Hamiltonian Complexity in Many-Body Quantum Physics

James David Watson

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

of

University College London.

Department of Computer Science

University College London

November 19, 2022

I, James David Watson, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the work.

Abstract

The development of quantum computers has promised to greatly improve our understanding of quantum many-body physics. However, many physical systems display complex and unpredictable behaviour which is not amenable to analytic or even computational solutions. This thesis aims to further our understanding of what properties of physical systems a quantum computer is capable of determining, and simultaneously explore the behaviour of exotic quantum many-body systems.

First, we analyse the task of determining the phase diagram of a quantum material, and thereby charting its properties as a function of some externally controlled parameter. In the general case we find that determining the phase diagram to be uncomputable, and in special cases show it is $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$-complete. Beyond this, we examine how a common method for determining quantum phase transitions — the Renormalisation Group (RG) — fails when applied to a set of Hamiltonians with uncomputable properties. We show that for such Hamiltonians (a) there is a well-defined RG procedure, but this procedure must fail to predict the uncomputable properties (b) this failure of the RG procedure demonstrates previously unseen and novel behaviour.

We also formalise in terms of a promise problem, the question of computing the ground state energy per particle of a model in the limit of an infinitely large system, and show that approximating this quantity is likely intractable. In doing this we develop a new kind of complexity question concerned with determining the precision to which a single number can be determined.

Finally we consider the problem of measuring local observables in the low energy subspace of systems — an important problem for experimentalists and theorists alike. We prove that if a certain kind of construction exists for a class of Hamiltonians, , the results about hardness of determining the ground state energy directly implies hardness results for measuring observables at low energies.

Impact Statement

It is widely believed that the advent of quantum computers will enable researchers to tackle computational problems previously thought to be intractable. Tasks concerning modelling physical systems, such as those relevant to material science and chemistry, are thought to be more amenable to quantum computers relative to their non-quantum counter-partners.

This thesis analyses to what extent this is true, and furthers our characterisation of what problems quantum computers will be able to efficiently solve. There are many groups working not only in academia, but also in industry attempting to use quantum computers for these purposes. As such this work will allow researchers to focus their efforts on finding algorithms for physical systems which are more likely to demonstrate a comparative speed-up for quantum computers. Furthermore, the understanding of what makes a physical system hard to "solve" helps inform algorithm design when solving other systems.

Aside from practical relevance to algorithms, the research contributes to a broader understanding of quantum complexity theory and the power of logical machines, and how quantum mechanics can be utilised by such machines. In particular here we study a unique form of computation problem that may lead to important insights into complexity theory.

Our results may have additional impact in the understanding of complex behaviour such as phase transitions in many-body quantum systems from a condensed matter perspective. In particular, we expect our methods to be of interest to scientists working to demonstrate rigorous results about complex physical systems such as spin glasses. In addition, some of the results here put hard limits on how effectively algorithms can be used to analyse many-body quantum systems — such as limiting the effectiveness of extrapolating numerical information for small systems to larger systems.

The results in this thesis have been distributed via academic talks at conferences, universities, and workshops as well as on social media.

Acknowledgements

I would like to give an enormous thank you to Toby Cubitt, who in addition to having two young children, running a start-up, and being a senior academic still managed to find time to be a fantastic PhD supervisor. His insight, support, and guidance has been without equal.

The CDT programme at UCL has enormously enhanced my PhD experience. No doubt this is due in part to the firm and competent leadership of Dan Browne and Paul Warburton, as well as Lopa Murgai, whose support and unique sense of humour has been a highlight. I would also like to thank all the members of my CDT cohort, whose company I have treasured throughout my time at UCL, and who I have subjected to my own unpolished sense of humour. These include Diego Aparicio, Gioele Consani, Abdulah Fawaz, Matt Flinders, Edward Grant, Tom Hird, Gareth Jones, Tamara Kohler, Conor McKeever, TK Le, Jamie Potter, James Seddon, and Maria Stasinou. A particular thanks goes to Tamara Kohler, who has gracefully shared an office with me for 3 years and provided many interesting discussions, as well as for her assurances that I am not a complete moron.

I am deeply indebted to the other members of UCL's quantum physics groups, including Emilio Onorati, Lluis Masanes, Tom Farshi, Joel Klassen, Charlie Derby, Laura Clinton, Jonathan Oppenheim, Zach Weller-Davies, and many others with whom I have shared valuable discussions and ideas. I also owe a great deal to George Pooley, Matthew Jilani, Joshua Sant, Jordan Harvey, Graeme Lye, and Gabriel Mead for keeping me in good spirits throughout my PhD. I would like to give a particular mention to Joe Sullivan and Oskari Timgren, who although not in the UK, have entertained and advised me from afar.

Throughout my PhD I have had the great joy of collaborating with Johannes Bausch, Emilio Onorati, and Sev Gharibian. The fountain of wisdom that Johannes has provided over the last few years has been greatly appreciated. No doubt my work would be much diminished without his keen insight. Emilio's deep knowledge of quantum information and Italian food has also been greatly appreciated, and enjoyed.

Sev's unparalleled understanding of quantum complexity theory has not only been of great interest, but has corrected many grave errors that I would have happily ignored.

I thank my parents Diane and Stephen for their tireless support without which I could not have done without, particularly during the COVID-19 pandemic. Also to my brother Charles for his unique motivation and encouragement as well as his quality meme content. I am enormously thankful to have had the support and company of my girlfriend, Flora, throughout my PhD.

Finally, I would like to thank my many school teachers for the their enthusiastic work and encouragement in going into science. These include M. Walker, K. Bridge, J. Sapsford, K. Parsons, I. Devereux, I. Nash, N. Rolands, R. Baker-Glenn, J. Lidiard, C. Dawes, and A. Smith. In particular, Mary Walker whose teaching ultimately inspired me to take up a career in physics.

Publications and Preprints

The work presented in this thesis contains results from the following papers/ preprints:

- [BCW21] J. Bausch, T. S. Cubitt, and J. D. Watson *"Uncomputability of phase diagrams"*, Nature Communications 12, 452 (2021)
- [WBG20] J. D. Watson, J. Bausch, and S. Gharibian, *"The Complexity of Translationally Invariant Problems beyond Ground State Energies"*, arXiv:2012.12717
- [WOC21] J. D. Watson, E. Onorati, T. S. Cubitt, *"Uncomputably Complex Renormalisation Group Flows"*, arXiv: 2102.05145.
- [WB21] J. D. Watson and J. Bausch, *"The Complexity of Approximating Critical Points of Quantum Phase Transitions"*, arXiv:2105.13350
- [WC21] J. D Watson and T. S. Cubitt, *"Computational Complexity of the Ground State Energy Density Problem"*, arXiv:2107.05060; and a condensed version is published in [WC22] *"Computational Complexity of the Ground State Energy Density Problem"* STOC 2022: Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing, June 2022 Pages 764–775.

We note for the reader that although results from all of these papers have been included, not all of the results from each of these papers have been included. Some results have been removed where they had significant input from coauthors or where they were too long to entirely fit in this thesis.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Quantum mechanics is a physical theory describing the microscopic physics of bodies, such as atoms and molecules, whereas classical physics describes physics at a larger scale, including almost all physical processes that a person might typically experience. Compared to classical physics, quantum mechanics displays a huge array of apparently mysterious and unfamiliar behaviour such as entanglement, de-localisation, and interference.

Much of the physics of the 20^{th} and 21^{st} century has been devoted to understanding the behaviour of systems described by quantum mechanics, in particular systems comprised of multiple interacting particles — known as many-body physics. That is, furnished with a description of a set of particles and the interactions between them, can we determine a given property e.g. the electrical conductivity? Intuitively, it might seem that, provided the interactions between particles are not too complicated, the problem is straightforward. Regrettably, this is often not the case. Although quantum mechanics may give a framework for studying microscopic systems, determining physical properties from first principles is often not an easy task and in most cases simple, closed-form solutions simply do not exist. This is succinctly summed up by the Philip W. Anderson quote "More is different", which is to say that large collections of particles, even if they have simple interactions between them, can display enormously complex behaviour — far more complicated than one might reasonably expect [And08]. In general, the behaviour of these systems cannot be straightforwardly extracted from the underlying interactions and requires

a completely different toolkit to analyse them — a reductionist viewpoint can only have limited success[1].

Some of those systems which have been most frustrating to physicists are so called "condensed matter systems", which are essentially ensembles of particles which can interact with each other in all manner of ways, such as in solid materials. Indeed, many unusual and perplexing properties of these many-body systems have been discovered and explained over course the past century which require the description of quantum mechanics including the low-temperature superconductivity and the quantum Hall effect. Many more have been observed but their mechanism remain unaccounted for, such as high temperature superconductivity [GK82].

Of course, understanding the properties of microscopic quantum systems is important, not just from a theoretical perspective but also for practical uses such as designing new materials, new drugs, and studying chemical reactions. Often we find these problems are simply intractable on classical computers. That is solving for these properties — even if great simplifications are made — would take too long to be practical.

Recently the development of quantum computers — computers which exploit the properties of quantum mechanics to improve information processing — promise to help tackle this problem. Quantum computers are capable of implementing a range of algorithms which (are thought to) solve problems far more efficiently than classical computers, especially problems related to physical systems such as simulating the time evolution of molecules. But, even for quantum computers, some physical systems and their properties are believed to remain highly non-trivial to determine.

In this thesis we are concerned with developing and improving the understanding of which physical systems and their properties can be determined efficiently, inefficiently, or not at all using either classical or quantum computers; a field called Hamiltonian complexity.

Hamiltonian complexity borrows many techniques and concepts from computa-

[1] In some sense this is obvious. Humans are capable of performing incredibly complex processes, but are built from elementary particles, the laws governing which can be written down on a piece of paper.

tional complexity theory which is a field devoted to characterising how efficiently computational problems can be solved. In particular, complexity theory typically asks whether an efficient algorithm exists for a particular problem. For example an efficient algorithm exists for adding two numbers, but is not thought to exist for calculating matrix permanents. Perhaps surprisingly, these tools from complexity theory are able not only to classify which quantum mechanical systems may or may not be solvable efficiently using algorithms (or a variety of other methods), but also can give us insight into the physical properties of such systems. This classification approach stands in contrast to how much of physics is typically done — ideally we would like a simple, closed-form solution, or perhaps a convergent perturbative series which solves our problem. Hamiltonian complex is an admission that, in general, this is not possible and many-body quantum physics is much more complicated and messy than we would like. While Einstein, Dirac and other giants of physics have produced beautiful and succinct theories describing nature, Hamiltonian complexity gently tells us that much of physics won't be like that. Instead, we'll have to wrestle with it numerically, and even then there will be many systems we cannot deal with.

Although Hamiltonian complexity has resulted in a rich classification of systems and properties which are easy or hard to determine, there exists little understanding or insight into where this boundary between simple and complex systems occurs. It remains to be seen whether it is possible to tease out a boundary where physical systems switch from simple to complex based only on the parameters describing the system.

Furthermore, there is a noticeable dearth of results relating to systems in the thermodynamic limit (i.e. systems that are infinitely large). This is an important limit for several reasons. Perhaps most importantly, the thermodynamic limit gives the "bulk" behaviour of materials and removes any finite-sized effects. Often this is the regime condensed matter physicists wish to study materials — reflecting the fact we usually deal with macroscopic systems of $> 10^{23}$ particles where finite sized effects are assumed to be irrelevant. Moreover, important phenomena such as quantum phase transitions only technically occur in this limit. In this thesis we seek to prove

results about such systems, thus bringing Hamiltonian complexity closer to problems studied by condensed matter physicists.

1.1 Structure of this Thesis

Chapter 2 gives preliminaries on notation, as well a more in-depth introduction to quantum physics, complexity theory and Hamiltonian complexity. It also introduces some of the relevant previous literature.

Chapter 3 explores the difficulty of computing the phase diagram at zero temperature for a general Hamiltonian on a 2D lattice. In particular, we show that the problem is uncomputable for a general Hamiltonian. That is, there exists no general algorithm, no matter how inefficient, which determines the phase diagram of a Hamiltonian. The results in this chapter are published in [BCW21].

Chapter 4 looks a new computational problem which characterises how difficult it is to estimate the energy per particle in a 2D infinite lattice. Not only is this one of the first *complexity* results in the thermodynamic limit, it introduces a new type of complexity theory problem in of which we study a fixed property of a fixed Hamiltonian and vary the precision. These results were first shown in [WC21].

Chapter 5 looks a set techniques for computing phases materials called the "Renormalisation Group" methods. We take a Hamiltonian which has uncomputable properties (the Hamiltonian from [CPGW15a]) and explicitly construct a renormalisation group flow for it. In doing so, we unveil an entirely new set of renormalisation group behaviour, and demonstrate why such techniques must fail for Hamiltonians with uncomputable properties. In particular, it shows the reason that the RG scheme must fail to "solve" the system is not because it impossible to construct an RG scheme, but instead because the RG scheme demonstrates pathological behaviour. These results were first shown in [WOC21].

Chapter 6 examines the problem of estimating the expectation value of local measurements on low temperature quantum systems. This is important as the only access an experimentalist has to a low temperature quantum system is through making local measurements. In particular, we prove a lifting lemma showing that if

determining the ground state energy of a family of Hamiltonian is computationally intractable, then computing the expectation values of local observables is too. These results were first given in [WBG20].

Finally, in chapter 7 we outline results from [WB21], where we considered the phase diagrams of physical systems which are simpler than those in chapter 3 — these systems are such that they only undergo a single phase transition, and if one takes a sufficiently large (and computable) section of the lattice then their phase is the same phase as in the thermodynamic limit. We show even for these simpler class of Hamiltonians, it remains computationally intractable to determine where a phase transitions takes place.

Chapter 2

Background and Previous Work

2.1 Notation

Let $\mathcal{B}(\mathcal{H})$ be the space of bounded linear operators on a complex Hilbert space $\mathcal{H}$. For $A \in \mathcal{B}(\mathcal{H})$, Null$(A)$ is the null-space/kernel of A. Define $\Lambda(L \times W) := \{1, \ldots, L\} \times \{1, \ldots, W\}$ to be the square lattice of length L, width W, with $L, W \in \mathbb{N}$. We attach to each site $i \in \Lambda(L \times W)$ in the lattice a Hilbert space $\mathcal{H}_i \cong \mathbb{C}^d$ for a constant d which we call the local Hilbert space dimension. The overall Hilbert space on of a set of lattice points $S \subset \Lambda(L)$, or indeed any graph in which the points are embedded, is then the tensor product of the all the local Hilbert spaces $\bigotimes_{i \in S} \mathcal{H}_i$.

Given a string $x \in \{0,1\}^n$, then $|x| = n$ will denote the binary length of the string. Systems with a local Hilbert space $\mathbb{C}^2$ will often be called "qubits" and systems with Hilbert space $\mathbb{C}^d$ will often be called qudits. A qudit will often be referred to as a "spin with local dimension d", or a "spin" where the local Hilbert space dimension is left implicit. Given some quantum state $|\psi\rangle \in \mathcal{H}$ and a Hermitian operator $A \in \mathcal{B}(\mathcal{H})$, the expectation value of A with respect to $|\psi\rangle$ is $\langle A \rangle := \langle \psi | A | \psi \rangle$.

A Hamiltonian is a Hermitian matrix describing the physics of a physical system. We will generally denote Hamiltonians with the letter H. We say a Hamiltonian $H \in \mathcal{B}((\mathbb{C}^d)^{\otimes N})$ is *k-local* if it can be written as:

$$H = \sum_i h_i \tag{2.1}$$

if each h_i acts non-trivially on at most k spins, $h_i \in \mathcal{B}((\mathbb{C}^d)^{\otimes k})$. That is to say, each

h_i is implicitly tensored with an identity on all the qudits it acts trivially on. For example, if given a 1D line of qudits of length N, the term $h_{2,3}$ in implicitly means:

$$\mathbb{1}_1 \otimes h_{2,3} \otimes \mathbb{1}_4 \otimes \mathbb{1}_5 \otimes \ldots \otimes \mathbb{1}_N, \tag{2.2}$$

where $h_{2,3} \in \mathcal{B}((\mathbb{C}^d)^{\otimes 2})$. Here the $h_i \in \mathcal{B}((\mathbb{C}^d)^{\otimes k})$ are called the local interaction terms. Typically we will require $\|h_i\| \leq 1$, which is to say that the norm of the local terms is independent of the total size of the physical system.

All Hamiltonians observed in the physical world are local Hamiltonians. Note that locality in this sense *does not* necessarily imply any sense of geometric or physical locality: the particles which interact may be separated by large physical distances. When specifying terms which reflect that interactions typically become weaker with distance, we will refer to this as geometric locality. Examples include nearest neighbour interactions, which only couple adjacent spins, or power-law interactions where the interaction strength between spins distance r away decays as $1/r^\alpha$ for some constant $\alpha \geq 1$.

For a given Hamiltonian H, we will denote its eigenvalues as $\lambda_i(H)$, such that $\lambda_{\min}(H) := \lambda_0(H) \leq \lambda_1(H) \leq \lambda_2(H) \leq \ldots$. The spectral gap, $\Delta(H)$, of a Hamiltonian H is then defined as:

$$\Delta(H) := \lambda_1(H) - \lambda_0(H). \tag{2.3}$$

Given a lattice $\Lambda(L \times W)$, a Hamiltonian $H = \sum_i h_i$ is *nearest-neighbour* if $h_i \in \mathcal{B}(\mathbb{C}^d \otimes \mathbb{C}^d)$ such that each h_i acts non-trivially only on neighbouring pairs of lattice sites. We write nearest neighbour sites as $\langle i, j \rangle$. Furthermore, on a D-dimensional hypercube lattice, *translational invariance* implies $h_{i,i+1} = h_{i+k,i+1+k}$ for any $k \in \mathbb{N}$, along every cardinal direction of the lattice (this is sometimes called shift invariance in other fields). We note that the interactions along any cardinal direction may be different from the interaction along any other cardinal direction[1].

