

With this, we can now prove the following:

Lemma 52 (Robinson + Gottesman-Irani Hamiltonian). *Let $h_c^{\text{row}}, h_c^{\text{col}} \in \mathcal{B}(\mathbb{C}^{\mathbb{C}} \otimes \mathbb{C}^{\mathbb{C}})$ 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 $h_q \in \mathcal{B}(\mathbb{C}^{\mathbb{Q}} \otimes \mathbb{C}^{\mathbb{Q}})$ 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 $h^{\text{row}}, h^{\text{col}} \in \mathcal{B}(\mathbb{C}^{\mathbb{C}^{(Q+1)}} \otimes \mathbb{C}^{\mathbb{C}^{(Q+1)}})$ such that, for any L, H , the ground state energy $\lambda_0(H^{\wedge(L \times H)})$ on a lattice of size $L \times H$ is contained in the interval

$$\left[\sum_{n=1}^{\lfloor \log_4(L/2) \rfloor} \left(\left\lfloor \frac{H}{2^{2n+1}} \right\rfloor \left(\left\lfloor \frac{L}{2^{2n+1}} \right\rfloor - 1 \right) \right) \lambda_0(4^n), \sum_{n=1}^{\lfloor \log_4(L/2) \rfloor} \left(\left(\left\lfloor \frac{H}{2^{2n+1}} \right\rfloor + 1 \right) \left\lfloor \frac{L}{2^{2n+1}} \right\rfloor \right) \lambda_0(4^n) \right] \tag{6.3}$$

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, $E(0 \text{ defects}) = \sum_{\ell \in \mathcal{L}} \lambda_0(|\ell|)$, is contained in the interval (6.3). (Recall that we only use the red 4^n -segments to define line-segments on which h_q acts non-trivially in the ground state.)

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

$$E(d \text{ defects}) = d + \sum_{\ell \in \mathcal{L}} \lambda_0(|\ell|) \geq d + \sum_{n=1}^{\lfloor \log_4(L/2) \rfloor} \left(\left\lfloor \frac{H}{2^{2n+1}} \right\rfloor \left(\left\lfloor \frac{L}{2^{2n+1}} \right\rfloor - 1 \right) - 2d \right) \lambda_0(4^n). \tag{6.4}$$

Since $\sum_{n=1}^{\infty} \lambda_0(4^n) < 1/2$ by assumption for a Gottesman-Irani Hamiltonian (see Definition 50), for all $d > 0$ this is in turn lower-bounded by

$$E(d \text{ defects}) \geq \sum_{n=1}^{\lfloor \log_4(L/2) \rfloor} \left(\left\lfloor \frac{H}{2^{2n+1}} \right\rfloor \left(\left\lfloor \frac{L}{2^{2n+1}} \right\rfloor - 1 \right) \right) \lambda_0(4^n). \tag{6.5}$$

The Lemma follows. □

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

Proposition 53 (Diverging g.s. energy). *Choose any $\beta \in \mathbb{Q}$, as small as desired. 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 $\leq 1/2$ and algebraic (hence computable) matrix entries, and there exist strictly positive (uncomputable) functions $\delta_1(n), \delta_2(n)$ such that either $\lambda_0(H_u^{\Lambda(L)}(n)) \leq -L \beta/2$ for all L , or $\lambda_0(H_u^{\Lambda}(n)) \geq L^2 \delta_2(n) - L \delta_1(n)$ for all $L \geq L_0(n)$ ($L_0(n)$ uncomputable), but determining which is undecidable.*

Moreover, the interactions can be taken to have the following form: $h_u^{(1)}(n) = -\beta(1 + \alpha_2(n))\mathbb{1}$ where $\alpha_2(n)$ is an algebraic number, $h_u^{col}(n)$ is $\{0, \beta\}$ -valued and independent of n and

$$h_u^{row}(n) = \beta \left(A + e^{i\pi\varphi} B + e^{i\pi 2^{-|\varphi|}} C \right) + h.c. \tag{6.6}$$

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

Proof. Let h_{q0} be the Hamiltonian obtained by applying Theorem 32 with $K = 3$ to the QTM from Theorem 10 with a proper reversible universal TM dovetailed after it. The Hamiltonian $h_q(n)$ in Lemma 52 will then be $h_q(n) = h_{q0}(n) + |\top\rangle\langle\top| \otimes \mathbb{1} + \mathbb{1} \otimes |\top\rangle\langle\top|$, where $|\top\rangle$ is the halting state of the universal TM. Clearly, h_q has the form given in part (vi) of Theorem 32. 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 32. In the initialisation sweep, the $\overset{\rightarrow}{\ominus}$ normally sweeps once from left-to-right along the chain, turns around at the $\overset{\circ}{\ominus}$ to become a $\overset{\leftarrow}{\ominus}$, 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 $\overset{\circ}{\ominus}$ is missing, then when the $\overset{\rightarrow}{\ominus}$ 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 $\overset{\circ}{\ominus}$ is missing, there is no further forward transition once the $\overset{\leftarrow}{\ominus}$ 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 32. For $r > 2$, we have