[1]For example, on a 2D lattice, a translationally invariant Hamiltonian will have the same interactions along all the rows, but may have a different set of interaction along the columns

By a *classical Hamiltonian*, we mean a Hamiltonian which is diagonal in the standard basis.

A special subclass of Hamiltonians are *frustration free* Hamiltonians. A Hamiltonian $H = \sum_i h_i$ is a frustration free Hamiltonian if $\{h_i\}_i$ are all positive semi-definite and the ground state energy of the overall Hamiltonian is zero. This implies that the ground state of H is also the ground state of each h_i individually.

Occasionally we will make use of the Pauli x, y and z operators which we denote as X, Y and Z respectively, where:

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

2.1.1 Quantum Mechanics

All quantum systems have their state at a given time described by a complex vector in some Hilbert space $\mathcal{H}$ which will evolve with time, $|\psi(t)\rangle$. The physics of the system — the way in which the particles interact — is described by a Hermitian linear operator known as the Hamiltonian H. These quantum systems then evolve according to the Schrödinger equation:

$$\frac{\partial}{\partial t} |\psi(t)\rangle = iH |\psi(t)\rangle ,$$

such that after time t', assuming the Hamiltonian remains fixed, the system's state is updated as

$$|\psi(t'+t)\rangle = e^{-iHt'} |\psi(t)\rangle .$$

The Hamiltonian also gives the energy of the system as

$$\langle\psi| H |\psi\rangle = E.$$

Remarkably, once the Hamiltonian for the system is known, we have a complete description of the system and can use this to predict its properties — although we

shall see this is often a difficult task! We also see that any local observable can be represented by a Hermitian operator A, such that when a measurement is made the outcome must always be an eigenvalue of A, and the state which results is the corresponding eigenvalue of A.

2.1.1.1 Phases of Matter

The reader will most likely be familiar with phases of matter in the context of thermal phases, such as ice melting to water and evaporating to steam. More generally, many-body systems are often characterised by a set of parameters, such as coupling strengths between atoms, temperature, or the strength of an applied magnetic field. For certain ranges of parameters, the system may display particular collective properties, and then undergo a rapid transition at a critical point to form a new phase with different properties. For example, in the transition from water to ice, the relevant parameters are pressure and temperature.

The mechanisms for phase transitions have confounded physicists since the early 20th Century (and probably before) due to the often surprising collective behaviour of the relatively simple constituent parts. Early models for phase transitions include the van der Waals gas model of hard sphere molecules, and the famous Ising model of atoms modelled as tiny magnetic moments.

If one restricts the system to zero temperature, then these phase transitions can still happen as a function of non-thermal parameters (e.g. applied magnetic field or chemical composition), and are known as quantum phase transitions [Sac11]. While so called thermal phase transitions can usually be understood a being driven by a change in temperature — the energy of the constituent atoms/molecules overcomes some energy barrier — quantum phase transitions are best thought of as being driven by quantum fluctuations which become relevant only near absolute zero. Quantum phase transitions describe the many phenomena such as low temperature superconductivity and the quantum hall effect.

Definition 2.1 (Quantum Phase Transition (QPT), from [Sac11])**.** *Consider a local Hamiltonian* $H(\varphi) = \sum h_i(\varphi)$*, where the matrix entries of* $h_i(\varphi)$ *are analytic in* φ*. In the thermodynamic limit, a* quantum phase transition *occurs where there is a*

non-analytic change in the characteristics of the ground state as a function of φ (e.g. the ground state energy $\lambda_{\min}(H(\varphi))$). If the matrix elements of a Hamiltonian are functions of multiple parameters $h_i(\varphi_1, \theta_2, \dots)$, then there is a phase transition at $\varphi = \varphi^*$ if there is a non-analytic change in the ground state energy at φ^* when all other parameters are held constant.

Quantum phase transitions can be extraordinarily complex: the phase diagram for the phase Hall effect consists of a fractal "Hoftstadter butterfly" with an infinite number of phases and phase transitions [OA01]. Moreover QPT characterise important behaviour in physics, such as the metal-insulator phase transitions [Voj00] or the structure of atomic nuclei [Elh+16]. Despite the relative importance of quantum phase transitions in explaining collective behaviour of systems, their complexity means they remain poorly understood.

If an experimentalist is working in a lab, it is likely difficult to measure non-analyticities in the ground state energy. Instead, it is often useful to work in terms of an order parameter which characterises some property of the phase. In the example of the transverse 1D Ising model, the order parameter corresponds to the magnetisation along the Z axis. There is one phase which has zero expected magnetisation, and another with ± 1 magnetisation. Throughout this thesis we will find it useful to work with order parameters.

2.1.2 Spin Systems

A key model for condensed matter physicists are the spin lattice Hamiltonians. These are idealised models of real materials in which a set of spins (which replicate the magnetic moments of atoms or molecules) are arranged in a lattice and given interactions between them. The interactions depend on the particular physical system they are trying to model, and under different guises they are used as models for everything from superconductivity [Hub63] to quantum memories [Den+02]. Remarkably, despite their simplicity, these models can accurately display a wide range of physical phenomena, whether it be magnetism (quantum or classical) or topological phenomena. This thesis will be concerned primarily with studying this class of Hamiltonians.

For the most part we will assume spin systems can accurately be described by local Hamiltonians, usually with some sense of geometric locality.

2.2 Computability and Computational Complexity

Computational complexity is a subfield of theoretical computer science that seeks to describe how difficult a given computational task is. Intuitively, it is obvious that some tasks are more difficult that others. An example the reader can test for themselves is the following: take two numbers of length 10 bits and add them together. Now consider a second problem: choose a random 20 bit integer and try to find its prime factors. The reader should have found the factoring significantly harder, despite the information input to both problems being 20 bits.

Complexity theory attempts to formalise how hard these computational problems are, and has lead to the development of so-called complexity classes. Complexity classes classify computational problems (such as addition, factorisation, route finding, etc.) in terms of the "amount of computational resources" need to solve them. By "computational resources" we mean the amount of time or space a Turing Machine needs to solve this class of problems. Turing Machines (TM) are a model of computation consisting of a control "head" and an infinite tape, such that the tape is divided up into cells. The head has several internal states that it can be in, and can read the tape and write symbols on the tape. Upon reading a cell in the tape, the TM head can write over it, transition to a new internal state and then move left or right — the particular action it take at a given point depends on the contents of the cell it is reading and the state the head is in. Turing Machines have become the standard model of computation in theoretical computer science. We give a rigorous definition of a TM a little later in definition 2.2. Remarkably, it is believed that all models of computation are equivalent to Turing Machines, and hence Turing Machines provide a natural way of characterising computation complexity (known as the Church-Turing thesis).

We have been deliberately imprecise in the previous paragraph, which we clarify here. Computational complexity primarily deals with "decision problems"; those

which can be formulated as having a yes or no answer. Thus, rather than asking for the prime factors of a number, instead the decision version of the problem asks whether a number has a prime factor above or below some integer k. It is clear that one can convert easily between the two (i.e. the decision version of factoring can be used to solve find the prime factors of a number by repeated applications).

The best known complexity classes are P and NP. The former, describes the set of decision problems which are solvable by polynomial time algorithms run on a TM: the algorithm runs for time at most $O(n^c)$ for an input of length n and a constant c. The class NP is the set of problems which can be verified in polynomial time: that is if one is given the answer to a computational problem, then one can check it is correct in polynomial time.

The factoring example above is a (notorious) example of a problem in NP, but thought not to be in P. Finding the prime factors of an n bit problem is thought to be difficult enough that no polynomial algorithm exists for the task. However, if one is given a set of integers, and asked to check they are the prime factors of some larger number, it is easy to do so by multiplying them together. Hence it has a polynomial time verification algorithm.

A fundamental problem in computer science is the question of whether $\mathsf{P} = \mathsf{NP}$ or $\mathsf{P} \neq \mathsf{NP}$. Remarkably, this apparently simple problem is still open. Unfortunately, this is a problem endemic to complexity theory: proving that complexity classes are equal to each other (or not) is extremely difficult, and the exact relationships between most important classes are unknown.

To compare the hardness of decision problems we often talk about reductions. We say that one computation problem A reduces to another B under a polynomial time reduction, if there is a polynomial time Turing Machine which takes the output of B and then uses it to solve A. In this sense problem B is at least as hard as problem A. If the reduction goes both ways then the two problems are equally hard.

Perhaps surprisingly, for most complexity classes there exist special subsets of problems called "hard" problems such that *any* problem in the complexity class can be reduced to a hard problem. If a problem is hard for a complexity class and

contained in that same complexity class we call it "complete" for that class. Well known NP-complete problems include the Travelling Salesman Problem and the Satisfiability problem.

2.2.1 Quantum Complexity Classes

During the 1980s and beyond, it was speculated that Turing Machines running according to the rules of quantum mechanics may show an advantage in a variety of areas [Fey82]. Some initial evidence was supplied by the Deutsch's algorithm which was capable of solving a problem with fewer queries than a deterministic classical TM, and then by a series of improvements which showed a quantum TM can use exponentially fewer queries than a classical TM [DJ92]. Convincing evidence of a quantum speed-up came with the arrival of Shor's algorithm which demonstrated that a quantum logic circuit could solve the factoring problem in polynomial time [Sho97], whereas the best known classical algorithms take quasi-exponential time despite decades of effort searching for improvements.

As such, there developed a parallel theory of complexity theory for quantum computers [BV97]. These classes are typically formulated as the set of problems which can be solved by quantum Turing Machines or quantum circuits. Quantum circuits are simply logic circuits where the logical gates must be unitaries chosen from some fixed, predetermined gate-set. Quantum Turing Machines (QTMs) are Turing Machines in which the transition rules are unitaries [BV97]. We give a rigorous definition of QTM later in definition 2.7.

Since quantum mechanics is inherently probabilistic, the complexity classes are typically phrased in terms of what set of problems can be solved *with high probability*, as opposed to those which can be solved deterministically. As such, the natural analogue of P is a class known as BQP. A problem is in BQP if it the correct answer can be found with probability $> 2/3$. Similarly, the natural quantum analogue of NP is QMA: the set of problems for which given a proof of the answer, can verify whether or not it is correct with probability $> 2/3$ [KSV02]. We note here that the bound $2/3$ is arbitrary and that by running the problem multiple times, one can amplify the output probability to get exponentially close to 1 [NC10].

We note that BQP and QMA are technically classes of *promise problems*, in which one is guaranteed some property of the problem one is studying holds, and cases where the promise does not hold are not studied.

As in the classical case, the relations between quantum and classical complexity classes are unknown. The fact factoring is in BQP provides evidence that the inclusion $\mathsf{P} \subset \mathrm{BQP}$ is strict, but it is not expected that $\mathsf{NP} \subseteq \mathrm{BQP}$ or $\mathrm{BQP} \subseteq \mathsf{NP}$ [RT19]. Beyond factoring there exist other problems which are efficiently solvable on a quantum computer which appear to be classically difficult. Two important ones are simulating physical quantum systems [Llo96] and solving linear systems of equations [HHL09].

2.2.2 Computability

If complexity theory is the study of how hard it is for a TM to solve a problem, it is natural to ask whether there are problems which cannot be solved by any TM in any finite amount of time — they are undecidable or uncomputable[2]. Turing demonstrated one such problem: determining whether an arbitrary TM halts on some input is undecidable [Tur37]. That is to say, there exists no finite time algorithm which will take as input a description of a Turing Machine and output, in a finite amount of time, determine whether that TM eventually halts or not. This was then extended by Rice to show that any non-trivial, semantic property of a Turing Machine is undecidable [Ric53].

Since Turing's original proof, a wide range of problems in computer science, mathematics and physics have been shown to be undecidable. These include determining whether a finite set of tiles will tile a plane [Rob71], determining the Kolmogorov complexity of a string [LV93], or determining whether a player has a winning strategy in a game of *Magic: The Gathering* [CBH19], as well as many others. Notably for physicists, the problems of determining the spectral gap of a Hamiltonian on an infinite lattice is undecidable [CPGW15b], as is determining the trajectory of a particle in a potential [Moo90]. We will discuss the spectral gap

[2]Technically, undecidability applies only to decision problem whereas uncomputability applies to functions.

problem in more detail later.

While quantum computers promise to solve some problems faster than classical computers, a quantum computer can be simulated by a classical computer with exponential overhead, hence the set of decidable problems is the same for classical and quantum computers — thus preserving the Church-Turing thesis. We note that there exists a strong version of the Church-Turing thesis (conveniently known as the Strong Church-Turing thesis) which postulates that all "reasonable" models of computation are equivalent to each other up to polynomial overhead. As such, quantum computation appears to violate this stronger version of the thesis.

2.2.3 Complexity Definitions

Here we give the formal definitions of the concepts we will need for complexity theory. We include the definition of Turing Machines here for completeness, but we will rarely use the full definition.

Definition 2.2 (Turing Machine [BV97])**.** *A deterministic 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 two final states* $q_f \neq q_0$*, and* δ*, the deterministic transition function, is a function*

$$\delta : Q \times \Sigma \to \Sigma \times Q \times \{L, R\} \tag{2.4}$$

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 or instantaneous description 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 tape cells may contain non-blank symbols. The initial contents of the tape is known as the input.*

For any configuration c *of TM* M*, the successor configuration* c_0 *is defined by applying the transition function to the current state* q *and currently scanned symbol* σ *in the obvious way. We write* $c \to_M c_0$ *to denote that* c_0 *follows from* c *in one step. By convention, we require that the initial configuration of* M *to satisfies the following*

conditions: the tape head is in cell 0*, called the start cell, and the machine is in state* q_0*. An initial configuration has input* $x \in (\Sigma - \#)^*$ *if* x *is written on the tape in positions* $0, 1, 2, \ldots$*, and all other tape cells are blank. The TM halts on input* x *if it eventually enters the final state* q_f*, from which there are no further transition rules. 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 which halts on all inputs therefore computes a function from* $(\Sigma - \#)^*$ *to* Σ^**.*

The main objects of study in classical computer science are decision problems, also known as languages, which we define below:

Definition 2.3 (Decision Problem or Language [AB10])**.** *Consider a Boolean function* $f : \{0,1\}^* \to \{0,1\}$*. We identify such a function* f *with the set* $L_f = \{x : f(x) = 1\}$ *and call such sets languages or decision problems (we use these terms interchangeably). We identify the computational problem of computing* f *(i.e., given* x *compute* $f(x)$*) with the problem of deciding the language* L_f *(i.e., given* x*, decide whether* $x \in L_f$ *).*

On an instance of a decision problem/language $x \in \{0,1\}^*$*, we say that a TM* M *accepts if it outputs* 1 *(i.e.* $M(x) = 1$*), and rejects if it outputs* 0 *(i.e.* $M(x) = 0$*).*

We note that sometimes rather than being given a single final state q_f, TMs are sometimes defined as having a pair of final states q_A, q_R which signal whether a TM is accepting or rejecting on some computation (e.g. in [Pap94]). Similar to single halting state, there are no transitions out of q_A or q_R.

We now define the class NP. This is the class of decision problems for which, if there is a "yes" output for the problem, then there is a witness or proof w which can be efficient verified by a classical TM.

Definition 2.4 (NP)**.** *A language* $L \subseteq \{0,1\}^*$ *is in* NP *if there exists polynomial* p*, and a deterministic polynomial time Turing Machine* M *such that for each instance* $x \in \{0,1\}^*$ *either:*

Accepts: if $x \in L$*,* $\exists w \in \{0,1\}^{p(|x|)}$ *such that* $M(x,w) = 1$ *with probability 1.*

Rejects: if $x \notin L$ *then* $\forall w \in \{0,1\}^{p(|x|)}$, $M(x,w) = 0$ *with probability 0.*

We will also be interested in the exponential version of this class:

Definition 2.5 (NEXP)**.** *A language* $L \subseteq \{0,1\}^*$ *is in* NEXP *if there exists polynomial p, and a deterministic exponential time Turing Machine M such that for each instance* $x \in \{0,1\}^*$ *either:*

Accepts if $x \in L$, $\exists w \in \{0,1\}^{2^{p(|x|)}}$ *such that* $M(x,w) = 1$ *with probability 1.*

Rejects if $x \notin L$ *then* $\forall w \in \{0,1\}^{2^{p(|x|)}}$, $M(x,w) = 0$ *with probability 0.*

Clearly NEXP is a much larger class of problems than NP and NP $\subset$ NEXP. We also define the concept of (Karp) reducibility:

Definition 2.6 (Polynomial Time Reductions [AB10])**.** *We say a language* $A \subset \{0,1\}^*$ *is polynomial-time Karp reducible to a language* $B \subseteq \{0,1\}^*$ *(sometimes shortened to just "polynomial-time reducible") if there is a polynomial-time computable function* $f : \{0,1\}^* \to \{0,1\}^*$*, such that for every* $x \in \{0,1\}^*$, $x \in A$ *if and only if* $f(x) \in B$.

As mentioned, we can also define a quantum Turing Machine which can run quantum computations. We note for simplicity, many formulations of quantum computation use quantum circuits as they are conceptually easier to deal with. However, we will make use of QTMs in this thesis and so include the definition for completeness.

Definition 2.7 (Quantum Turing Machine [BV97])**.** *Call* $\tilde{\mathbb{C}}$ *the set consisting of* $\alpha \in \mathbb{C}$ *such that there is a deterministic algorithm that computes the real and imaginary parts of* α *to within* 2^{-n} *in time polynomial in n. A Quantum Turing Machine M is described by a tuple* (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*

$$\delta : Q \times \Sigma \to \tilde{\mathbb{C}}^{\Sigma \times Q \times \{L,R\}}. \tag{2.5}$$

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 S be the inner-product space of finite complex linear combinations of configurations of M with the Euclidean norm. We call each element $|\phi\rangle \in S$ a superposition of M. The QTM M defines a linear operator $U_M : S \to S$, called the time evolution operator of M, as follows: if M starts in configuration c with current state p and scanned symbol σ, then after one step M will be in superposition of configurations $|\psi\rangle = \alpha P$, where each non-zero α_i corresponds to a transition $\delta(p, \sigma, \tau, q, d)$, and c_i is the new configuration that results from applying this transition to c. Extending this map to the entire space S through linearity gives the linear time evolution operator U_M.

For decision or promise problems run on a QTM, it is not immediately clear how to define acceptance or rejection.

We now rigorously define the quantum analogue[3] of the class P.

Definition 2.8 (BQP)**.** *A promise problem $L = (L_{YES}, L_{NO}) \subset \{0,1\}^*$ is in* BQP *if there exists a polynomial q and a Quantum Turing Machine M such that for each instance $x \in L$, $M(x)$ halts in $q(x)$ steps. Furthermore, let $\Pi = |1\rangle\langle 1|$ be a projector onto the first qubit of the QTM's track, then if a M produces an output state $|\phi\rangle$:*

Accepts: $\langle\phi|\,\Pi\,|\phi\rangle > 2/3$ *if* $x \in L_{YES}$.

Rejects: $\langle\phi|\,\Pi\,|\phi\rangle < 1/3$ *if* $x \in L_{NO}$.

The natural analogue of NP is the class of promise problems for which a witness can be given as a quantum state and can be efficiently checked using a quantum circuit or QTM.

Definition 2.9 (QMA)**.** *A promise problem $L = (L_{YES}, L_{NO}) \subset \{0,1\}^*$ is in* QMA *iff there exists a a Quantum Turing Machine M such that for each instance x and any $|\psi\rangle$ on $O(p(|x|))$ qubits, on input $(x, |\psi\rangle)$, M halts in $O(q(|x|))$ steps, for polynomials p, q. Furthermore, let $\Pi = |1\rangle\langle 1|$ be a projector onto the first qubit of the QTM's track, then if a M produces an output state $|\phi\rangle$:*

[3] Technically, BQP is closer to the classical class BPP, but "morally" BQP plays the same role as P.

Accepts: *if* $x \in L_{YES}$, $\exists\, |\psi\rangle$ *such that* M *runs on input* $(x, |\psi\rangle)$ *and the output satisfies* $\langle\phi|\,\Pi\,|\phi\rangle > 2/3$.

Rejects: *if* $x \in L_{NO}$, *then* $\forall\, |\psi\rangle$, M *runs on input* $(x, |\psi\rangle)$ *and the output satisfies* $\langle\phi|\,\Pi\,|\phi\rangle < 1/3$.

Problems which are QMA-complete are believed to be intractable for a quantum computer. We can also define the associated exponential class.

Definition 2.10 ($\mathsf{QMA}_{\mathsf{EXP}}$)**.** *A promise problem* $L = (L_{YES}, L_{NO}) \subset \{0,1\}^*$ *is in* $\mathsf{QMA}_{\mathsf{EXP}}$ *iff there exists a* k *and a Quantum Turing Machine* M *such that for each instance* x *and any* $|\psi\rangle$ *on* $O(2^{|x|^k})$ *qubits, on input* $(x, |\psi\rangle)$, M *halts in* $O(2^{|x|^k})$ *steps. Furthermore, let* $\Pi = |1\rangle\langle 1|$ *be a projector onto the first qubit of the QTM's track, then if a* M *produces an output state* $|\phi\rangle$*:*

Accepts: *if* $x \in L_{YES}$, $\exists\, |\psi\rangle$ *such that* M *runs on input* $(x, |\psi\rangle)$ *and the output satisfies* $\langle\phi|\,\Pi\,|\phi\rangle > 2/3$.

Rejects: *if* $x \in L_{NO}$, *then* $\forall\, |\psi\rangle$, M *runs on input* $(x, |\psi\rangle)$ *and the output satisfies* $\langle\phi|\,\Pi\,|\phi\rangle < 1/3$.

We note that some previous literature has made use of the abbreviation QMAEXP, which we avoid in this thesis.

Throughout, we will make use of oracle classes: these are the set of problems solvable by a Turing Machine with access to an oracle solving some problem (or class of problems).

Definition 2.11 (Oracle Turing Machines [AB10])**.** *An oracle Turing machine is a TM,* M*, that has a special read/write tape we call* M*'s oracle tape and three special states* q_{query}, q_{yes}, q_{no}*. To execute* M*, we specify in addition to the input a language* $O \subset \{0,1\}^*$ *that is used as the oracle for* M*. Whenever during the execution* M *enters the state* q_{query}*, the machine moves into the state* q_{yes} *if* $q \in O$ *and* q_{no} *if* $q \notin O$*, where* q *denotes the contents of the special oracle tape. Note that, regardless of the choice of* O*, a membership query to* O *counts only as a single computational step. If*

M is an oracle machine, $O \subset \{0,1\}^*$ a language, and $x \in \{0,1\}^*$, then we denote the output of M on input x and with oracle O by $M^O(x)$.

Definition 2.12 (Oracle Classes [AB10])**.** *For every $O \subset \{0,1\}^*$, P^O is the set of languages decided by a polytime deterministic TM with oracle access to O and NP^O is the set of languages decided by a polytime nondeterministic TM with oracle access to O.*

There are often definitional problems with oracle classes which arise. We will flag these and discuss them when necessary.

2.3 Hamiltonian Complexity and Current State of the Field

We've briefly discussed condensed matter physics and complexity theory, but how do these apparently vastly different fields relate to each other? While some physical systems can be solved to give a neat set of formula for every property — such as the 2D classical Ising model — many systems of interest are simply too complex to be nicely characterised by a set of closed-form equations. Sometimes these systems are amenable to statistical arguments, but in general one must resort to using numerics to study them. For such systems, the question then becomes "can we develop an algorithm which will efficiently solve this system and output the desired properties?" This very naturally becomes "can we characterise the computational complexity of Hamiltonians and their properties?". This is the field of Hamiltonian Complexity.

As mentioned previously, studying physical systems is far from just an academic interest: modelling Hamiltonians and predicting their properties is essential for a wide range of commercial interests including battery development, optimising chemical processes, and drug development [MBK15]. Classical algorithms for solving physical properties of systems have been widely employed both for academic and commercial purposes. Often these algorithms lack rigorous theoretical backing: they appear to work "well enough", are motivated by some physical intuition, or make some simplifying assumptions to make the problem more tractable. Famous

examples of such methods include the density matrix renormalisation group [Whi92], and density functional theory [Par85]. By studying the complexity of Hamiltonians, we can see to what extent we hope to find efficient algorithms for such problems.

2.3.1 The Local Hamiltonian Problem

Given some quantum system, described by a Hamiltonian, often the first properties a physicist is interested in are its energy levels and the corresponding eigenstates. The most important of these is the ground state (the lowest energy state of the Hamiltonian), which characterises the low temperature properties of the Hamiltonian.

The first attempt to study the complexity of finding the ground state energy (or indeed any property of a physical systems) was made by Barahona [Bar82], where it was demonstrated that approximating the ground state energy of a classical Ising model is NP-complete by giving a reduction from MAX-SAT to the ground state energy of the Ising model. As a result, it is believed that finding the ground state of spin-glasses is intractable, even if restricted to entirely classical Hamiltonians. However, if the interactions of the Ising model are restricted to be a planar graph, then the problem suddenly becomes tractable and there is a polytime algorithm [Had75].

A foundational result for quantum systems was proved by Kitaev [KSV02], where it was shown that estimating the ground state energy of a 5-local Hamiltonian was QMA-complete, and in the process demonstrated the class QMA is the natural quantum analogue to NP. More formally, the result shows that the following promise problem, is QMA-complete:

Definition 2.13 (Local Hamiltonian problem (LH))**.**
Input: *A k-local Hamiltonian $H = \sum_i^m h_i$, acting on N qudits, and $m = O(\text{poly}(N))$, with two parameter α, β such that $\beta - \alpha = \Omega(1/\text{poly}(N))$.*
Output:

YES: if $\lambda_0(H) < \alpha$

NO: if $\lambda_0(H) > \beta$.

Promise: $\lambda_0(H) \notin [\alpha, \beta]$.

A key tool in Kitaev's proof of QMA-hardness is a technique known as the circuit-to-Hamiltonian mapping, which we review in the following subsection. The QMA-completeness of the Local Hamiltonian problem implies that we do not expect to efficiently estimate ground state energy of local Hamiltonians, even with access to quantum computers. It is worth noting that the results [Bar82] and [KSV02] only apply to spin-glasses: collections of particles for which the interactions can be chosen independently. Furthermore, for these results there are no geometric constraints on how the particles interact with each other. As such, these result do not necessarily prove QMA-completeness of the Local Hamiltonian problem for systems with interactions similar to those found in nature.

2.3.1.1 Circuit-to-Hamiltonian Mappings

The circuit-to-Hamiltonian mapping is the central tool for proving complexity results, and can itself be viewed as a quantum version of the much celebrated Cook-Levin theorem. The idea is to construct a Hamiltonian $H \in \mathcal{B}((\mathbb{C}^d)^{\otimes N})$ such that its ground state is a superposition of some computation at all different steps of the computation. Such Hamiltonians are also referred to as *Feynman-Kitaev Hamiltonians* or *History State Hamiltonians*.

We let the Hilbert space of the Hamiltonian have two parts:

$$\mathcal{H} = \mathcal{H}_C \otimes \mathcal{H}_R$$

where $\mathcal{H}_C$ will contain a set of "clock states" which record which time step the register is in, and $\mathcal{H}_R$ which is the computational register of the computation. In particular, let there be a computation starting with an initial state $|\psi_0\rangle \in \mathcal{H}_R$, and lasting for T steps such that the computational register after t steps is $|\psi_t\rangle = U_t U_{t-1} \dots U_1 |\psi_0\rangle$, then this can be encoded in a so called *history state*:

Definition 2.14 (History state)**.** *A* history state $|\Psi_{hist}\rangle \in \mathcal{H}_\mathrm{C} \otimes \mathcal{H}_\mathrm{R}$ *is a quantum state of the form*

$$|\Psi_{hist}\rangle = \frac{1}{\sqrt{T}} \sum_{t=1}^{T} |t\rangle_\mathrm{C} |\psi_t\rangle_\mathrm{Q}, \tag{2.6}$$

where $\{|1\rangle, \ldots, |T\rangle\}$ *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* $|\psi_0\rangle \in \mathcal{H}_R$ *and set of unitaries* $U_i \in \mathcal{B}(\mathcal{H}_R)$.

$\mathcal{H}_C$ *is called the* clock register *and* $\mathcal{H}_R$ *is called the* computational register. *If* U_t *is the unitary transformation corresponding the* t^{th} *step of a quantum computation — which may be a gate in a quantum circuit or a QTM transition — then* $|\psi_t\rangle$ *is the state of the computation after t steps. We say that the history state* $|\Psi\rangle$ encodes *the evolution of the quantum computation.*

To construct a Hamiltonian with $|\Psi_{hist}\rangle$ as a ground state, Kitaev develops a Hamiltonian which is the sum of three parts. The first is the propagation Hamiltonian.

$$H_{prop} := \sum_{t=1}^{T} \left(|t\rangle \langle t-1|_C \otimes U_t + |t-1\rangle \langle t|_C \otimes U_t^\dagger \right). \tag{2.7}$$

This forces the ground state to be the evolution of some quantum state through a particular quantum circuit. Kitaev also introduces an initialisation term:

$$H_{init} := |0\rangle\langle 0|_C \otimes \Pi_{init}. \tag{2.8}$$

This is a projector onto the first time step which ensures the computational register is initialised in whatever we want the initial state to be (which will usually be the input to the computation). There is a third term H_{Clock} designed to enforce the correct propagation of the clock. How the clock is implemented depends on the particular construction and so we leave it nondescript here. The resulting Hamiltonian $H = H_{init} + H_{prop} + H_{clock}$ has a history state ground state as per eq. (2.6).

Proving QMA-hardness

To prove QMA-hardness of the Local Hamiltonian problem, another term must be added to the Hamiltonian above. This term is chosen to penalise rejected outputs. In particular we assume that there is an output qubit which onto which the output of the computation is placed. Then add the following term to the Hamiltonian:

$$H_{out} := |T\rangle\langle T| \otimes |0\rangle\langle 0|_{out},$$

to give:

$$H = H_{init} + H_{prop} + H_{out} + H_{clock}. \tag{2.9}$$

To prove QMA-hardness of the Local Hamiltonian problem for this Hamiltonian, we allow part of the initial computational register to be unconstrained; this part of the register will be taken as a witness state. We then use the circuit-to-Hamiltonian mapping to encode a verification circuit which takes as input some $x \in \{0,1\}^n$ and some witness state $|w\rangle \in (\mathbb{C}^2)^{\otimes O(\text{poly}(n))}$, runs for time $T = O(\text{poly}(n))$, and then outputs either $|0\rangle$ or $|1\rangle$ with high probability.

Remarkably, one can show using a basis transformation that the Hamiltonian in eq. (2.9) has the same energy as a particle propagating along a line and receiving an energy penalty on the final step, where the energy penalty is proportional to $Tr(|\psi_T\rangle\langle\psi_T| \mathbb{1} \otimes |0\rangle\langle 0|_{out})$ [KSV02]. This allows us to reduce the eigenvalue analysis to that of solving a tridiagonal Toeplitz matrix. The energy is then either $\lambda_0(H) \leq \alpha$ or $\lambda_0(H) \geq \beta$ if the computation accepts or rejects respectively, for some pair (α, β) such that $\beta - \alpha = \Omega(1/\text{poly}(N))$. Thus the Local Hamiltonian problem is QMA-hard.

Towards More Natural Systems

The initial construction by Kitaev to prove hardness of the Local Hamiltonian problem was for a 5-local, non-geometrically local, highly non-physical Hamiltonian. In natural physical systems, interactions between bodies tend to be 2-local, bodies typically couple more strongly with particles which are geometrically closer to them, and there is often some underlying interaction structure (e.g. particles are on a lattice or some other graph).

The more "natural" or simple a Hamiltonian is while also displaying hard-to-compute properties, the more insight this gives us into the intrinsic complexity of nature and into the structure of Hamiltonians as a whole. While Kitaev's initial hardness proof prevents one from finding an efficient algorithm for solving ground state energies of 5-local Hamiltonians with no geometric structure, it leaves open the possibility that the same problem on a 1D chain of qubits with nearest neighbour interactions efficiently solvable.

The Hamiltonians observed in real physical systems possess a wide range of properties which one might make them simpler to solve. First of all, almost all interactions seen in nature are 2-local, and very often shielding factors mean the interactions between particles decays exponentially with distance. Condensed matter systems often possess symmetries, such as a regular lattice, and can be effectively modelled as Hamiltonians with only nearest neighbour interactions. Furthermore, the local interactions between particles are often the same between ever pair of neighbouring particles, and hence the interactions are both rotationally and translationally invariant. Beyond the properties of the local interactions, properties such as having a large spectral gap, or being frustration free can be shown to make systems easier to solve.

Perhaps intuitively one can understand this by the fact that by imposing simpler interactions and symmetries on a Hamiltonian, and other conditions on a physical system, there are fewer degrees of freedom in which one can "encode" complexity or computation. Alternatively, more degrees of freedom reduces size of the solution space on must search through.

An important goal of Hamiltonian complexity is to examine increasingly simple systems and determine where the boundary between complex and simple lies. With this in mind there has been a huge number of results proving QMA-completeness of Hamiltonians for increasingly simple systems. The two primary methods for this are improvements to the *circuit-to-Hamiltonian mapping* and through *perturbation gadgets*. We review these advances separately.

Adaptations of the Circuit-to-Hamiltonian Mapping

Over many works, new variants of circuit-to-Hamiltonian mappings have been developed to apply the method to more realistic Hamiltonians which have properties similar to those we find in nature.

One of the first improvements was made in [KR03a] which improved locality to 3-local, and [KKR06], who reduced the locality to 2-local Hamiltonians. In [Aha+07] it was shown that the Local Hamiltonian problem was QMA-hard for 1D Hamiltonians of local Hilbert space dimension 13. Remarkably, this was then improved upon in

[GI09], where it was shown how a 1D, translationally invariant, nearest neighbour Hamiltonian can be used to encode a circuit-to-Hamiltonian mapping [4]. This result is especially important as the only input to the Hamiltonian is the length of the chain it is specified on — all other parts of the Hamiltonian remain fixed. This allows $\mathsf{QMA}_{\mathsf{EXP}}$-hardness of the Local Hamiltonian problem to be proven for 1D, translationally invariant, nearest neighbour Hamiltonians.

However, the Hamiltonian in [GI09] has a huge local Hilbert space dimension which is highly unlikely to be observed in nature. [BCO17] use Quantum Thue Systems to reduce the local Hilbert space dimension to 42. [BP17a] proves $\mathsf{QMA}_{\mathsf{EXP}}$-hardness on a 3D lattices with face-centred cubic unit cells, but for a local Hilbert space dimension of 4, and for 4-local Hamiltonians.

The Space-time Circuit-to-Hamiltonian Mapping

There is an similar, but slightly conceptually different version of the circuit-to-Hamiltonian mapping that has been used throughout the literature, often known as the *space-time circuit-to-Hamiltonian mapping*. The construction presented in the previous sections has a ground state which consists of the computational register tensored with a global clock, and then puts these in superposition. Space-time circuit-to-Hamiltonian mappings forgo a global clock construction, and instead give each qubit a local clock variable. The ground state is then a superposition of all valid configurations of these clocks. The construction is based on those in first utilised in [Zur90] and [Jan07]. This construction has been used to prove QMA-completeness for interacting fermions [BT14], to develop better quantum error-correcting codes [Boh+19], and give an alternative proof to the equivalence between adiabatic computation and the circuit model [MLM07; GTV15].