$$\begin{aligned} \lambda_0(r) &\leq \langle \psi | H_q(r) | \psi \rangle = \langle \psi | \left(\sum_i h_{q0}^{(i,i+1)}(n) + |\top\rangle\langle\top|_i \otimes \mathbb{1}_{i+1} + \mathbb{1}_i \otimes |\top\rangle\langle\top|_{i+1} \right) | \psi \rangle \\ &= \sum_{t=1}^T \frac{1}{T} \langle \psi_t | \left(\sum_i |\top\rangle\langle\top|_i \otimes \mathbb{1}_{i+1} + \mathbb{1}_i \otimes |\top\rangle\langle\top|_{i+1} \right) | \psi_t \rangle \leq O\left(\frac{1}{|\Sigma \times Q|^r}\right), \end{aligned} \tag{6.7}$$

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 $N(n) := \max\{\|\tilde{h}_u^{\text{row}}(n)\|, \|\tilde{h}_u^{\text{column}}(n)\|\}$, and fix any rational number $\beta \leq \frac{1}{N(n)}$ for all n . Such a β exists by the form of h_q guaranteed by part (vi) in Theorem 32 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 10 has sufficient tape and time to finish, and we can be sure that the reversible universal TM starts. Consider first the case in which the universal TM does *not* halt on input n . In that case, for all $r \geq |n| + 6$ we have that $\lambda_0(r) = 0$. Let $L_0(n)$ denote the minimal L such that the modified Robinson tiling of $\Lambda(L)$ necessarily contains a 4^m -segment of size $4^m \geq |n| + 6$. If we take $L \geq L_0(n)$, thanks to Lemma 52 we can bound the ground state energy for $H^{\Lambda(L)}(n)$:

$$\lambda_0(H^{\Lambda(L)}) \leq \beta \sum_{\substack{1 \leq r \leq |n|+6 \\ r=4^m, m \in \mathbb{N}}} \left\lfloor \frac{L}{2r} \right\rfloor \left(\left\lfloor \frac{L}{2r} \right\rfloor + 1 \right) \lambda_0(r) \tag{6.8a}$$

$$\begin{aligned} &= \beta L^2 \left[\sum_{\substack{1 \leq r \leq |n|+6 \\ r=4^m, m \in \mathbb{N}}} \frac{\lambda_0(r)}{4r^2} \right] + \beta L \left[\sum_{\substack{1 \leq r \leq |n|+6 \\ r=4^m, m \in \mathbb{N}}} \frac{\lambda_0(r)}{2r} \left(1 - 2 \text{frac} \left(\frac{L}{2r} \right) \right) \right] \\ &\quad - \beta \left[\sum_{\substack{1 \leq r \leq |n|+6 \\ r=4^m, m \in \mathbb{N}}} \lambda_0(r) \text{frac} \left(\frac{L}{2r} \right) \left(1 - \text{frac} \left(\frac{L}{2r} \right) \right) \right] \end{aligned} \tag{6.8b}$$

$$=: \beta \left(L^2 \alpha_2(n) + L \alpha_1(n, L) - \alpha_0(n, L) \right), \tag{6.8c}$$

where $\text{frac}(x) := x - \lfloor x \rfloor$ denotes the fractional part of x . Since the number of terms in the sum is bounded by $|n|$ and for any fixed r the quantity $\lambda_0(r)$ is an eigenvalue of a finite-dimensional matrix, $\alpha_2(n)$ is an algebraic number. Moreover, $\alpha_2(n) \leq \beta$.

We now shift the ground state energy by adding $\beta \mathbb{1} \otimes \mathbb{1}$ to h_{row} and adding a 1-body term $-\beta(1 + \alpha_2(n)) \mathbb{1}$. Overall, this adds the Hamiltonian $\sum_{\text{rows}} \sum_{c \in \text{cols}} (\beta \mathbb{1}_c \otimes \mathbb{1}_{c+1}) - \sum_{i \in \Lambda(L)} \beta(1 + \alpha_2(n)) \mathbb{1}_i = -\beta(L^2 \alpha_2(n) + L) \mathbb{1}_{\Lambda(L)}$. Note that this commutes with all the other terms in the Hamiltonian. After this shift, and using the fact that $\alpha_0(n, L) \geq 0$, the ground state energy in the non-halting case becomes

$$\lambda_0(H^{\Lambda(L)}) \leq \beta L (\alpha_1(n, L) - 1) \leq -\frac{\beta}{2} L, \tag{6.9}$$

since

$$\alpha_1(n, L) = \sum_{\substack{1 \leq r \leq |n|+6 \\ r=4^m, m \in \mathbb{N}}} \frac{\lambda_0(r)}{2r} \left(1 - 2 \operatorname{frac} \left(\frac{L}{2r} \right) \right) \leq \sum_{\substack{1 \leq r \leq |n|+6 \\ r=4^m, m \in \mathbb{N}}} \frac{\lambda_0(r)}{2r} \leq \frac{1}{2}. \tag{6.10}$$