To illustrate how this works, we give an example of how the computation is encoded in [BT14]. Assume we want to encode a depth D circuit of n qubits. Then we give each qubit a local clock register of $D-1$ states, $|t\rangle$, $t \in \{0,1,2,\dots D\}$ which records the time-step in the circuit. Let $U_t^1[q]$ be a 1-qubit gate acting on the qubit q on the t^{th} time-step. Then for each single-qubit gate there is a term in H_{trans} of the

[4] Technically Gottesman and Irani encode a QTM's evolution in a Hamiltonian rather than a circuit, so this should more properly be called a "QTM-to-Hamiltonian mapping".

form:

$$H_t^1[q] = (U_t^1[q] \otimes |t\rangle - |t-1\rangle)(U_t^{1\dagger}[q] \otimes \langle t| - \langle t-1|).$$

Denote two qubit gates acting on qubits p and q at time-step t as $U_t^2[q,p]$. Then such a gate is encoded in the Hamiltonian as:

$$H_t^2[q,p] = (U_t^2[q,p] \otimes |t,t\rangle - |t-1,t-1\rangle)(U_t^{2\dagger}[q,p] \otimes \langle t,t| - \langle t-1,t-1|).$$

The D-dimensional qudit can then be mapped down to a set of $D+1$ qubits as $|t\rangle = |0\rangle_1 |0\rangle_2 \ldots |0\rangle_t |1\rangle_{t+1} |0\rangle_{t+2} \ldots |0\rangle_{D+1}$. Thus the time dimension of the circuit is mapped to a spacial dimension, where rows of qubits represent a particular time step (hence the inclusion of "space-time" in the name of the construction). Penalties can then be added to the Hamiltonian such that only valid time configurations occur in the history state (i.e. so that a gate on one qubit line does not get applied in an incorrect order). Input and output penalties are then added on appropriately.

We also note an innovative construction used to prove the Local Hamiltonian problem is QMA-complete for the Bose-Hubbard model [CGW13]. Rather than encoding the desired computation in the interactions, it keeps the strength and form of the interactions constant and instead uses the interaction graph to encode computations. Importantly, the interactions here are comparatively simple when compared to the techniques above which encode circuits or Turing Machines (though the interaction geometry is comparatively complex). These techniques were then adapted to prove QMA-completeness for the XY-model with magnetisation [CGW16].

Although in this section we have focused on circuit-to-Hamiltonian mappings for QMA-hardness construction, they have been used to prove a variety of other results. Consequently there are a variety of constructions useful for different systems which are often trying to improve on different properties of the system or construction [CLN18; BC18a; GC18].

2.3.1.2 Perturbative Gadgets

Beyond circuit-to-Hamiltonian mappings, there exists another common method of proving hardness results for Hamiltonians, known as perturbative gadgets.

First developed in [KKR06] as an alternative method to prove QMA-hardness of the Local Hamiltonian problem, perturbative gadgets are additional qubits (or qudits) added to the Hamiltonian which allow the form and type of the Hamiltonian's interaction terms to be rearranged, while preserving the spectrum up to some small error. Using perturbative gadgets, it is possible to take an initial Hamiltonian with complicated interactions and high locality, and find a new Hamiltonian which has simple, 2-local interactions with lower spacial dimensionality (e.g. of the form of the Heisenberg Hamiltonian [CM13; CMP18]) such that the low energy spectrum is preserved. The fundamental idea to change the form of interactions is the following: suppose there is a 3-local interaction between qubits q_1, q_2, q_3 that we wish to reduce to a set of 2-local interactions. We call this 3-local Hamiltonian H_{target}. We then introduce a new gadget Hamiltonian with auxiliary qubits $\tilde{H} = H + V$, where H is diagonal and has spectral gap Δ, and where V is a 2-local perturbation with off diagonal terms. The new qubits then "mediate" the interaction between the original qubits q_1, q_2, q_3. By expanding $\tilde{H}$ in the low energy subspace using a Feynman-Dyson series (other expansions can be used), then for sufficiently large Δ and an appropriate choice of V the difference $||\tilde{H} - H_{target}||$ can be made sufficiently small.

The introduction of perturbative gadgets sparked a cascade of results: interacting qubit Hamiltonians on a 2D lattice are QMA-complete [OT08], 2-local qubit systems with restricted types of interactions are QMA-complete [BL08], and that density functional theory is QMA-complete [SV09] (this results also proves the LH problem for the Hubbard model in an external magnetic field is QMA-complete). Recently this result has been utilised to prove that determining the ground state energy of interacting electron orbitals is QMA-complete [O'G+21]. Finally, perturbative gadgets have been used to providing a complexity classification for all qubit Hamiltonians [CM13]. Indeed, Cubitt and Montanaro demonstrate that the Local Hamiltonian problem for a given class of Hamiltonian is complete for one of the following classes: P, NP,

StoqMA, or QMA. This result was then generalised to qudit Hamiltonians [PM18].

We note that one of the disadvantages of using perturbative gadgets to prove complexity results is that one needs to add large energy factors to the local interaction terms in the new Hamiltonians. For large systems, these factors can be enormous and arguably non-realistic. Work has been done to address this criticism — making the gadgets more "physical" — but these alternative constructions usually contain undesirable trade-offs (e.g. additional ancillary qubits) [CN15; Bau19].

2.3.2 Problems Beyond Ground State Energies

The ground state energy of a Hamiltonian is far from the only property of interest for physicists, and the Hamiltonian complexity literature reflects this. Indeed, the problem of determining many physical properties can be formulated as a computational problem and its complexity studied. We examine some examples here:

Local Observables: Ambainis considers the problem of approximating the expectation value of $\log(n)$-local measurements on ground states and low energy subspace — a natural extension beyond the Local Hamiltonian problem [Amb14]. Formally, given an observable A, one is asked whether for states $\delta = O(1/\text{poly}(N))$ close to the ground state if the expectation value $\langle A \rangle > \beta$ or $\langle A \rangle < \alpha$, for $\beta - \alpha = \Omega(1/\text{poly}(N))$. In particular, it is demonstrated that the problem is $\mathsf{P}^{\mathsf{QMA}[\log]}$-complete. Gharibian and Yirka improve on this result to $O(1)$-local Hamiltonians [GY19]. This is further improved to $\mathsf{P}^{\mathsf{QMA}[\log]}$-completeness for 1D systems, and for physically realistic systems such as the Heisenberg model using perturbative gadgets [CMP18]. Part of this thesis furthers the study of this problems to new classes of Hamiltonians (see chapter 6).

Low Energy Subspace Structure: Gharibian and Sikora examine the problem of determining if parts of the ground state subspace have an energy barrier between them [GS18]: this is important for quantum memories as for a quantum memory to remain stable over time, it must have an energy barrier between codewords. This problem can be formalised as a computational problem GSCON and was shown to be QCMA-complete for 5-local Hamiltonians, for 21-local commuting Hamiltonians [GMV16], and for 21-local stoquastic Hamiltonians [Nag+21]. The

result in [GMV16] is particularly important as most designs for error correcting codes involve commuting Hamiltonians (albeit not with 21-local interactions). In [WBG20], it was shown that even for a 2-local, nearest neighbour, translationally invariant Hamiltonian, the problem remains QCMA_{EXP}-complete (or QCMA-complete if the Hamiltonian is not translationally invariant). Furthermore, this work developed a "lifting lemma" showing that if a circuit-to-Hamiltonian mapping exists for a family of Hamiltonians (subject to certain constraints) then the GSCON problem is QCMA or QCMA_{EXP}-hard.

Spectral Gaps: The spectral gap of a Hamiltonian (the difference in energies between the first excited state and the ground state energy) is essential to determining the properties of Hamiltonians: systems which have large spectral gaps tend to be significantly easier to solve for[5]. Indeed, there are efficient approximations to states and local observables for gapped Hamiltonians in certain cases [Has07a; Has07b; LVV15; DB19].

In [Amb14] it is shown that the problem of estimating the spectral gap is contained in $\mathsf{P}^{\mathsf{QMA}[\log]}$, and is $\mathsf{P}^{\mathsf{UQMA}[\log]}$-hard. Beyond the spectral gap's importance physically, it also useful in characterising the complexity of the Local Hamiltonian problem for different parameter regimes [DGF20].

Consistency of Local Density Matrices: The Consistency of Local Density Matrices (CLDM) is another important problem which widely occurs in chemistry. Here one is given a set of reduced k-local density matrices $\{\rho_1, \rho_2 \dots \rho_m\}$, where each ρ_i describes the reduced state on some set of qudits S_i, and are asked to determine whether there is a global density matrix of a system such that $|Tr_{S_i}(\sigma) - \rho_i| < \epsilon$ where Tr_{S_i} traces out the entire system except S_i. This was shown to be QMA-complete under Turing reductions by Liu [Liu06], which was then improved to Karp reductions [BG19]. A natural variation of this is the N-representability problem, in which the Hamiltonian is restricted to only 2-local interactions between fermions, and the aim to to determine the overall density matrix describing the system. This has also been shown to be QMA-complete [LCV07]. For stoquastic Hamiltonians, the the CLDM

[5] In this paragraph we only discuss the spectral gap with respect to finite systems. We discuss the thermodynamic limit later.

problem has been shown to be StoqMA-hard for a slightly different definition of CLDM than that used in previous results [Liu07].

Time Evolution: The above properties are concerned primarily with the static properties of a Hamiltonian. Naturally one might ask what happens if we evolve a state under the time dynamics of a system? In general, predicting the time dynamics of a Hamiltonian is BQP-complete problem [Fey84], and hence we do not expect there to be classical algorithm which efficiently simulates time dynamics for quantum systems (in contrast quantum computers are capable of this [Llo96; BCK15]). One might expect that Hamiltonians for which predicting ground state properties (for example) are complex, also have difficult to determine time evolutions. However, this is not the case in general. Examples of this include the existence of Hamiltonians for which the Local Hamiltonian problem is StoqMA-complete, but the time dynamics are BQP-complete. Even more strikingly is the ferromagnetic Heisenberg model which has a 0 energy ground state, but BQP-complete time dynamics [CGW13].

Partition Functions and Free Energies The partition function and free energy give the thermal properties of a Gibbs states of some systems — that is, the properties of a system in thermal equilibrium at some temperature. For both classical and quantum systems it is well known that for sufficiently high temperatures, there exists efficient algorithms for approximating the partition function [Bar18; MH21]. Classically, it is well known that computing the partition function exactly is $\#P$-hard, and that within some relative error it is $\mathrm{BPP}^{\mathrm{NP}}$-hard in general [Sto83]. In the quantum setting, the free energy has been studied under certain conjectures, and the problem of computing the partition function has been shown to be equivalent to counting the number of witness states accepted by a QMA verifier [Bra+21]. However, the complexity of computing the free energy of a quantum Hamiltonian polynomial additive precision remains unknown[6]. A more limited version of this problem where try to compute the partition function normalised by a factor 2^{-n}, to $1/\mathrm{poly}(n)$ precision, has been shown to be DQC1-complete. Here DQC1 is a quantum complexity class which is thought to be weaker than BQP, but still stronger than P [Bra08]. However, general

[6]This corresponds to computing the partition function to relative precision.

hardness results for the non-normalised quantum case remain illusive.

Others Problems: Properties such as the energy of excited states [JGL10], determining the critical point of phase transitions [WB21], determining ground state energies of free fermion models with impurities [BG17][7], and calculating the density of states [BFS11] have all been studied from a Hamiltonian complexity perspective. We will examine some of these in this thesis.

We also, note that Hamiltonian complexity techniques have been important in formalising and understanding the idea of analogue Hamiltonian simulation. This is the idea of designing a Hamiltonian which replicates the behaviour of another Hamiltonian of interest (sometimes called the "target" Hamiltonian) [BK02; VC05; CMP18; PB20; Koh+20].

2.3.3 Hamiltonian Complexity in the Thermodynamic Limit

The problems in the previous sections we have been dealing with Hamiltonians defined on a finite number of particles. However, much of condensed matter physics is concerned with properties of Hamiltonians in the thermodynamic limit of an infinite number of particles. The thermodynamic limit is employed to study the bulk properties of materials — those which are independent from finite-size effects. The bulk properties are typically those which occur when looking at macroscopic behaviour, where the number of particles is of the order $\sim 10^{23}$. Such systems are assumed to be large enough that their behaviour is considered independent of finite-size effects. Furthermore, particular phenomena such as phase transitions are only rigorously defined in the thermodynamic limit [Sac11]. As a result it is important to study this limit theoretically.

A series of computability and complexity results have been for developed such systems. Notably, in a seminal piece of work, Cubitt, Perez-Garcia, and Wolf used techniques from Hamiltonian complexity to demonstrate that determining whether a system has a constant spectral gap, or is gapless in the thermodynamic limit is undecidable [CPGW15a; CPGW15b]. Their construction applies to 2D, nearest neighbour, translationally invariant Hamiltonians. Furthermore, their result implies a

[7] This result proves containment in QCMA, but does not prove hardness.

"Rice's Theorem" for Hamiltonian ground states: (almost) any non-trivial property of a ground state is undecidable in the thermodynamic limit. This includes determining expectation values of local observables, correlation functions, entanglement measures etc.

This undecidability of the spectral gap result was then improved to 1D systems [Bau+18b]. Building on these works, a similar set of authors proved the existence of "size-driven phase transitions": phase transitions where the driving parameter is the system size [Bau+18a]. Their system undergoes a phase transition at some finite, but arbitrarily large and uncomputable system size, demonstrating that finite-sized effects are in some sense always present. It has also been shown that thermalisation of an observable on 1D translationally invariant Hamiltonians from an initial state is undecidable [SM21]. [8]

A key consequence of these results is that a system at a finite size may have properties which *do not* necessarily reflect the properties of the system in the thermodynamic limit. Importantly, this means that using an algorithm to determine some property of an $N \times N$ system (for example) and then extrapolating to the thermodynamic limit is not guaranteed to work (or rather, in some circumstances it is guaranteed not to work!). This is often a working assumption for doing numerics in condensed matter physics, and hence this result potentially undermines this assumption.

An alternative formulation of the above paragraph is that finite-sized effects are always present, and can radically alter the behaviour of the system even at large sizes. Thus, in some sense using the thermodynamic limit as a mathematical "trick" to study the bulk properties of some medium is not always well defined.

We note, however, that there are systems which are known to have computable and well defined properties in the thermodynamic limit. There are also well known results that allow the determination of the spectral gap in the thermodynamic limit from finite sized systems. The Knabe bound shows that for frustration free

[8] Although not directly relevant to the discussion above, we note an the following result for the sake of interest: it has been proven that there exists supersymmetric quantum field theories for which determining whether or not the theory breaks supersymmetric is undecidable [Tac22].

Hamiltonians, if a spectral gap decreases slowly enough then the system must be gapped in the thermodynamic limit [Kna88]. Bounds such as these have been used to show that for frustration-free, translationally invariant, nearest neighbour spin-1/2 Hamiltonians in 1D the spectral gap problem is decidable [BG15].

Quantum Phase Transitions Though a large number of properties of Hamiltonians have been studied, a somewhat glaring gap exists in the literature: there is a limited understanding of phase transitions from a Hamiltonian complexity perspective. The undecidability of the spectral gap results [CPGW15a; Bau+18b], although hinting at results about phase transitions (since the closing of the spectral gap is necessary for a quantum phase transitions), do not actually prove results about phase transitions for reasons that will be discussed later. This is notable given that quantum phases, and quantum phase transitions are responsible for a huge amount of interesting, low-temperature physics. Indeed, a huge amount of effort has been devoted to determining phases: from a theoretical perspective the renormalisation group methods were designed to better understand exactly this problem (we will discuss this more later). Experimentally, analogue simulation has been used to probe systems which exhibit QPTs [LS17; Eba+21], however, there are limitations on how effectively some systems and properties can be simulated [CPD21]. In chapter 3 and chapter 7 we seek to improve our understanding QPTs from a complexity perspective and understand to what extent numerics and computation able to predict quantum phase transitions.

Complexity (not Computability) in the Thermodynamic Limit The Local Hamiltonian problem plays a foundational role in Hamiltonian complexity, there is little literature on the equivalent problem in the thermodynamic limit. In the thermodynamic limit, the ground state energy necessarily diverges, and hence is not a meaningful quantity, hence it is natural to ask about the energy per particle, better known as the ground state energy *density*. On a lattice $\Lambda(L \times W)$ this is the energy per particle, defined as

$$\mathcal{E}_\rho(L,W) := \frac{\lambda_0(H^{\Lambda(L\times W)})}{LW}, \tag{2.10}$$

and the quantity in the thermodynamic limit is found by taking $L, W \to \infty$. In chapter 4 we explore the complexity of computing this quantity.

2.3.4 The Renormalisation Group

So far we have focused on results demonstrating that properties of Hamiltonians are intractable or uncomputable, especially at low energies/temperatures. But of course, there are a wide range of results proving that in certain cases properties can be efficiently computed.

Renormalisation group (RG) techniques are a broad range of numerical and analytic methods which typically work by iteratively simplifying the systems at different length scales or momentum scales while preserving physical properties. The idea being that after many iterations, the small scale "microscopic" details will have been integrated out, while the observable physics will remain. Based on ideas by Wilson [Wil71; WK74] who used the techniques to solve the Kondo problem, a huge number of techniques have emerged to determine the phase of physical systems. Renormalisation group techniques and the associated ideas have become a fundamental part of condensed matter physics and high energy physics since their development.

Importantly for this thesis, RG methods have been of fundamental importance in describing phase transitions (of both the quantum and thermal kind) and associated phenomena such as universality. Given a description of a Hamiltonian in terms of the couplings between spins, it is not at all clear a phase should emerge. But by iterating an appropriate renormalisation group scheme, we can remove the necessary "irrelevant physics" while preserving important properties about the phase (e.g. whether the phase is spontaneously magnetised). This has been applied with incredible success to a huge range of systems, including the 1D Ising model where RG techniques can be used to solve the system exactly. Indeed, many of the tensor network techniques such as MPS, PEPS and MERA were originally inspired by RG methods.

Despite their important role in modern condensed matter physics, the effectiveness of RG techniques on the exotic systems that occur in Hamiltonian complexity has not been analysed. In particular, what happens when we apply RG techniques

to Hamiltonians with undecidable properties. We know they *must* fail, but it is not immediately clear how they do so, or if even one can meaningfully apply RG techniques to Hamiltonians with uncomputable properties. In chapter 5 we investigate this.

2.4 Authors and Contributions

Chapter 3 is co-authored by Johannes Bausch and Toby Cubitt, and has been published here [BCW21]. The problem of the computability of phase diagrams was initially suggested by Toby Cubitt. Most of the proof idea is my own work, although many of the details of the construction were added to, corrected and worked out by Johannes Bausch.

Chapter 4 is co-authored with Toby Cubitt. The problem of estimating the ground state energy density was initially suggested by Toby Cubitt. Most of what is in this thesis is my own work — a section about the robustness of the tiling which has significant input from Toby Cubitt have been omitted. A full version of the work can be found in [WC21], and a condensed version in [WC22].

Chapter 5 is joint work with Emilio Onorati and Toby Cubitt and has been submitted for publication. A preprint can be found here [WOC21] The idea of applying renormalisation group techniques to the uncomputable Hamiltonian of [CPGW15a] was Toby Cubitt's. Emilio Onorati's contributions are primarily in the section concerning the renormalisation of the classical tiling Hamiltonian which has been excluded from this thesis.

Chapter 6 is joint work with Johannes Bausch and Sevag Gharibian. A preprint can be found here [WBG20] The work was my idea, and initially inspired by a search for a problem which was $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$-complete, which was then further refined and generalised on by my co-authors. All authors of this work were heavily involved in proving these results. We note that only some of the results from this work have been included here.

Chapter 7 is joint work with Johannes Bausch. A preprint can be found here [WB21] The idea was inspired by a conversation with Ashley Montanaro at a visit to

the University of Cambridge. The proof idea and details were refined many times and have equal contributions from the authors. For the purposes of brevity, some of the main proofs have only been sketched out rather than being included in full.

Figures in chapters which are joint work with Johannes Bausch were designed by him, and in all other chapters are my own work.

Chapter 3

Uncomputability of Phase Diagrams

3.1 Introduction

In this part of the thesis we address the problem of determining the phase diagram of a condensed matter system in the thermodynamic limit. In particular can we find an algorithm which will determine such a thing?

As described in chapter 2, quantum phase transitions describe a large number of important physical phenomena. Examples include the 2D Hubbard model which is thought to describe the behaviour of the high-temperature superconducting cuprates but remains poorly understood [PKC15]; the structure of atomic nuclei [Elh+16]; nucleation processes in QCD [CK92]; and many more processes.

Some of these phase transitions are well understood: classic toy models of phase transitions include the 1-dimensional transverse field Ising Model which is known to have a transition from an unordered to ordered phase at a critical magnetic field strength [Sac11]. On the other hand, in general numerical simulations are computationally difficult, and may even be intractable [SMS13; SV09]. Moreover, quantum phase diagrams can be highly complex. Experimentally and computationally one of the best studied, the 2D electron gas—a model for free electrons in semiconductors—is well known to have a complex phase diagram; the system undergoes a large number of phase transitions, most notably those associated with the quantum hall effect. Indeed, the phase diagrams of such systems are known to be incredibly rich with some producing Hoftstadter butterfly patterns with an infinite number of phases

[OA01].

There are a variety of methods that are heuristically used to calculate phases, or at least predict some properties of phases, but a general algorithm for determining the phase diagrams of a system is unknown. As outlined in section 2.3.2, renormalisation group methods have played a key role in this study, but are often applied in an ad hoc way. That is, given a Hamiltonian, it is not clear whether a particular RG method should be successful. The decimation procedure which is successful for 1D classical Ising model [Car96] does not take into account factors such as entanglement, and hence one should not expect it to succeed for highly entangled quantum systems (and indeed, typically it does not).

As per definition 2.1, given a Hamiltonian $H(\varphi)$ parametrised by φ, a quantum phase transition occurs where there is a non-analytic change in the ground state energy in terms of φ [Sac11]. So a vanishing spectral gap is a necessary (though not always sufficient) condition for a phase transition to occur. Cubitt, Perez-Garcia, and Wolf [CPGW15a; CPGW15b] showed that determining whether the spectral gap is zero or some constant is undecidable. Specifically, given a (finite) description of a translationally invariant, nearest neighbour Hamiltonian on a 2D square lattice, they prove that deciding whether it has a spectral gap or not is at least as hard as solving the Halting Problem. This result was recently strengthened to the case of 1D Hamiltonians [Bau+18b].

In both the 1D and 2D undecidability constructions mentioned above, the Hamiltonian $H(\varphi)$ depends on an external parameter φ. Whether the Hamiltonian is gapped or gapless depends on the value of this parameter, with the former corresponding to a non-critical phase and the latter to a critical phase. Hence these results give Hamiltonians with highly complex phase diagrams [CPGW15b]. However, the Hamiltonian $H(\varphi)$ is not a continuous function of φ. Specifically, the Hamiltonian contains some terms which depend continuously on the value of φ, but also others which depend on the number of bits $|\varphi|$ in the binary expansion of φ. Clearly, the latter takes integer values, and has discontinuous jumps as φ varies smoothly.

This limitation significantly restricts the implications one can draw from the spectral gap undecidability results, in particular for quantum phase diagrams, which are one of the main reasons for caring about spectral gaps in first place. For example, the definition of a quantum phase transition is a non-analytic change in the ground state energy, however, if a model is allowed to have non-analytic changes in the parameters of the Hamiltonian engineering a quantum phase transition is trivial. Take the 1D transverse Ising model, and rather than having the magnetic field as a parameter, take the binary length of the magnetic field strength as a coefficient so that the Hamiltonian is $H = \sum_{\langle i,j \rangle} Z_{(i)} Z^{(j)} + |\varphi| \sum_i X_{(i)}$. It is trivial to show that the ground state can vary non-analytically moving from one value of φ ot the next.

Furthermore, it is not "natural" for the spectral gap of a Hamiltonian to depend on the length of φ's binary expression in the sense that no real system is likely to have properties like this. Finally, the constructions in [CPGW15a] do not allow anything to be said about the spectral gap if φ and $|\varphi|$ are decoupled: this means it is not possible say anything about the full phase diagram of these models, except along a disconnected collection of line-segments in the full 2D phase diagram.

It also means that the undecidability critically relies on fine-tuning the value of one parameter in the Hamiltonian to precisely the integer value that matches $|\varphi|$; for an arbitrarily small deviation from this precise value, the proof techniques cannot say anything about the spectral gap, let alone about the phase diagram. The fact that this discontinuous dependence on both φ and $|\varphi|$ is fundamental to the proof approach raises the possibility that undecidability and its consequences may not apply to the continuous families of Hamiltonians traditionally considered in condensed matter models, and may have no real consequences for understanding quantum phase diagrams, even in principle.

The result proved in this chapter directly implies that the phase diagram of systems such as this can be uncomputable: there provably does exist any procedure or algorithm for determining the phase diagram of the system, even given a complete description of the parameters of the model.

3.2 Preliminaries

The quantum many-body systems we will consider are translationally invariant, nearest-neighbour, 2D spin lattice models. Since we are interested in phase transitions— which strictly speaking can only occur in the thermodynamic limit of infinitely large lattices—we will take the thermodynamic limit by letting the lattice size L go as $L \to \infty$. An alternative definition to a non-analyticity in the ground state energy is that we can recognise phase transitions by a discontinuous change of a macroscopic observable $O_{A/B}$, Both conditions will be satisfied by our construction. The resulting Hamiltonian over the entire lattice is then

$$H^{\Lambda(L)} := \sum_{i=1}^{L}\sum_{j=1}^{L-1} h^{\text{row}}_{(i,j),(i+1,j)} + \sum_{i=1}^{L-1}\sum_{j=1}^{L} h^{\text{col}}_{(i,j),(i,j+1)}. \tag{3.1}$$

As in [CPGW15a], we then define a Hamiltonian to be gapped or gapless as:

Definition 3.1 (Gapped, from [CPGW15a])**.** *We say that $H^{\Lambda(L)}$ of Hamiltonians is gapped if there is a constant $\gamma > 0$ and a system size $L_0 \in \mathbb{N}$ such that for all $L > L_0$, $\lambda_0(H^{\Lambda(L)})$ is non-degenerate and $\Delta(H^{\Lambda(L)}) \geq \gamma$. In this case, we say that* the spectral gap is at least γ.

Definition 3.2 (Gapless, from [CPGW15a])**.** *We say that $H^{\Lambda(L)}$ is gapless if there is a constant $c > 0$ such that for all $\epsilon > 0$ there is an $L_0 \in \mathbb{N}$ so that for all $L > L_0$ any point in $[\lambda_0(H^{\Lambda(L)}), \lambda_0(H^{\Lambda(L)}) + c]$ is within distance ϵ from* $\operatorname{spec} H^{\Lambda(L)}$.

We note that not all Hamiltonians are gapped or gapless as per this definition, but these stronger definitions will help remove any ambiguities. Throughout the paper we will be using the notion of a *continuous family of Hamiltonians*, defined as:

Definition 3.3 (Continuous family of Hamiltonians)**.** *We say that a Hamiltonian $H(\varphi) = \sum_j h_j(\varphi)$ depending on a parameter $\varphi \in \mathbb{R}$, made up of a sum over local terms $h_j(\varphi)$ each acting on a local Hilbert space $\mathcal{H}$, is* continuous *if each $h_j(\varphi) : \mathbb{R} \longrightarrow \mathcal{B}(\mathcal{H})$ is a continuous function. We say that a family of Hamiltonians $\{H_i(\varphi)\}_{i\in I}$ for some index set I is a* continuous family *if each $H_i(\varphi)$ is continuous.*

3.3 Results

In this chapter we will explicitly construct a one-parameter continuous family of Hamiltonians, such that for all values $\varphi \in \mathbb{R}$ of the external parameter, the system is guaranteed to be in one of two possible phases, distinguished by an order parameter given by the ground state expectation value of a translationally-invariant macroscopic observable/order parameter $O_{A/B}$. The two phases are also distinguished by the spectral gap of the Hamiltonian H^Λ. However, determining which phase the system is in is undecidable, hence the phase diagram of the system as a function of φ is uncomputable. More precisely, we prove the following theorem:

Theorem 3.1 (Phase Diagram Uncomputability)**.** *For any given Turing Machine TM, we can construct explicitly a dimension $d \in \mathbb{N}$, $d^2 \times d^2$ matrices a, a', b, c, c' and a $d \times d$ matrix m with the following properties:*

(i) *a, c and m are diagonal with entries in $\mathbb{Z}$*

(ii) *a' is Hermitian with entries in $\mathbb{Z} + \frac{1}{\sqrt{2}}\mathbb{Z}$,*

(iii) *b has integer entries.*

(iv) *c' is Hermitian with entries in $\mathbb{Z}$.*

(v) *For any real number $\varphi \in \mathbb{R}$ and any $0 \leq \beta \leq 1$, which can be chosen arbitrarily small, setting*

$$h^{\mathrm{col}} := c + \beta c' \qquad \text{independent of } \varphi,$$
$$h^{\mathrm{row}}(\varphi) := a + \beta \left(a' + e^{\mathrm{i}\pi\varphi} b + e^{-\mathrm{i}\pi\varphi} b^\dagger \right),$$

we have $\|h^{\mathrm{row}}(\varphi)\| \leq 2$, $\|h^{\mathrm{col}}(\varphi)\| \leq 1$.

Define $H^{\Lambda(L)}$ as in eq. (3.1), and let $O_{A/B} := L^{-2} \sum_{i \in \Lambda} m_i$. Then, given $\varphi \in [2^{-\eta}, 2^{-\eta} + 2^{-\eta-\ell})$ with $\eta \in \mathbb{N}$, the following statements hold:

1. *If TM halts on input η, then for some $\ell \geq 1$, $H^\Lambda(\varphi)$ is gapless in the sense of definition 3.2, with a ground state that is critical (i.e. with algebraic decay of*

correlations), and for all eigenstates $|\Psi_B\rangle$ *with energy* $\langle\Psi_B|H^\Lambda(\varphi)|\Psi_B\rangle \leq 1$ *it holds that* $\langle\Psi_B|O_{A/B}|\Psi_B\rangle = 0$.

2. *If TM is non-halting on input* η *and* $\ell = 1$*, then* $H^\Lambda(\varphi)$ *is gapped in the sense of definition 3.1, with a unique, product ground state* $|\Psi_A\rangle$ *with* $\langle\Psi_A|O_{A/B}|\Psi_A\rangle = 1$.

Undecidability of which of the two cases pertains follows immediately from undecidability of the Halting Problem, by choosing TM to be a universal Turing Machine. For simplicity we will refer to the phases A and B determined by the value for the macroscopic observable $O_{A/B}$ as the gapped and gapless phase respectively.

As a consequence of the new Hamiltonian construction in this paper, we also obtain the following result:

Corollary 3.1. *For all* $\varphi \in [0,1]$*,* $H^\Lambda(\varphi)$ *is either in a phase with a product ground state and a spectral gap* ≥ 1*, or it is in a gapless phase with algebraic decay of correlations, where the two phases are distinguished by the expectation value of a macroscopic observable* $O_{A/B}$*. Moreover, there exists a subset* $S \subset [0,1]$ *with Borel measure* $\mu(S) > 0$*, such that even for computable* $\varphi \in S$*, determining the phase that* $H^\Lambda(\varphi)$ *is in is uncomputable.*

A less precise but simple interpretation of the above corollary is:

Corollary 3.2 (informal)**.** *The phase diagram of* $H^\Lambda(\varphi)$ *as a function of its parameter* φ *is uncomputable.*

A set of schematic phase diagrams is shown in fig. 3.1.

3.4 Proof Overview

As shown in section 2.3.1.1 it is possible to construct a quantum many-body system whose lowest-energy eigenstate represents the evolution of any desired computation [KSV02]. If we introduce a local term in the Hamiltonian that gives additional energy to any state with overlap with the halting state of the computation, we can arrange for states representing computations that halt to pick up additional energy

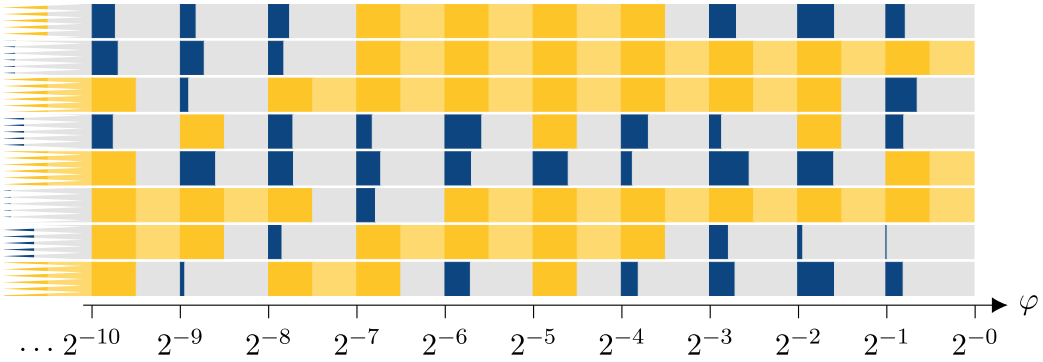

Figure 3.1: A selection of sample phase diagrams of the continuous family $\{H^{\Lambda(L)}(\varphi)\}_{L,\varphi}$ written for a series of possible universal encoded Turing machines varying from top to bottom, plotted against φ on the x-axis (note the log scaling). Blue means gapless (which is where the TM halts asymptotically on input φ), yellow gapped (TM runs forever). At the points $2^{-\eta}$ for $\eta \in \mathbb{N}$ we can have a phase transition between gapped and gapless phases, depending on the behaviour of the encoded TM; there is a positive measure interval above these points where the phase behaviour is consistent. The grey sections are parameter ranges which we do not evaluate explicitly; there will be a phase transition at some point within that region if the bounding intervals have different phases. The lighter yellow area indicates a changing gapped instance. In our construction the gapless behaviour is more intricately-dependent on φ; but the TM can be chosen such that both halting and non-halting phases cover an order one area of the phase diagram.

relative to states representing computations that do not halt, and open up a gap in the spectrum. In this way, Turing's well-known Halting Problem can be transcribed into a property of the quantum many-body system, namely whether or not it has a spectral gap. Thus determining whether the system has a spectral gap is at least as hard the Halting Problem. Since the Halting Problem is known to be undecidable, determining whether the Hamiltonian is gapped or gapless is also undecidable. Conceptually, this is how [CPGW15b; CPGW15a; Bau+18b] proved undecidability of the spectral gap.

The starting point for our construction is also undecidability of the Halting Problem [Tur37]: in brief, this states that determining whether a universal (classical) Turing machine (UTM) halts or not on a given input is, in general, undecidable. In the quantum computation setting, [CPGW15a] showed how an input can be extracted from a phase in a quantum gate such as $U = \mathrm{diag}(1, \exp(2\pi \mathrm{i}\varphi))$, using quantum phase estimation (QPE, [NC10]) which outputs a binary expansion of φ. The latter can then be fed as input to a UTM. Thus this combination of QPE and UTM runs

the universal Turing Machine on any desired input encoded in φ, and the Halting Problem for this combination is undecidable.

How do we reduce this QTM-based Halting Problem to a result about phases in a many-body system? This is a culmination of the following techniques from previous works. However, for each one of them, significant obstacles must be overcome to prove uncomputability of phase diagrams.