Consider now the case in which the universal TM *does* halt on input n . Take r of the form 4^m large enough so that the TM has sufficient time and tape to halt. Let $r_1(n)$ denote the minimal such r . The computational history state encoding the evolution necessarily has support on $|\top\rangle$ and is the unique ground state of $h_{q0} \geq 0$. Thus $\lambda_0(r) > 0$ and by Lemma 52 the ground state energy $\lambda_0(H^{\wedge(L)})$ is, after the energy shift, lower-bounded by

$$\begin{aligned} \lambda_0(H^{\wedge(L)}) &\geq -\beta L^2 \alpha_2(n) - \beta L + \beta \sum_{\substack{1 \leq r \leq |n|+6 \\ r=4^m, m \in \mathbb{N}}} \left\lfloor \frac{L}{2r} \right\rfloor \left(\left\lfloor \frac{L}{2r} \right\rfloor - 1 \right) \lambda_0(r) \\ &\quad + \beta \sum_{\substack{r \geq r_1(n) \\ r=4^m, m \in \mathbb{N}}} \left\lfloor \frac{L}{2r} \right\rfloor \left(\left\lfloor \frac{L}{2r} \right\rfloor - 1 \right) \lambda_0(r) \end{aligned} \tag{6.11a}$$

$$\begin{aligned} &= \beta L^2 \left(\sum_{\substack{r \geq r_1(n) \\ r=4^m, m \in \mathbb{N}}} \frac{\lambda_0(r)}{4r^2} \right) - \beta L \left(1 + \sum_{\substack{|r| \leq n+6 \text{ or } r \geq r_1(n) \\ r=4^m, m \in \mathbb{N}}} \frac{\lambda_0(r)}{2r} \left(2 \operatorname{frac} \left(\frac{L}{2r} \right) + 1 \right) \right) \\ &\quad + \beta \left(\sum_{\substack{|r| \leq n+6 \text{ or } r \geq r_1(n) \\ r=4^m, m \in \mathbb{N}}} \lambda_0(r) \operatorname{frac} \left(\frac{L}{2r} \right) \left(\operatorname{frac} \left(\frac{L}{2r} \right) + 1 \right) \right) \end{aligned} \tag{6.11b}$$

$$=: \beta \left(L^2 \delta_2(n) - L \delta_1(n, L) + \delta_0(n, L) \right) \tag{6.11c}$$

$$\geq \beta \left(L^2 \delta_2(n) - L \delta_1(n) \right), \tag{6.11d}$$

where

$$\delta_1(n) := \left(1 + \sum_{\substack{|r| \leq n+6 \text{ or } r \geq r_1(n) \\ r=4^m, m \in \mathbb{N}}} \frac{3\lambda_0(r)}{2r} \right) \geq \delta_1(n, L), \tag{6.12}$$

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

Summarising, the ground state energy of $H^{\wedge(L)}$ in the non-halting case is bounded by

$$\lambda_0(H^{\wedge(L)}) \leq -L \frac{\beta}{2} \tag{6.13}$$

whereas in the halting case we have the bound

$$\lambda_0(H^{\wedge(L)}) \geq L^2 \delta_2(n) - L \delta_1(n). \tag{6.14}$$

The Proposition follows from undecidability of the Halting Problem. □

The following corollary follows immediately by letting $L(n)$ be the minimal L in Proposition 53 such that, in the halting case, $L^2\delta_2(n) - L\delta_1(n) \geq 1$. This will be useful shortly.

Corollary 54 (Undecidability of g.s. energy with promise). *There exists a family of interactions $h_u^{row}(n), h_u^{col}(n) \in \mathcal{B}(\mathbb{C}^U \otimes \mathbb{C}^U)$ and $h_u^{(1)}(n) \in \mathcal{B}(\mathbb{C}^U)$ with operator norm $\leq 1/2$ and algebraic (hence computable) matrix entries, and an (uncomputable) function $L(n)$, such that either $\lambda_0(H_u^{\Lambda(L)}(n)) \leq 0$ for all L , 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 5 (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 (hence computable) matrix entries, and local interaction strengths bounded by 1. Then determining whether $E_\rho = 0$ or $E_\rho > 0$ is an undecidable problem.*

6.2. 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. The overall intuition behind the proof is illustrated in Figure 13.

Recall that for each natural number n , we define $\varphi = \varphi(n)$ to be the rational number whose binary fraction expansion contains the digits of n in reverse order after the decimal.²⁷

Theorem 3 (restated). *For any given universal Turing Machine UTM, we can construct explicitly a dimension d , $d^2 \times d^2$ matrices A, A', B, C, D, D' , a $d \times d$ diagonal projector Π and a rational number β which can be as small as desired, with the following properties:*

- (i) A is diagonal with entries in \mathbb{Z} .
- (ii) A' is Hermitian with entries in $\mathbb{Z} + \frac{1}{\sqrt{2}}\mathbb{Z}$,
- (iii) B, C have integer entries,
- (iv) D is diagonal with entries in \mathbb{Z} ,
- (v) D' is hermitian with entries in \mathbb{Z} .

For each natural number n , define:

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

Then:

- (i) The local interaction strength is bounded by 1, i.e. $\max(\|h_1(n)\|, \|h_{row}(n)\|, \|h_{col}(n)\|) \leq 1$.
- (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.