1. The first necessary ingredient is a QTM-to-Hamiltonian mapping which allows the construction of local, translationally-invariant couplings which result in a 1D spin chain Hamiltonian whose ground state energy is exactly zero if the encoded QTM does halts within a certain time interval; or otherwise is positive [GI09]. Using such an QTM-to-Hamiltonian mapping, a QTM running the QPE + UTM computation described above is encoded into the spin chain Hamiltonian, with φ now appearing as a parameter of the resulting Hamiltonian. However, the energy difference between the halting and non-halting cases decreases as the time interval increases, meaning we need further techniques to obtain a non-zero energy difference in the thermodynamic limit.

2. A second ingredient is amplifying this penalty. In [CPGW15a] this is done by combining the QTM-to-Hamiltonian mapping with an aperiodic tiling Hamiltonian, thereby ensuring that, for each length of computation, a fixed density of such circuit-to-Hamiltonian instances are distributed across the spin lattice. In this way, the ground state energy density is zero iff the QTM-to-Hamiltonian mapping always has zero energy, and thus depends on whether the QPE+UTM computation ever halts.

3. In [Bau+18b] point 2 is replaced by a so-called Marker Hamiltonian. This in combination with a circuit-to-Hamiltonian construction results in a ground state which partitions the spin chain into segments just large enough for the UTM to halt, if it halts. Here the segments do not have a fixed length, but instead self-adjust to find their own length. In contrast to [CPGW15a], this has the effect that either all encoded instances of computations halt, or none do.

4. The final step is the addition of an Ising-type coupling as in [CPGW15a; Bau+18a], which breaks the local Hilbert space up into subspaces $\mathcal{H}_A \oplus \mathcal{H}_B$, and which ensures that the low-energy spectrum is contained either entirely in the A or B subspace, depending on the ground state energy density just constructed. Since determining the ground state energy density is uncomputable, it is also uncomputable to determine whether the system is in phase A or B with respect to the Hamiltonian parameter φ.

As mentioned, we require significant alterations to this collection of ingredients. Concretely, the issue is that if we encode an input φ to be extracted using error-free QPE then we require the circuit gates to depend explicitly on the binary length of φ, denoted $|\varphi|$. Consequently the resulting matrix elements of the Hamiltonian will also explicitly depend on $|\varphi|$—a discontinuous function of φ. To remove this dependence we instead perform the QPE procedure approximately, by using a universal gate set that approximates all the gates depending on $|\varphi|$ [DN05a]. However, [CPGW15a; Bau+18b]'s construction crucially relies on the QPE expansion of φ to be performed exactly; any errors destroy the construction.

To overcome this obstacle, we first encode this approximate QPE plus the evolution of a UTM in a QTM-to-Hamiltonian mapping, which has positive energy iff the QPE + UTM computation does not halt. We label the resulting Hamiltonian H_{comp}. This is outlined in sections 3.4.1, 3.4.2 and 3.4.3 where the QTM-to-Hamiltonian mapping and the computation it encodes are explained, respectively. A significant novel technical contribution of this work is then a proof that the Marker Hamiltonian used in [Bau+18b] does, in fact, allow for some leeway in the precision to which QPE is performed and can be used to provide a correction for the energy penalty picked up as a result of any errors in the QPE.

To generate the required energy correction, we consider a 2D spin lattice and construct an underlying classical Hamiltonian, which we denote H_{cb}, that partitions the lattice into a uniform grid of squares. We note that the method from [CPGW15a] would be inappropriate for this construction as it would lead to an accumulation of energies we cannot correct for without matrix elements depending explicitly on

$|\varphi|$. Within each square, the ground state encodes the evolution of a classical Turing machine (encoded as a tiling problem akin to the ones used in [Ber66; Rob71]) which will calculate the energy correction necessary to offset the error introduced by approximately performing QPE. The classical Hamiltonian is then coupled to the Marker Hamiltonian. We denote this combination $H^{(\boxplus)}$.

Section 3.4.5 describes the ground state of the resulting Hamiltonian $H^{(\boxplus)}$ such that the halting or non-halting behaviour together with the Marker Hamiltonian determines whether the energy density of the constructed Hamiltonian is non-negative (in the non-halting case), or negative (in the halting case). Crucially, it is now robust with respect to the errors present in the expansion of φ from the approximate QPE procedure. Finally, in section 3.4.6 we show how $H_{\text{comp}}, H_{\text{cb}}, H^{(\boxplus)}$ are combined mathematically to lift this undecidability of the ground state energy density, to uncomputablity of the phase diagram, using now-standard techniques [CPGW15a; Bau+18a; Bau+18b]. For a mathematically rigorous derivation we refer the reader to the later sections from section 3.5 up to and including section 3.9.

3.4.1 Encoding Computation in Hamiltonians

As explain in section 2.3.1.1, it is possible to encode the evolution of a computation the ground state of a Hamiltonian. The ground state energy of the Hamiltonian can be made dependent on aspects of the computation by adding a projector that penalises certain computational states, and the resulting energy is known to high precision [BC18a; Wat19].

3.4.2 The Encoded Computation

As in [CPGW15a; Bau+18b], the computation we wish to encode via such a QTM-to-Hamiltonian mapping will be a pair of QTMs running in succession: the first will run quantum phase estimation on a quantum gate U_φ which outputs a number in binary, and the second will be a UTM which takes the output of the QPE as input. The gate U_φ is encoded in the transition unitary U describing the QTM, which is in turn encoded in the matrix elements of the Hamiltonian. The energy of the Hamiltonian encoding the computation will then be made dependent on whether the computation

halts or not, allowing us to relate its ground state energy to the halting property.

Phase Estimation. Given a unitary matrix $U_\varphi = \left(\begin{smallmatrix} 1 & 0 \\ 0 & e^{i\pi\varphi} \end{smallmatrix}\right)$, the quantum phase estimation (QPE) algorithm takes as input the eigenvector corresponding to the eigenvalue $e^{i\pi\varphi}$, and outputs an estimate of φ in binary. If the number of qubits on which the phase estimation is performed is smaller than the number of bits required to express φ in full, the algorithm is only approximate [NC10]. Furthermore, if a finite gate set is used, some of the required unitary gates in the algorithm must be approximated rather than performed exactly [DN05a]. Hence from phase estimation we get an output state consisting of a superposition over binary strings:

$$|\chi(\varphi)\rangle = \sum_{x\in\{0,1\}^n} \beta_x |x\rangle, \tag{3.2}$$

where the amplitudes β_x are concentrated around those values for which $x \approx \varphi$ and rapidly drop off away from φ. Details are in section 3.5.

Universal QTM. We then feed the output $|\chi(\varphi)\rangle$ of this phase estimation into the input of a universal Turing Machine, as in [CPGW15a], which then runs a computation which may or may not halt. By the well-known undecidability of the Halting Problem [Tur37], determining whether the QTM halts for a given string is undecidable.

3.4.3 From QTM to Hamiltonian

Using the QTM-to-Hamiltonian mapping described in section 3.4.1, the computation outlined above is mapped to a one-dimensional, translationally-invariant, nearest-neighbour Hamiltonian $H_{\text{comp}}(\varphi)$ [GI09], with a penalty for the non-halting case. It can be shown that the ground state energy of $H_{\text{comp}}(\varphi)$ scales as

$$\lambda_{\min}(H_{\text{comp}}(\varphi)) \sim \epsilon(L)/\text{poly}\, L, \tag{3.3}$$

where

$$\epsilon(L) = \sum_{x\in S(L)} |\beta_x|^2. \tag{3.4}$$

The β_x are the QPE coefficients in eq. (3.2), and $S(L)$ is the set of inputs for which the universal TM does not halt within time $T(L)$.

Since the β_x are concentrated around the binary expansion of φ, if the latter encodes a halting instance there will be a length L_0 for which $\epsilon(L) \approx 0$ for all $L > L_0$; otherwise $\epsilon(L) \approx 1$ for all L. This immediately yields a Hamiltonian for which the ground state energy is halting-dependent (and hence uncomputable). We refer the reader to section 3.5 for details.

3.4.4 Tiling and Classical Computation

In eq. (3.3) we see that the difference between the Hamiltonian's ground state energy in the case where $\epsilon(L)$ from eq. (3.4) is approximately 1 or 0 decreases with the system size L. Thus the energy gap between the two cases goes to zero irrespective of whether φ encodes a halting or non-halting instance. To amplify this gap so that there is a finite energy gap in the thermodynamic limit (as per points 2 and 3), we will combine the Feynman-Kitaev Hamiltonian with a classical Hamiltonian based on a Wang tiling that partitions the space suitably to ensure a fixed density of computation instances is spawned across the lattice. The result we achieve with this is an energy gap opening up as L grows between the cases where ϵ takes different values.

It is well known that there exist Wang tile sets that encode the evolution of a classical TM [Ber66; Rob71] within a square grid: TM tape configurations are represented by rows, such that adjacent rows represent successive time steps of the TM (fig. 3.2).

We combine both Wang tiles and the Turing Machine tiling ideas by constructing a tile set whose valid tilings have the following properties:

1. A tiling pattern that creates a square grid across the lattice Λ (much like a checkerboard). The side length of the grid squares can be varied (provided all grid squares are the same size) and still correspond to a valid tiling (depicted in fig. 3.3).

2. Within each square of the grid we use the tiles to encode a TM which first counts the size of the square it is contained in, and then outputs a marker • on

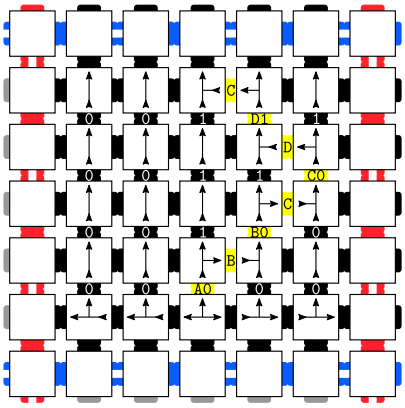

Figure 3.2: The evolution of a classical TM can be represented by Wang tiles, where colours of adjacent tiles have to match, and arrow heads have to meet arrow tails of the appropriate kind. Here the evolution runs from the bottom of the square to the top, where it places a marker • on the boundary as explained in section 3.4.5. In this image, the TM's evolution is contained in an individual square in the checkerboard grid shown in the figure below.

the top border of the square, where the placement of this marker is a function of the size of the square (depicted in in fig. 3.2).

Having developed this tiling, we map it to a corresponding tiling Hamiltonian using the mapping described in section 3.7.1, which we denote H_{cb}, such that its ground states retain the properties of the valid tilings listed above. The reader is referred to section 3.7 for details.

3.4.5 Classical Tiling with Quantum Overlay

We now want to combine the classical Hamiltonian encoding the Wang tiles, and the quantum Hamiltonian encoding the Halting Problem computation, to create an overall Hamiltonian which has a large ground state energy difference between the halting and non-halting cases, without the $\sim 1/\mathrm{poly}(L)$ decay in eq. (3.3). To do so, we split the local Hilbert space of each lattice spin into a classical part $\mathcal{H}_{\mathrm{c}}$ and a quantum part $\mathcal{H}_{\mathrm{e}} \oplus \mathcal{H}_{\mathrm{q}}$ giving $\mathcal{H} = \mathcal{H}_{\mathrm{c}} \otimes (\mathcal{H}_{\mathrm{e}} \oplus \mathcal{H}_{\mathrm{q}})$, where $\mathcal{H}_{\mathrm{e}} = \{|e\rangle\}$ just contains a filler state $|e\rangle_{\mathrm{e}}$. The ground state can then be designed to be a product state of the form $|C\rangle_{\mathrm{c}} \otimes |\psi_0\rangle_{\mathrm{eq}}$, where $|C\rangle_{\mathrm{c}}$ is a valid classical tiling configuration—as described in section 3.4.4—and $|\psi_0\rangle_{\mathrm{eq}}$ is a quantum state with the following properties:

1. We use the 1D Marker Hamiltonian from [Bau+18b], and couple its negative

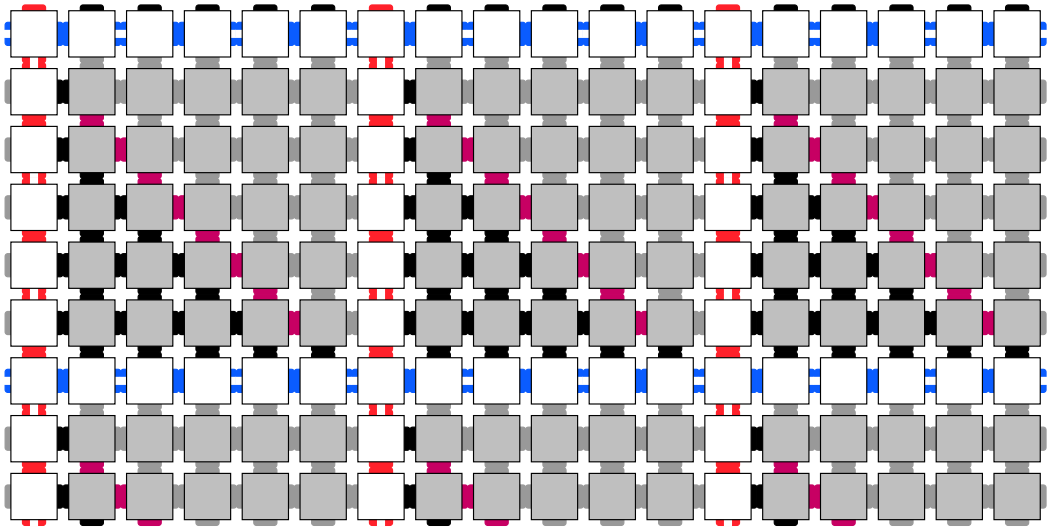

Figure 3.3: Section of the checkerboard tiling Hamiltonian's ground state. The white squares form borders, and in the interior we place tiles simulating the evolution of a classical Turing Machine.

energy contribution to the size of each grid square in the classical tiling and the placement of the • marker. The negative energy each square contributes is a determined by where the • marker is placed, and thus by the action of the classical TM. We denote this combined Hamiltonian $H^{(\boxplus)}$.

2. We effectively place the ground state of a Hamiltonian H_{comp} encoding the quantum phase estimation plus universal Turing machine along the top edge of the square, by adding additional penalty terms to the Hamiltonian that penalize the classical and quantum layers to occur in this configuration elsewhere.

3. Everywhere not along the horizontal edge of a grid square in $\mathcal{H}_{\text{c}}$ is in the zero-energy $|e\rangle_{\text{e}}$ filler state in $\mathcal{H}_{\text{e}} \oplus \mathcal{H}_{\text{q}}$.

As mentioned, the patterns in the degenerate ground space of H_{cb} are checkerboard grids of squares with periodicity $w \times w$, where the integer square size w is not fixed.

By choosing the classical TM encoded in the tiling to place a • marker at an appropriate point, we are able to tune the ground state energy of $H^{(\boxplus)}$ such that the total energy of a single $w \times w$ square A in the checkerboard pattern is:

$$\lambda_{\min}(w) := \lambda_{\min}\left(H^{(\boxplus)}|_A + H_{\text{comp}}|_A\right)\begin{cases}\geq 0 & \text{if } \epsilon(w) \geq \epsilon_0(w)\ \forall w. \\ < 0 & \text{if } \epsilon(w) < \epsilon_0(w)\ \forall w \geq w_0\end{cases} \tag{3.5}$$

where $\epsilon_0(w)$ is some cut-off point, w_0 is the halting length (recall from the previous section that the runtime of the computation encoded in the ground state depends on the size of the available tape, i.e. the size of the checkerboard square edge that the TM runs on), and where $\lambda_{\min}(w_0) = -\delta(w_0) < 0$ for the halting length w_0 is a small negative constant.

The Marker Hamiltonian's energy bonus thus compensates for the QPE approximation errors by lowering the energy by just enough such that a halting instance has negative energy. On the other hand, the energy of the non-halting instance remains large enough that the energy of a single square remains positive [Bau+18b; Wat19].

Thus, provided $\epsilon(w)$ is sufficiently small, the ground state of $H_{\text{cb}} + H_{\text{comp}} + H^{(\boxplus)}$(+coupling terms) is a checkerboard grid of squares with a constant but negative energy density. Otherwise the ground state energy density of the lattice is lower-bounded by zero. Which of the two cases holds depends on determining whether $\epsilon(w) \geq \epsilon_0(w)$ or $< \epsilon_0(w)$, which is undecidable; undecidablity of the ground state energy density follows. The reader is referred to section 3.8 for more details.

We define the Hamiltonian formed by H_{cb}, H_{comp}, $H^{(\boxplus)}$ and the coupling terms as $H_u(\varphi)$. Assume φ encodes a halting instance and set $w = \text{argmin}_s\{\lambda_{\min}(s) < 0\}$, and A is a single square of size $w \times w$. Then the ground state energy $H_u(\varphi)$ on a grid Λ of size $L \times H$ is given by

$$\lambda_{\min}(H_u(\varphi)) = \left\lfloor \frac{L}{w} \right\rfloor \left\lfloor \frac{H}{w} \right\rfloor \lambda_{\min}(H_u(\varphi)|_A). \tag{3.6}$$

3.4.6 Uncomputability of the Phase Diagram

To go from undecidability of the ground state energy density, demonstrated at the end of section 3.4.5, to undecidability of the phase (and spectral gap) we follow the approach of [CPGW15a; CPGW15b] by combining $H_u(\varphi)$ with a trivial state $|0\rangle$

such that $|0\rangle^{\otimes\Lambda}$ has zero energy, and the spectrum of $H_u(\varphi)$ is shifted up by +1 (see lemma 3.20). From this shift and eq. (3.6) it can be shown:

$$\lambda_{\min}(H_u(\varphi)) \begin{cases} \geq 1 & \text{in the non-halting case, and} \\ \longrightarrow -\infty & \text{otherwise.} \end{cases} \tag{3.7}$$

Let $h_u(\varphi)$ denote the local terms of $H_u(\varphi)$, let h_d be the local terms of the critical XY-model, and let $|0\rangle$ be a zero energy state, such that the total Hilbert space is $(\mathcal{H}_1 \otimes \mathcal{H}_2) \oplus \{|0\rangle\}$. Then the local terms of the total Hamiltonian $H^\Lambda(\varphi)$ are defined as

$$h_{i,i+1}(\varphi) \coloneqq |0\rangle\langle 0|^{(i)} \otimes (\mathbb{1} - |0\rangle\langle 0|)^{(i+1)} + (\mathbb{1} - |0\rangle\langle 0|)^{(i)} \otimes |0\rangle\langle 0|^{(i+1)} \tag{3.8}$$

$$+ h_u^{i,i+1}(\varphi) \otimes \mathbb{1}_2^{(i)} \otimes \mathbb{1}_2^{(i+1)} + \mathbb{1}_1^{(i)} \otimes \mathbb{1}_1^{(i+1)} \otimes h_d^{i,i+1}. \tag{3.9}$$

The result is the following: the overall spectrum of the Hamiltonian is

$$\operatorname{spec}(H^\Lambda(\varphi)) = \{0\} \cup (\operatorname{spec}(H_u(\varphi)) + \operatorname{spec}(H_d)) \cup G,$$

where $G \subset [1, \infty)$. To understand this, we consider the two cases. If a non-halting instance is encoded $\lambda_{\min} \geq 0$ in eq. (3.5), then $H_u(\varphi)$—from eq. (3.7)—has ground state energy lower-bounded by 1; the ground state of the overall Hamiltonian is the trivial zero energy classical product state $|0\rangle^{\otimes\Lambda}$, and H^Λ has a constant spectral gap. If $\lambda_{\min} < -|\delta|$, then $H_u(\varphi)$ has a ground state with energy diverging to $-\infty$. We further note that the critical XY-model has a dense spectrum with zero ground state energy, hence from the H_d term we obtain a dense spectrum above the ground state [LSM61]. As a result the Hamiltonian becomes gapless and has a highly entangled ground state with algebraically decaying correlations.

Since existence of a halting length w_0 in eq. (3.5) is undecidable, discriminating between $\lambda_{\min} \geq 0$ or $< -|\delta|$ is also undecidable. This implies determining whether

the Hamiltonian is in the critical, quantum-correlated phase or the trivial product state gapped phase is undecidable as well.

As $H^\Lambda(\varphi)$ is a continuous function of φ, there exist finite measure regions for which all values of φ have the same ground state and for which there is no closing of the spectral gap, which delineates the two phases. Setting $\Pi_i \coloneqq |0\rangle\langle 0|^{(i)}$, the observable $O_{A/B} = L^{-2}\sum_{i\in\Lambda}\Pi_i$ has expectation value 1 when in the state $|0\rangle^{\otimes\Lambda(L)}$, and 0 in the other case. This is true even if the observable is restricted to a finite geometrically local subset of the lattice.

We direct the reader to section 3.9 for details.

3.5 Modified Quantum Phase Estimation

We start by rigorously proving the results about quantum phase estimation (QPE) described in the proof overview.

3.5.1 The State of the Art

In [CPGW15a], the QPE on a unitary U_φ with eigenvalue $e^{i\pi\varphi}$ can output φ exactly: this is due to the fact that, in their construction, the QTM has access to a perfect gate set that is sufficient to expand precisely $|\varphi|$ digits—in particular, the standard QPE algorithm requires performing small controlled rotation gates R_n with angles $2^{i\pi 2^{-n}}$ for $n = 1,\dots,|\varphi|$, and since $|\varphi|$ is explicitly encoded in the local terms of the Hamiltonian, this circuit can be performed.

Furthermore, in [Bau+18b], one can detect when the binary expansion of φ is too long for the tape available to the QTM and penalize said segment lengths accordingly—the Marker Hamiltonian then has as a ground state a partition of the spin chain into segments of length just long enough to perform QPE on φ and for the dovetailed TM to halt—if it halts.

In our new construction the situation is fundamentally different. Since the local terms of our Hamiltonian $H^\Lambda(\varphi)$ do not explicitly depend on $|\varphi|$ any more, we cannot provide the QPE with a set of rotation gates sufficient to perform an exact quantum Fourier transform. This means that we cannot guarantee the parameter we are estimating is short enough to be written on the tape available.

We therefore have to change the construction in two key ways. First, our encoding of φ will be in unary instead of binary. Since this is a undecidability result we are not constrained by poly-time reductions—or indeed any finite computational resources; any runtime overhead is acceptable. Secondly, we will perform some gates in the QPE only approximately. The gate approximation uses standard gate synthesis algorithms from Solovay-Kitaev [DN05a], where we gear the precision of the algorithm such that it suffices to obtain a large enough certainty on the first j digits of φ, given our tape has said length. The error resulting from truncating φ to j digit is more involved, as QPE yields a superposition of states close *in value* to φ, which can for example mean that it rounds an expansion like 0.00001111 to 0.00010. We will circumvent this issue by choosing an encoding which lets us easily discover and penalize a too-short expansion, similar to the one in [Bau+18b].

3.5.2 Notation for QPE

Throughout we will denote the binary expansion of a number x as $\bar{x}$, and the first j digits of such an expansion as $\bar{x}_{\dots j}$. A question mark ? will denote a digit that can either be a 0 or a 1. The j^{th} digit of $\bar{x}$ will then be $\bar{x}_j$. For a given number x, we define $\operatorname{clz} x$ to be the count of leading zeros until the first 1 within $\bar{x}$—where we set $\operatorname{clz} 0 = \infty$. Similarly, we define the string $\operatorname{pfx} x$ to be the prefix of the string $\bar{x}$ such that $\operatorname{pfx} x = 0^{\times \operatorname{clz} x}1$, i.e. $\bar{x} = (\operatorname{pfx} x)??\dots$.

Within this section, we will further denote by U_φ a local unitary operator with eigenvalue $e^{i\pi\varphi}$, and will refer to φ as the phase to be extracted.

Finally, let $\mathcal{M}$ be a universal reversible classical TM that takes its input in unary, i.e. as a string $00\dots0100\dots$; everything past the first leading one will be ignored; we lift $\mathcal{M}$ to a quantum TM by standard procedures [BV97].

In the following analysis we first start with an encoding scheme and analyse how the approximate QPE behaves on it; we finally show that each encoded parameter φ admits a small ϵ-ball around it where the system behaves in an identical fashion, making the behaviour of gapped vs. gapless robust and showing that our family of Hamiltonians is undecidable on a non-zero-measure set over the entire parameter range $\varphi \in [0, 1]$. We do not make a claim of knowing how the construction behaves

for *any* choice of parameter. That is, given a particular value of φ, even if the halting behaviour of $\mathcal{M}$ on input $\operatorname{clz}\varphi$ were known, this would not always be sufficient to determine the behaviour of the Hamiltonian at this point.

3.5.3 Exact QPE with Truncated Expansion

We deal with the expansion error of our phase estimation first. As already mentioned, we need to choose an encoding that lets us detect and penalize expansion failure.

Definition 3.4 (Unary Encoding)**.** *Let $\eta \in \mathbb{N}$ be the input we wish to encode. Then*

$$\varphi = \varphi(\eta) := \mathtt{0.}\underbrace{\mathtt{000}\cdots\mathtt{0}}_{\eta-1 \text{ digits}}\mathtt{100}\cdots \equiv 2^{-\eta}.$$

As mentioned, it is unclear a priori how much overlap the post-QPE state has with binary strings that encode the same number in unary (i.e. the string with the same number of leading 0 digits). The benefit of using the above encoding is that phase estimation tends to *round* numbers that are too short to be expanded in full. Since we are encoding small numbers (assuming a little Endian bit order), this rounding will produce a large overlap with the all-zero state $\left|\bar{0}\right\rangle$. If we then penalise this outcome—e.g. by defining the dovetailed TM to move right forever on a zero input, which means it does not halt—we can ensure that the tape length will be extended until the input can be read in full, at which point there is no further expansion error to deal with.

As a first step we analyse the approximate quantum phase estimation procedure and compare the associated error with the perfect case, meaning that for now we give the QTM access to the same operations as in [CPGW15a] and [Bau+18b], which includes access to the unitary U_φ and rotation gates $R_n = 2^{\mathrm{i}\pi 2^{-\eta}}$, which suffice to perform phase estimation exactly. We then do the QPE algorithm identically to that laid out in [CPGW15a]; as this is the standard QPE algorithm from [NC10], we phrase the following lemma in a generic way.

Lemma 3.1. *Let $\varphi(\eta) \in \mathbb{R}$ be a unary encoding of $\eta \in \mathbb{N}$ as per definition 3.4. On t qubits of precision, QPE is performed on the unitary U_φ encoding $\varphi(\eta)$ defined in*

definition 3.4; denote the QPE output by $|\chi\rangle$. Then either:

1. $t \geq |\varphi|$, *and* $|\chi\rangle = |\bar{\varphi}\rangle$,

2. $t < |\varphi|$, *and*

$$|\chi\rangle = \sum_{x\in\{0,1\}^t} \beta_x |x\rangle \quad \textit{with} \quad |\beta_0| \geq \frac{1}{2}.$$

Proof. The first case is clear—we have a perfect gate set and sufficient tape, hence QPE is performed exactly. For the second case where $t < |\varphi|$ the β_x are given in [NC10, eq. 5.25],

$$\beta_x = \frac{1}{2^t} \frac{1-\exp\big(2\pi\mathrm{i}(2^t\varphi - (b+x))\big)}{1-\exp(2\pi\mathrm{i}(\varphi-(b+x)/2^t))}, \tag{3.10}$$

where b is the best t bit approximation to φ *less* than ϕ, i.e. $0 \leq \varphi - 2^{-t}b \leq 2^{-t}$. By definition 3.4 we have $b \equiv 0$, and therefore here

$$\beta_x = \frac{1}{2^t} \frac{1-\exp\big(2\pi\mathrm{i}(2^{t-a} - x)\big)}{1-\exp(2\pi\mathrm{i}(2^{-a}-x/2^t))}$$

and thus

$$\beta_0 = \frac{1}{2^t} \frac{1-\exp\big(2\pi\mathrm{i}2^{t-a}\big)}{1-\exp(2\pi\mathrm{i}2^{-a})} = \frac{1}{2^t} \frac{\sin\big(\pi 2^{t-a}\big)}{\sin(\pi 2^{-a})}.$$

The claim then follows from $x/2 \leq \sin(x) \leq x$ for $x \in [0, \pi/2)$. □

Corollary 3.3. *Take some $\eta \in \mathbb{N}$ and $\varphi(\eta)$ as defined in definition 3.4. Running the same quantum phase estimation QTM as in [CPGW15a] to precision m bits yields an output state $|\chi\rangle$ given in lemma 3.1, such that either*

1. $m \geq \eta$ *and* $|\chi\rangle = |\bar{\varphi}(\eta)\rangle$, *or*

2. $m < \eta$ *and* $|\langle\chi|0\rangle| \geq 1/2$.

What if $\varphi(\eta)$ is not exactly given by the encoding in definition 3.4? It is clear that $|\chi\rangle$ is still a superposition of bit strings $|x\rangle$, weighted by β_x as in eq. (3.10). But our encoding allows us to derive a variant for corollary 3.3 that applies to an interval around the correctly-encoded inputs. Here we prove that we still have a large overlap with the all zero if the phase φ is not expanded fully.

Corollary 3.4. *Let $\eta \in \mathbb{N}$, and $\varphi(\eta)$ as in definition 3.4. Take a perturbed phase $\varphi' \in [\varphi(\eta), \varphi(\eta) + 2^{-\eta-\ell})$ for some $\ell \in \mathbb{N}$, $\ell \geq 1$. Running the same quantum phase estimation QTM as in [CPGW15a] to precision m bits yields an output state $|\chi\rangle$ given in lemma 3.1, such that either*

1. *$m \geq \eta$ and $|\langle\chi|\bar{\varphi}(\eta)\rangle| \geq 1 - 2^{-\ell}$, or*
2. *$m < \eta$, and $|\langle\chi|0\rangle| \geq 1/4$.*

Proof. We start with the first case. Take β_x from eq. (3.10). Assume for now that $m = \eta$; for increasing m the overlap with $\bar{\varphi}(\eta)$ can only increase. It is clear that the best m bit approximation to φ' less than φ' is given by $b = 2^m\bar{\varphi}(\eta)$ (as the first η digits of both are identical, and $\bar{\varphi}'_{\eta+1} = 0$ by assumption). Then

$$\beta_0 = \frac{1}{2^m}\frac{1-\exp(2\pi\mathrm{i}(2^m\varphi' - b))}{1-\exp(2\pi\mathrm{i}(\varphi' - b/2^m))} = \frac{1}{2^m}\frac{\sin(\pi 2^m\epsilon)}{\sin(\pi\epsilon)} \geq 1 - 2^{-\ell},$$

where $\epsilon = \varphi' - b/2^m$, and the last inequality follows from $\sin(x)/x \geq 1 - x$ for $x \in [0, 1]$.

The second claim follows analogously: here again $b = 0$, and at most $2^m\varphi' \in [0, 3/4)$; the final bound is obtained by applying $x/4 \leq \sin(x) \leq x$ for $x \in [0, 3\pi/4)$, via

$$\frac{1}{2^m}\frac{\sin(\pi 2^m\varphi')}{\sin(\pi\varphi')} \geq \frac{1}{4}.$$

□

3.5.4 Solovay-Kitaev Modification to Phase Estimation

The second step in our QPE analysis is to approximate the small rotation gates that were previously allowed in corollary 3.3. We construct a QTM which only uses a standard gate set and U_φ for some $\varphi = \varphi(\eta) = 2^{-\eta}$, to run Quantum Phase Estimation (QPE) on U_φ and output a state which is very close in fidelity to the expansion of φ if done without error (i.e. if all gates were exact).

First note that all steps of the QPE procedure as described in [CPGW15a] can be done exactly up to applying the phase gradient and locating the least significant bit—i.e. up until Section 3.6. However, after this, controlled rotation gates of the form $R_n = 2^{\mathrm{i}\pi 2^{-n}}$, for $1 \leq n \leq |\bar{\varphi}| = \eta$, need to be applied to perform the inverse

QFT. In [CPGW15a], this was done by further giving the QTM access to the gate $2^{i\pi 2^{-\eta}}$. To circumvent this necessity, we approximate small rotation gates using the Solovay-Kitaev algorithm.

3.5.4.1 Solovay-Kitaev QTM

First we introduce the standard statement for the existence of a TM which outputs a high precision approximation to the gate $R_n = 2^{i\pi 2^{-\eta}}$ using the Solovay-Kitaev algorithm.

Lemma 3.2 (SK Machine [DN05b])**.** *There exists a classical TM which, given an integer k and maximum error ϵ, outputs an approximation $\tilde{R}_k$ to the gate $R_k \in \mathrm{SU}(2)$ such that $\|\tilde{R}_k - R_k\| < \epsilon$. The TM runs in time and space $\mathrm{O}(\log^{c_1}(1/\epsilon))$ for some $3.97 < c_1 < 4$.*

Part way through the quantum phase estimation procedure, we need to apply the inverse QFT. However, we do not have access to gates of the form $2^{i\pi 2^{-\eta}}$, and our entire QTM will be limited to space L. As a result, whenever the procedure requires a $2^{i\pi 2^{-\eta}}$-gate or a power of such a gate, we run the Solovay-Kitaev algorithm to generate an approximation. As there is $\mathrm{O}(\eta^2)$ many gates to be approximated overall, the procedure will have to be repeated this many times.

However, since we are performing the QPE on a finite length tape, we only have L qubits onto which we can write out the output of the Solovay-Kitaev algorithm; this limits the precision we can achieve using this technique.

Inverting the space bound in lemma 3.2 with respect to the error ϵ, the best approximation obtainable is thus

$$\left\|\tilde{R}_k - R_k\right\| \leq e^{-\mathrm{O}(L^{1/c_1})} \leq 2^{-c_2 L^{1/c_1}}, \tag{3.11}$$

where we wrote the constant in the exponent as c_2. Both Solovay-Kitaev constants c_1 and c_2 can be written down explicitly.

3.5.4.2 Approximation Error for Output State

The gates used in the inverse QFT in the previous section were only performed up to a finite precision and hence there will be an error associated with the output state relative to the case with perfect gates. We will see that the output is then a state that is exponentially close to what would be expected in the case with perfect gates.

Let $\tilde{R}_n$ be the approximation to the rotation gate $R_n = 2^{\pi \mathrm{i} 2^{-\eta}}$ such that $\|\tilde{R}_n - R_n\| < \epsilon$, where $\epsilon = 2^{-c_2 L^{1/c_1}}$ is given by the Solovay-Kitaev theorem, eq. (3.11) and lemma 3.2.

Lemma 3.3. *Let U_{QPE} be the unitary describing the implementation of QPE by a QTM on m qubits with each gate performed exactly. Let $\tilde{U}_{\mathrm{QPE}}$ be the unitary describing the same QPE algorithm on m qubits, but where Solovay-Kitaev is used to approximate the rotation gates R_n to precision ϵ; all other gates are implemented exactly. Then the total error of the approximate QPE is*

$$\left\|\tilde{U}_{\mathrm{QPE}} - U_{\mathrm{QPE}}\right\| < \frac{m^2}{2}\epsilon = \frac{m^2}{2} 2^{-c_2 L^{1/c_1}}. \tag{3.12}$$

Proof. The first part of the phase estimation procedure—the phase gradient operations U_{PG}—can be done exactly in both the approximate and exact cases. If QPE is performed to m qudits, we see that there are $m^2/2$ applications of R_n gates during the inverse QFT procedure. As $\tilde{U}_{\mathrm{QPE}} = \tilde{U}^{\dagger}_{\mathrm{QFT}} U_{\mathrm{PG}}$, the claim follows from applying the triangle inequality $m^2/2$ times. □

3.5.5 Total Quantum Phase Estimation Error

We have seen previously that there will be errors from both the fact that the parameter φ may have a binary expansion longer than the tape length available, and from the Solovay-Kitaev (S-K) algorithm we use to approximate and apply the rotation gates. Here we combine the two errors and upper bound the total deviation introduced. We continue using m to denote the number of binary digits that φ is expanded to, and L is the full tape length.