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 $h_u^{(i,j)}$ be the two-body interactions obtained in Corollary 54 (keeping the one-body term $h_u^{(i)}(n)$ separate, for later convenience).

²⁷The reverse order being an unimportant artefact of the way we constructed the phase-estimation QTM in Section 3.

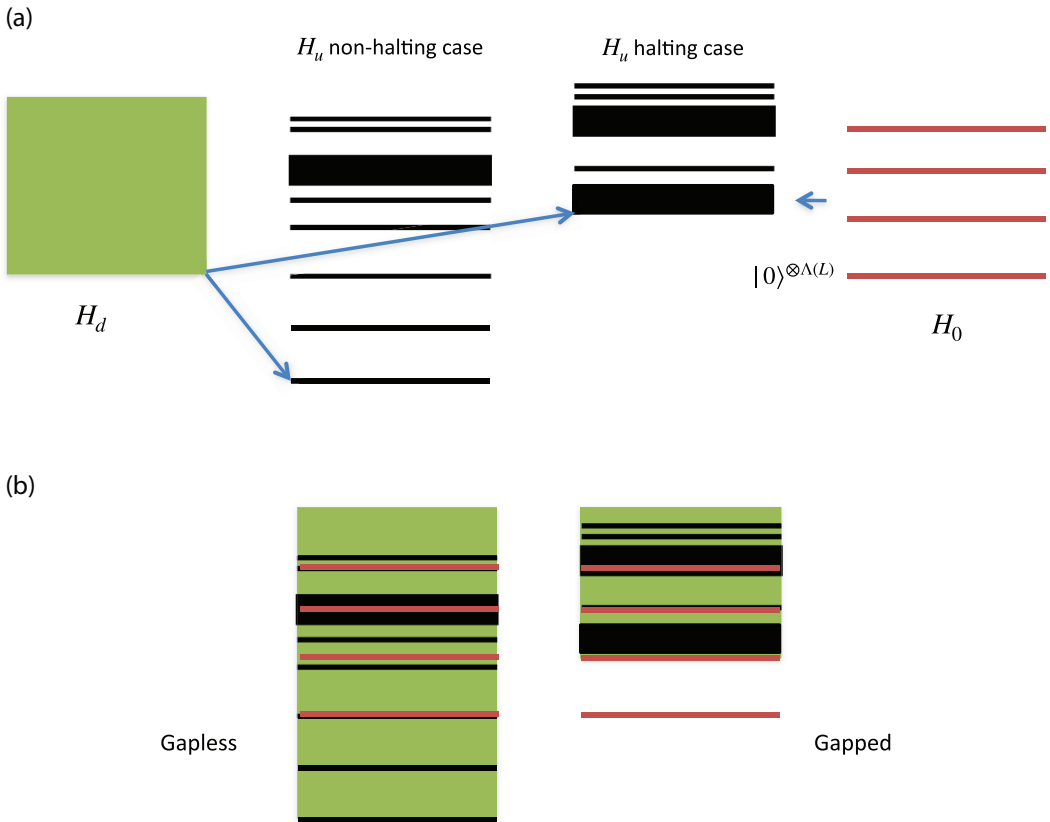


Figure 13. Starting with the Hamiltonian H_u constructed in Proposition 53, we consider two additional Hamiltonians, H_d and H_0 , with dense and trivial spectrum respectively. These are combined into a final Hamiltonian H such that the different spectra get combined as indicated by the arrows in the figure. This results in an overall Hamiltonian H with gapped or gapless behaviour, as shown in the bottom figure, depending on whether the TM encoded in Hamiltonian H_u halts or not.

Let h_d be the 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): $\lim_{L \rightarrow \infty} \text{spec } H_d^{\wedge(L)} \rightarrow [0, \infty)$. (For example [53], 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.) We normalise the interaction strength such that $\|h_d\| \leq \frac{1}{2}$.

Define the Hamiltonian $H(n)$ in terms of its two-body and one-body interactions $h(n)$ as follows (with $\alpha(n) = \alpha_2(n)$, $\Pi = \Pi_{ud}$):

$$h^{(i,j)}(n) := |0\rangle\langle 0|^{(i)} \otimes \Pi_{ud}^{(j)} + \Pi_{ud}^{(i)} \otimes |0\rangle\langle 0|^{(j)} \tag{6.15a}$$

$$+ \mathbb{1}_u^{(i)} \otimes \mathbb{1}_u^{(j)} \otimes h_d^{(i,j)}, \tag{6.15b}$$

$$+ h_u^{(i,j)}(n) \otimes \mathbb{1}_d^{(i)} \otimes \mathbb{1}_d^{(j)} \tag{6.15c}$$

$$h^{(i)}(n) := -\beta(1 + \alpha_2(n))\Pi_{ud}^{(i)}. \tag{6.15d}$$

Clearly, both have norm bounded by 1. Note that we can also rescale h_d so that $\|h_d\| \leq \beta$.

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 $H^{\wedge(L)} =: \tilde{H}_0 + \tilde{H}_d + \tilde{H}_u$, where $\tilde{H}_0, \tilde{H}_d, \tilde{H}_u$ are defined by taking the sum over sites separately for the expressions



Figure 14. Boundary tiles for periodic boundary condition construction.

in (6.15a), (6.15b), and (6.15c) + (6.15d), respectively. Note that the three terms commute with each other and

$$\text{spec } \tilde{H}_d = \text{spec } H_d, \quad \text{spec } \tilde{H}_u = \{0\} \cup \text{spec } H_u, \quad \text{spec } \tilde{H}_0 \subset \mathbb{Z}_{\geq 0} \tag{6.16}$$

Let us analyse the spectrum of $H^{\wedge(L)}$ in both the halting and non-halting cases. First consider the halting case, and take $L \geq L(n)$ as defined in Corollary 54. In that case, $\tilde{H}_d, \tilde{H}_u \geq 0$ and hence $H^{\wedge(L)} \geq \tilde{H}_0$. Since $|0\rangle^{\otimes \wedge(L)}$ is the unique ground state of \tilde{H}_0 , with energy 0, and is also a 0-energy state for $H^{\wedge(L)}$, we have that the spectral gap of $H^{\wedge(L)}$ is at least as large as the spectral gap of \tilde{H}_0 , which is 1.

In the non-halting case, it is clear from the structure of the Hamiltonians (6.15) that

$$\text{spec } \tilde{H}_u + \text{spec } \tilde{H}_d \subset \text{spec } H^{\wedge(L)} \subset \text{spec } \tilde{H}_u + \text{spec } \tilde{H}_d + \text{spec } \tilde{H}_0. \tag{6.17}$$

Since $\text{spec } \tilde{H}_0$ is contained in the set of non-negative integers, $\text{spec } \tilde{H}_d \subset [0, +\infty)$ and converges in the thermodynamic limit to $[0, +\infty)$, and $\lambda_0(\tilde{H}_u) \leq 0$ for all L by (6.16), statement (iii) in the Theorem follows. \square

6.3. 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.

To extend our result to periodic boundary conditions, we add to the modified Robinson tiles of Section 5 the three new tile types shown in Figure 14.

Let h_c be the standard tiling Hamiltonian of the modified Robinson tiles. We consider the Hamiltonian given by $3h_c$ plus the weighted terms given by the following tables (where $\boxed{\cdot}$ stands for any modified Robinson tile):

| | | | | | |
|-----------|-----------------|-----------------------------------------------------------------|---|-----------------------------------------------------------------|---|
| | | Right tile | | Tile above | |
| | | $\boxed{\cdot}$ $\boxed{\quad}$ $\boxed{\quad}$ $\boxed{\quad}$ | | $\boxed{\cdot}$ $\boxed{\quad}$ $\boxed{\quad}$ $\boxed{\quad}$ | |
| Left tile | $\boxed{\cdot}$ | 0 | 5 | 0 | 5 |
| | $\boxed{\quad}$ | 5 | 0 | 0 | 5 |
| | $\boxed{\quad}$ | 0 | 5 | 5 | 5 |
| | $\boxed{\quad}$ | 5 | 1 | 5 | 5 |
| | | | | Tile below | |
| | | | | $\boxed{\cdot}$ | 0 |
| | | | | $\boxed{\quad}$ | 0 |
| | | | | $\boxed{\quad}$ | 5 |
| | | | | $\boxed{\quad}$ | 5 |
| | | | | $\boxed{\quad}$ | 5 |

(6.18)

The resulting weighted tiling Hamiltonian \tilde{h}_c effectively turns the weighted tiling problem for an $L \times L$ square region with periodic boundary conditions (a torus), into the standard tiling problem for the modified Robinson tile set on a square region with open boundary conditions of size $(L - 1) \times (L - 1)$:

Lemma 55 (Periodic to open b.c.). *For any $L \in \mathbb{N}$, all minimum weight tilings of an $L \times L$ torus using the tile set described above consist of a single $\boxed{\quad}$ tile at an arbitrary location, a complete row of $\boxed{\quad}$ tiles extending horizontally from the $\boxed{\quad}$, a complete column of $\boxed{\quad}$ tiles extending vertically from the $\boxed{\quad}$, and a valid modified Robinson tiling of the remaining $(L - 1) \times (L - 1)$ region.*

Proof. The tilings described in the statement of 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. First, consider the case in which only modified Robinson tiles are used. By the aperiodicity of Robinson tiling, there are at least two mismatching tiles in an $L \times L$ torus, one in the vertical and one in the horizontal direction. Since we multiplied the original modified Robinson tiling Hamiltonian H_c by a factor of 3, each of these mismatches has weight 3 and we get a total weight of 6. 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 modified Robinson 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 modified Robinson 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 modified Robinson tiles is a square of size $(L - 1) \times (L - 1)$. It contributes weight 0 iff, in addition, they form a valid tiling. \square

Using this weighted tile set in Lemma 51, the proofs of Lemmas 51 and 52, Proposition 53, and Corollary 54 go through unchanged for square lattices with periodic boundary conditions. Undecidability of the ground state energy density (Theorem 5), and undecidability of the spectral gap (Theorem 3) follow easily.

Acknowledgments. 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.

The authors are very grateful to Norbert Schuch for pointing out that the aperiodic tiling itself can be exploited to both simplify and at the same time strengthen our original version of the periodic boundary condition results, and to James Watson for pointing out a number of errors in the draft manuscript.

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; also the organisers of the 2014 Seefeld Quantum Information Workshop, where another 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 some of this work was completed.

Conflict of Interest None.

Financial support. TSC is supported by the Royal Society. DPG acknowledges support from MINECO (grants MTM2014-54240-P, MTM2017-88385-P and ICMAT Severo Ochoa projects SEV-2015-0554 and CEX2019-000904-S), from Comunidad de Madrid (grants QUITEMAD, ref. S2013/ICE-2801 and S2018/TCS-4342), and the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement GAPS No 648913). DPG and MMW acknowledge support from the European CHIST-ERA project CQC (funded partially by MINECO grant PRI-PIMCHI-2011-1071). MMW acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC-2111 – 390814868.

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

- [1] I. Affleck, T. Kennedy, E. H. Lieb and H. Tasaki, ‘Valence bond ground states in isotropic quantum antiferromagnets’, *Commun. Math. Phys.* **115** (1988), 477–528.

- [2] D. Aharonov, D. Gottesman, S. Irani and J. Kempe, ‘The power of quantum systems on a line’, *Commun. Math. Phys.* **287** (2009), 41–65.
- [3] M. Aizenmann, ‘Open problems in mathematical physics’, 1998. <http://web.math.princeton.edu/~aizenman/OpenProblems.iamp>.
- [4] P. W. Anderson, ‘More is different’, *Science* **177** (1972), 393–396.
- [5] P. W. Anderson, ‘Resonating valence bonds: a new kind of insulator?’, *Mater. Res. Bull.* **8** (1973), 153–160.
- [6] I. Arad, A. Kitaev, Z. Landau and U. Vazirani, ‘An area law and sub-exponential algorithm for 1D systems’, Preprint, 2013, [arXiv:1301.1162](https://arxiv.org/abs/1301.1162).
- [7] I. Arad, Z. Landau, U. Vazirani and T. Vidick, ‘Rigorous RG algorithms and area laws for low energy eigenstates in 1D’, *Commun. Math. Phys.* **356** (2017), 65–105.
- [8] J. Bausch, T. Cubitt, A. Lucia and D. Perez-Garcia, ‘Undecidability of the spectral gap in one dimension’, *Phys. Rev. X* **10** (2020), 031038.
- [9] J. Bausch, T. Cubitt and J. D. Watson, ‘Uncomputability of phase diagrams’, *Nat. Commun.* **12** (2021), 452.
- [10] J. Bausch, T. S. Cubitt and M. Ozols, ‘The complexity of translationally-invariant spin chains with low local dimension’, *Ann. Henri Poincaré* **18** (2017), 3449–3513.
- [11] J. Bausch, T. S. Cubitt, A. Lucia, D. Perez-Garcia and M. M. Wolf, ‘Size-driven quantum phase transitions’, *Proc. Natl. Acad. Sci. (USA)* **115** (2018), 19–23.
- [12] A. Bendersky, G. Senno, G. de la Torre, S. Figueira and A. Acin, ‘Non-signaling deterministic models for non-local correlations have to be uncomputable’, *Phys. Rev. Lett.* **118** (2017), 130401.
- [13] C. H. Bennett, ‘Logical reversibility of computation’, *IBM J. Res. Develop.* **17** (1973), 525–532.
- [14] C. H. Bennett, ‘Undecidable dynamics’, *Nature* **346** (1990), 606–607.
- [15] R. Berger, ‘The undecidability of the domino problem’, *Mem. Amer. Math. Soc.* **66** (1966), 1–72.
- [16] E. Bernstein and U. Vazirani, ‘Quantum complexity theory’, *SIAM J. Comput.* **26**(5) (1997), 1411–1473.
- [17] L. Blum and S. Smale, ‘The Gödel incompleteness theorem and decidability over a ring’, in M. W. Hirsch, J. E. Marsden and M. Shub, eds., *From Topology to Computation: Proceedings of the Smalefest* (Springer, New York, 1993), 321–339.
- [18] S. Bravyi and D. Gosset, ‘Gapped and gapless phases of frustration-free spin-1/2 chains’, *J. Math. Phys.* **56** (2015), 061902.
- [19] S. Bravyi, M. Hastings and S. Michalakis, ‘Topological quantum order: stability under local perturbations’, *J. Math. Phys.* **51** (2010), 093512.
- [20] T. Cubitt, M. Gu, A. Perales, D. Perez-Garcia and M. M. Wolf, ‘Undecidability in physics: a quantum information perspective’, In preparation, 2022.
- [21] G. De las Cuevas, T. S. Cubitt, J. I. Cirac, M. M. Wolf and D. Pérez-García, ‘Fundamental limitations in the purifications of tensor networks’, *J. Math. Phys.* **57**(7) (2016), 071902.
- [22] T. S. Cubitt, D. Perez-Garcia and M. M. Wolf, ‘Undecidability of the spectral gap’, *Nature* **528** (2015), 207–211.
- [23] T. S. Cubitt, D. Perez-Garcia and M. M. Wolf, ‘Comment on “On the uncomputability of the spectral gap”’, Preprint, 2016, [arXiv:1603.00825](https://arxiv.org/abs/1603.00825).
- [24] W. De Roeck and M. Salmhofer, ‘Persistence of exponential decay and spectral gaps for interacting fermions’, *Commun. Math. Phys.* **365**(2) (2019), 773–796.
- [25] E. Domany and W. Kinzel, ‘Equivalence of cellular automata to Ising models and directed percolation’, *Phys. Rev. Lett.* **53**(4) (1984), 311.
- [26] J. Eisert, M. Cramer and M. B. Plenio, ‘Colloquium: area laws for the entanglement entropy’, *Rev. Modern Phys.* **82**(1) (2010), 277.
- [27] J. Eisert, M. P. Müller and C. Gogolin, ‘Quantum measurement occurrence is undecidable’, *Phys. Rev. Lett.* **108** (2012), 260501, 6.
- [28] D. Elkouss and D. Pérez-García, ‘Memory effects can make the transmission capability of a communication channel uncomputable’, *Nat. Commun.* **9** (2018), 1149.
- [29] C. Fernández-González, N. Schuch, M. M. Wolf, J. I. Cirac and D. Pérez-García, ‘Frustration free gapless Hamiltonians for matrix product states’, *Commun. Math. Phys.* **333** (2015), 299–333.
- [30] R. Feynman, ‘Quantum mechanical computers’, *Optics News* **11** (1985), 11–20.
- [31] E. Fredkin and T. Toffoli, ‘Conservative logic’, *Int. J. Theoret. Phys.* **21**(3) (1982), 219–253.
- [32] D. Gottesman and S. Irani, ‘The quantum and classical complexity of translationally invariant tiling and Hamiltonian problems’, in *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science, FOCS’09* (IEEE Computer Society, Washington, DC, 2009), 95–104.
- [33] B. Grünbaum and G. C. Shephard, *Tilings and Patterns* (W.H. Freeman and Company, New York, 1987).
- [34] M. Gu, C. Weedbrook, A. Perales and M. A. Nielsen, ‘More really is different’, *Physica D* **238** (2009) 835–839.
- [35] J. Haegeman, S. Michalakis, B. Nachtergaele, T. J. Osborne, N. Schuch and F. Verstraete, ‘Elementary excitations in gapped quantum spin systems’, *Phys. Rev. Lett.* **111**(8) (2013), 080401.
- [36] F. D. M. Haldane, ‘Nonlinear field theory of large-spin Heisenberg antiferromagnets: semiclassically quantized solitons of the one-dimensional easy-axis Néel state’, *Phys. Rev. Lett.* **50** (1983), 1153–1156.

- [37] T.-H. Han, J. S. Helton, S. Chu, D. G. Nocera, J. A. Rodriguez-Rivera, C. Broholm and Y. S. Lee, ‘Fractionalized excitations in the spin-liquid state of a Kagome-lattice antiferromagnet’, *Nature* **492** (2012), 406–410.
- [38] M. B. Hastings, ‘Lieb-Schultz-Mattis in higher dimensions’, *Phys. Rev. B* **69** (2004), 104431.
- [39] M. B. Hastings, ‘An area law for one-dimensional quantum systems’, *J. Stat. Mech.: Theory and Experiment* 2007(8) (2007), P08024.
- [40] M. B. Hastings, ‘Entropy and entanglement in quantum ground states’, *Phys. Rev. B* **76**(3) (2007), 035114.
- [41] M. B. Hastings, ‘The stability of free fermi Hamiltonians’, *J. Math. Phys.* **60** (2019), 042201.
- [42] M.B. Hastings and T. Koma, ‘Spectral gap and exponential decay of correlations’, *Commun. Math. Phys.*, **265** (2006), 781–804.
- [43] D. R. Hofstadter, ‘Energy levels and wave functions of Bloch electrons in rational and irrational magnetic fields’, *Phys. Rev. B* **14**(9) (1976), 2239–2249.
- [44] Y. Iqbal, D. Poilblanc and F. Becca, ‘Vanishing spin gap in a competing spin-liquid phase in the Kagome Heisenberg antiferromagnet’, *Phys. Rev. B* **89**(2) (2014), 020407.
- [45] I. Kanter, ‘Undecidability principle and the uncertainty principle even for classical systems’, *Phys. Rev. Lett.* **64** (1990), 332–335.
- [46] J. Kempe, A. Kitaev and O. Regev, ‘The complexity of the local Hamiltonian problem’, *SIAM J. Comput.* **35** (2006), 1070–1097.
- [47] A. Y. Kitaev, A. H. Shen and M. N. Vyalyi, ‘*Classical and Quantum Computation, Graduate Studies in Mathematics, Vol. 47* (American Mathematical Society, Providence, RI, 2002).
- [48] M. Kliesch, D. Gross and J. Eisert, ‘Matrix-product operators and states: NP-hardness and undecidability’, *Phys. Rev. Lett.* **113** (2014), 160503.
- [49] A. Komar, ‘Undecidability of macroscopically distinguishable states in quantum field theory’, *Phys. Rev.* **133** (1964), B542–B544.
- [50] D. Kozen, *Automata and Computability* (Springer Science & Business Media, New York, 1997).
- [51] Z. Landau, U. Vazirani and T. Vidick, ‘A polynomial-time algorithm for the ground state of one-dimensional gapped local Hamiltonians’, *Nat. Phys.* **11** (2013), 566–569.
- [52] M. Lemm, A. W. Sandvik and L. Wang, ‘Existence of a spectral gap in the Affleck–Kennedy–Lieb–Tasaki model on the hexagonal lattice’, *Phys. Rev. Lett.* **124**(17) (2020), 177204.
- [53] E. H. Lieb, T. D. Schultz and D. C. Mattis, ‘Two soluble models of an antiferromagnetic chain’, *Ann. Phys.* **16** (1961), 407–466.
- [54] S. Lloyd, ‘Quantum-mechanical computers and uncomputability’, *Phys. Rev. Lett.* **71**(6) (1993), 943.
- [55] S. Lloyd, ‘Necessary and sufficient conditions for quantum computation’, *J. Modern Opt.* **41**(12) (1994), 2503–2520.
- [56] S. Lloyd, ‘On the uncomputability of the spectral gap’, Preprint, 2016, [arXiv:1602.05924](https://arxiv.org/abs/1602.05924).
- [57] S. Michalakis and J. Pytel, ‘Stability of frustration-free Hamiltonians’, *Commun. Math. Phys.* **322**(2) (2013), 277–302.
- [58] J. Miękisz, ‘Stable quasicrystalline ground states’, *J. Stat. Phys.* **88** (1997), 691–711.
- [59] C. Moore, ‘Unpredictability and undecidability in dynamical systems’, *Phys. Rev. Lett.* **64** (1990), 2354–2357.
- [60] J. Morton and J. Biamonte, ‘Undecidability in tensor network states’, *Phys. Rev. A* **86** (2012), 030301.
- [61] B. Nachtergaele and R. Sims, ‘Lieb-Robinson bounds and the exponential clustering theorem’, *Commun. Math. Phys.* **265**(1) (2006), 119–130.
- [62] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2000).
- [63] R. Oliveira and B. Terhal, ‘The complexity of quantum spin systems on a two-dimensional square lattice’, *Quantum Inf. Comput.* **8** (2008), 0900–0924.
- [64] S. Omohundro, ‘Modelling cellular automata with partial differential equations’, *Physica D: Nonlinear Phenomena* **10**(1–2) (1984), 128–134.
- [65] R. Orús, ‘A practical introduction to tensor networks: matrix product states and projected entangled pair states’, *Ann. Phys.* **349** (2014), 117–158.
- [66] D. Osadchy and J. E. Avron, ‘Hofstadter butterfly as quantum phase diagram’, *J. Math. Phys.* **42**(12) (2001), 5665–5671.
- [67] N. Pomata and T.-C. Wei, ‘Demonstrating the AKLT spectral gap on 2D degree-3 lattices’, *Phys. Rev. Lett.* **124** (2020), 177203.
- [68] B. Poonen, ‘Undecidable problems: a sampler’, in J. Kennedy, ed., *Interpreting Gödel: Critical Essays* (Cambridge University Press, 2014), 211–241.
- [69] M. B. Pour-El and I. Richards, ‘The wave equation with computable initial data such that its unique solution is not computable’, *Adv. Math.* **39**(3) (1981), 215–239.
- [70] R. Robinson, ‘Undecidability and nonperiodicity for tilings of the plane’, *Invent. Math.* **12** (1971), 177–209.
- [71] W. Slofstra, ‘Tsirelson’s problem and an embedding theorem for groups arising from non-local games’, Preprint, 2016, *J. Amer. Math. Soc.* **33** (2020), 1–56.
- [72] A. Turing, ‘On computable numbers, with an application to the entscheidungsproblem’, *Proc. London Math. Soc.* **42** (1936), 230–265.

- [73] M. Van den Nest and H. J. Briegel, 'Measurement-based quantum computation and undecidable logic', *Found. Phys.* **38**(5) (2008), 448–457.
- [74] H. Wang, 'Proving theorems by pattern recognition', *Bell System Tech. J.* **40** (1961), 1–41.
- [75] M. M. Wolf, T. S. Cubitt and D. Perez-Garcia, 'Are problems in quantum information theory (un)decidable?', Preprint, 2011, [arXiv:1111.5425](https://arxiv.org/abs/1111.5425).
- [76] S. Yan, D. A. Huse and S. R. White, 'Spin-liquid ground state of the S=1/2 Kagome Heisenberg antiferromagnet', *Science* **332** (2011), 1173–1176.
- [77] A. Yedidia and A. Aronson, 'A relatively small Turing machine whose behavior is independent of set theory', *Complex Systems* **25** (2016), 297–327.
- [78] A. Yu, *Kitaev* (1995). <https://arxiv.org/abs/quant-ph/9511026>.