We emphasize that the two are not necessarily identical, as we can always cordon off a section of the tape to restrict the QPE to only work to within a more limited

precision—i.e. we can execute the QPE TM on a subsegment of size $m \leq L$ as in corollary 3.3, and approximate the latter with Solovay-Kitaev that itself can make use of the full tape space available, i.e. L. For now we treat L and m as independent quantities, regardless of how they are implemented, and we will choose their specific relation in due course.

Lemma 3.4. *Let $\eta \in \mathbb{N}$ and $\varphi(\eta) \in \mathbb{R}$ as in definition 3.4, and take $\tilde{U}_{\mathrm{QPE}}$ as the Solovay-Kitaev QPE unitary with output $|\tilde{\chi}\rangle$. Then either*

1. *$m \geq \eta$ and $|\langle\tilde{\chi}|\bar{\varphi}(\eta)\rangle| \geq 1 - \delta(L,m)$, or*
2. *$m < \eta$ and $|\langle\tilde{\chi}|0\rangle| \geq 1/2 - \delta(L,m)$.*

Here

$$\delta(L,m) < \frac{m^2}{2} 2^{-c_2 L^{1/c_1}}.$$

Proof. Immediate from lemma 3.3, eq. (3.11), and corollary 3.3. □

As before, we add an approximate variant for the case where $\varphi' \neq \varphi(\eta)$.

Lemma 3.5. *Let $\eta \in \mathbb{N}$, and $\varphi(\eta)$ as in definition 3.4. Take a perturbed phase $\varphi' \in [\varphi(\eta), \varphi(\eta) + 2^{-\eta-\ell})$ for $\ell \in \mathbb{N}$, $\ell \geq 1$, and consider the same setup as in lemma 3.4. Then either*

1. *$m \geq \eta$ and $|\langle\tilde{\chi}|\bar{\varphi}(\eta)\rangle| \geq 1 - 2^{-\ell} - \delta(L,m)$, or*
2. *$m < \eta$, and $|\langle\tilde{\chi}|0\rangle| \geq 1/4 - \delta(L,m)$.*

Proof. Analogously to lemma 3.4, but using corollary 3.4. □

The bound in terms of $\delta(L,m)$ is only useful for large L, in which case it is easy to see that since $m \leq L$, $\delta \to 0$ for $L \to 0$. Since we need δ to be small in due course, we capture a more precise bound in the following remark.

Remark 3.1. *For any $\delta_0 > 0$ there exists an $L_0 = L_0(c_1, c_2, \delta_0)$ such that $\delta(L,m) < \delta_0$ for all $L \geq L_0$, where $\delta(L,m)$ is defined in lemma 3.4, and c_1, c_2 are the Solovay-Kitaev constants from eq. (3.11).*

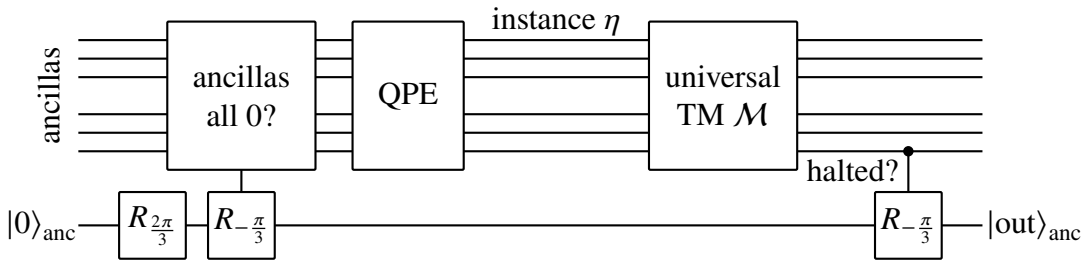

Figure 3.4: QPE and universal TM circuit. The construction uses one flag ancilla $|0\rangle_{\text{anc}}$ to verify that as many ancillas as necessary for the successive computation are correctly-initialized ancillas (e.g. $|0\rangle$), and if not rotating the single guaranteed $|0\rangle_{\text{anc}}$ flag by $\pi/3$. On some ancillas, the problem instance l is written out. Another rotation by $\pi/3$ is applied depending on whether the dovetailed universal TM $\mathcal{M}$ halts on η or not within the number of steps allowed by the clock driving its execution, which in turn is limited by the tape length.

3.6 QPE and Universal QTM Hamiltonian

In this section we examine how to encode the quantum Turing machine performing quantum phase estimation described in section 3.5 into a Hamiltonian on a spin chain of length L, such that the ground state energy of the Hamiltonian is non-negative if and only if a dovetailed universal Turing machine $\mathcal{M}$ halts on input $\varphi(\eta)$ and within tape length L. We further prove that this ground state energy remains non-negative (or negative, respectively) if instead of $\varphi(\eta)$ we are given a slightly perturbed phase $\varphi' \in [\varphi(\eta), \varphi(\eta) + 2^{-\eta-\ell})$, given $\ell \geq 1$ is large enough.

To this end, we first amend the computation slightly. In [CPGW15a], the authors used Gottesman and Irani's history state construction for a Turing machine with an initially empty tape. To ensure a correctly initialised tape, the authors use an initialization sweep; essentially a single sweep over the entire tape with a special head symbol, under which one can penalize a tape in the wrong state.

Instead of using an initialization sweep, we make do with a single ancilla (denoted with subscript "anc" in the following) which is initialized to $|0\rangle_{\text{anc}}$, and verify on a circuit level that all the other ancillas are correctly initialized. In order to achieve this, we first execute a single $R_{2\pi/3}$ rotation on $|0\rangle_{\text{anc}}$ to initialize it to a $R_{2\pi/3}|0\rangle_{\text{anc}}$-rotated state. Next, we execute a controlled $R_{-\pi/3}$ rotation in the opposite direction on $|\phi\rangle_{\text{anc}}$, where the controls are on all the ancillas we wish to ensure are in the right state. If and only if *all* of the controlling ancillas are in state $|1\rangle$—which we

can check e.g. with a multi-anticontrolled operation—will we perform a rotation by $R_{-\pi/3}$. After the controlled rotation, we apply X flips to all the ancillas we wish to initialize to $|0\rangle$.

This ancilla will carry another role: in case the dovetailed universal TM $\mathcal{M}$ from section 3.5 halts, we transition to a finalisation routine that performs another $R_{-\pi/3}$ rotation on it. The net effect of this circuit is that, after the entire computation ends, the ancilla is in state $|\text{out}\rangle_{\text{anc}}$ with overlap

$$\langle 1|\text{out}\rangle_{\text{anc}} = \begin{cases} 0 & \text{if all ancillas are correctly initialized } \textit{and } \mathcal{M} \text{ halted, or} \\ \frac{\sqrt{3}}{2} & \text{otherwise.} \end{cases} \tag{3.13}$$

This idea in the context of circuit-to-Hamiltonian mappings was introduced in [BCO17]; for completeness we give an overall circuit diagram of the entire computation to be mapped to a Hamiltonian in fig. 3.4. We remark that breaking down a multi-controlled quantum gate into a local gate set is a standard procedure described e.g. in [NC10]. We formalise the above procedure in the following lemma:

Lemma 3.6. *Consider an initial state*

$$|\psi_0\rangle = |0\rangle_{\text{anc}} \left(\alpha\,|1\rangle^{\otimes L} + \sqrt{1-\alpha^2}\,|\phi\rangle\right) \quad \textit{where} \quad |\phi\rangle \perp |1\rangle^{\otimes L}.$$

Assume the Turing machine $\mathcal{M}$ halts with probability ϵ when acting on an initial state $|0\rangle_{\text{anc}}\,|1\rangle^{\otimes L}$. Then, the final output state of the computation $|\psi_T\rangle$ satisfies

$$\langle\psi_T| \left[|1\rangle\langle 1|_{\text{anc}} \otimes \mathbb{1}^{\otimes L}\right] |\psi_T\rangle = \frac{3}{4}\left(1-\alpha^2\epsilon^2\right).$$

Proof. By explicit calculation, we have

$$
\begin{aligned}
|\psi_0\rangle &\xmapsto{R_{2\pi/3}} R_{2\pi/3}|0\rangle_{\text{anc}}\left(\alpha|1\rangle^{\otimes L}+\sqrt{1-\alpha^2}|\phi\rangle\right)\\
&\xmapsto{cR_{-\pi/3}} \alpha R_{\pi/3}|0\rangle_{\text{anc}}|1\rangle^{\otimes L}+\sqrt{1-\alpha^2}R_{2\pi/3}|0\rangle|\phi\rangle\\
&\xrightarrow{\mathcal{M}} \alpha R_{\pi/3}|0\rangle_{\text{anc}}\left(\epsilon|\psi_{\text{halt}}\rangle+\sqrt{1-\epsilon^2}|\psi_{\text{non-halt}}\rangle\right)\\
&\quad+\sqrt{1-\alpha^2}R_{2\pi/3}|0\rangle_{\text{anc}}\left(\epsilon'|\phi_{\text{halt}}\rangle+\sqrt{1-\epsilon'^2}|\phi_{\text{non-halt}}\rangle\right)\\
&\xrightarrow{cR_{-\pi/3}} \alpha\epsilon|0\rangle|\psi_{\text{halt}}\rangle+\alpha\sqrt{1-\epsilon^2}R_{\pi/3}|0\rangle|\psi_{\text{non-halt}}\rangle\\
&\quad+\epsilon'\sqrt{1-\alpha^2}R_{\pi/3}|0\rangle_{\text{anc}}|\phi_{\text{halt}}\rangle+\sqrt{1-\epsilon'^2}\sqrt{1-\alpha^2}R_{2\pi/3}|0\rangle_{\text{anc}}|\phi_{\text{non-halt}}\rangle\\
&= |\psi_T\rangle\,.
\end{aligned}
$$

Using

$$
|0\rangle \xmapsto{R_{\pi/3}} \cos\left(\frac{\pi}{3}\right)|0\rangle+\sin\left(\frac{\pi}{3}\right)|1\rangle \quad \text{and} \quad |0\rangle \xmapsto{R_{2\pi/3}} \cos\left(\frac{\pi}{3}\right)|0\rangle-\sin\left(\frac{\pi}{3}\right)|1\rangle
$$

this means that

$$
\langle\psi_T|\left[|1\rangle\langle 1|_{\text{anc}}\otimes\mathbb{1}^{\otimes L}\right]|\psi_T\rangle = \sin^2\left(\frac{\pi}{3}\right)\left(1-\alpha^2\epsilon^2\right). \qquad \square
$$

3.6.1 Feynman-Kitaev Hamiltonian

Given our quantum Turing machine from section 3.5 augmented with a single necessary "good" ancilla $|0\rangle_{\text{anc}}$ as just described, we apply the Gottesman and Irani construction from [GI09] to translate our desired computation in the ground state of a one-dimensional, nearest neighbour, translationally invariant Hamiltonian with open boundary conditions (cf. section 2.3.1.1 for a more detailed explanation). We summarize the core ideas to set up the notation used in this section, but refer the reader to [GI09; CPGW15a; BC18a] for details.

As discussed in section 3.5.5, the QPE Turing machine we devised has two meta parameters L and m. On a spin chain of length L, instead of expanding $L-3$ digits of φ as is the case in [CPGW15a], we allow the expansion to happen on a smaller sub-segment of length m of the chain. This can be done dynamically, i.e.

by adding a Turing machine before the QPE invocation which sections off a part $m = m(L)$ of the tape and places a distinct symbol Ⓘ there. Since it is obvious how to do this we will not go into detail here, and remark that in the final construction we will choose $m = L - 3$: an explicit construction for such a Turing machine is given in[Bau+18b, Lem. 15]. The QPE and dovetailed universal TM—augmented by the single-ancilla construction described at the start of this section—we will jointly call $\mathcal{M}' = \mathcal{M}'(L, m)$, i.e. such that there is L tape available; we emphasize that $\mathcal{M}'(L, m)$ has an identical set of symbols and internal states for all L and m.

In all of the following we will analyse the spectrum of the history state Hamiltonian within a "good" type of subspace, by which we mean a tape bounded by special endpoint states Ⓛ and Ⓡ. This subspace will, analogous to the 2D undecidability construction, be called *bracketed* states; on an overall local Hilbert space $\mathcal{H} = \mathcal{H}_a \oplus \mathcal{H}_b$ such that $|Ⓛ\rangle, |Ⓡ\rangle \in \mathcal{H}_b$, we set

$$\mathcal{S}_{\mathrm{br}}(m) := |Ⓛ\rangle \otimes \mathcal{H}_a^{\otimes L} \otimes |Ⓡ\rangle . \tag{3.14}$$

Since no transition rule for the history state Hamiltonian ever moves these boundary markers, the overall Hamiltonian we construct will be block-diagonal with respect to signatures determined by the brackets. A standard argument then shows that within this bracketed subspace, the history state Hamiltonian encoding the QPE Turing machine behaves as designed, and we can analyse the spectrum therein by analysing the encoded computation. Outside of the bracketed subspace, a variant of the Clairvoyance lemma allows us to always lower-bound the energy, such that it does not interfere with the rest of the construction.

In order to make all of this precise, we first define the full QPE history state Hamiltonian in the following theorem, which is adapted from [CPGW15a, Th. 10].

Theorem 3.2 (QPE history state Hamiltonian)**.** *Let $L, m \in \mathbb{N}$, $0 < m \leq L - 3$. Let there exist a Hermitian operator $h \in \mathcal{B}(\mathbb{C}^d \otimes \mathbb{C}^d)$, where the local Hilbert space contains special marker states $|Ⓛ\rangle$ and $|Ⓡ\rangle$ that define the bracketed subspace $\mathcal{S}_{\mathrm{br}}$ as in eq. (3.14), such that*

1. $h \geq 0$,
2. *d depends (at most polynomially) on the alphabet size and number of internal states of $\mathcal{M}'$,*
3. $h = A + e^{i\pi\varphi(\eta)}B + e^{-i\pi\varphi(\eta)}B^\dagger$, *where*
   - $B \in \mathcal{B}(\mathbb{C}^d \otimes \mathbb{C}^d)$ *independent of η and with coefficients in $\mathbb{Z}$, and*
   - $A \in \mathcal{B}(\mathbb{C}^d \otimes \mathbb{C}^d)$ *is Hermitian, independent of η, and with coefficients in $\mathbb{Z} + \mathbb{Z}/\sqrt{2}$;*

Furthermore, a spin chain of length L with local dimension d, the translationally-invariant nearest-neighbour Hamiltonian $H_{\mathrm{QTM}}(L) := \sum_{i=1}^{L-1} h^{(i,i+1)}$ has the following properties.

4. $H_{\mathrm{QTM}}(L)$ *is frustration-free, and*
5. *the unique ground state of $H_{\mathrm{QTM}}(L)|_{\mathcal{S}_{\mathrm{br}}(m)}$ is a computational history state as in definition 2.14 encoding the evolution of $\mathcal{M}'(L, m)$.*

The history state satisfies

6. $T = \Omega(\mathrm{poly}(L)2^L)$ *time-steps, in either the halting or non-halting case;*
7. *If $\mathcal{M}'$ runs out of tape within a time T less than the number of possible TM steps allowed by the history state clock, the computational history state only encodes the evolution of $\mathcal{M}'$ up to time T.*
8. *In either the halting or non-halting case, the remaining time steps of the evolution encoded in the history state leave the computational tape for $\mathcal{M}'$ unaltered, and instead the QTM runs an arbitary computation on a waste tape as described in [CPGW15a].*

Proof. Almost all of the above follows from [CPGW15a, Th. 10]. Item 3 differs only in that we have removed any dependence on $e^{i\pi 2^{-|\varphi|}}$ due to the new modified transition rules, as we now approximate the necessary rotations using the Solovay-Kitaev theorem (see section 3.5). □

3.6.1.1 Clock Construction

The history state Hamiltonian described above encodes an evolution of a computation for $T(L)$ steps, where $T(L)$ *does not* depend on the computation itself. This ensures that the history state will be a superposition over $T(L)$ time steps independent on whether $\mathcal{M}'$ halts on the tape of length $L-2$ and with cordonned-off subsection m. As mentioned previously, in the case of the computation halting, this is done by forcing the QTM head to switch to an additional "waste tape" where an arbitrary computation is performed until the clock finishes.

Theorem 7.9 uses the clock construction designed in [CPGW15a, sec. 4.2, 4.3, 4.4]. Bounds on the clock runtime are readily obtained: if $T(L)$ denotes the runtime of the clock on a tape of length L, we have

$$\Omega\left(L\xi^L\right) \leq T(L) \leq \mathrm{O}\left(L\xi^L \log(L)\right) \tag{3.15}$$

for some constant $\xi \in \mathbb{N}$.

3.6.1.2 QTM and Clock Combined

Theorem 7.9 combines the QTM and clock such that the QTM head only makes a transition when the oscillator from the clock part of the history state passes over the QTM head. Details can be found in [CPGW15a, sec. 4.6.1].

3.6.2 The Initialisation and Non-Halting Penalty

We now want to introduce a penalty term which will penalise computations that have not halted and not been initialised correctly.

Initialisation Penalty. In order to ensure that the single ancilla we require is correctly initialized, we introduce a projector that penalizes $|\psi\rangle_{\mathrm{anc}}$ in any state but $|0\rangle_{\mathrm{anc}}$ at the start of the computation. This can be done by a term of the form $|0\rangle\langle 0|_{\mathrm{C}} \otimes (\mathbb{1} - |1\rangle\langle 1|)_{\mathrm{anc}}$, which is local if and only if we can locally detect the initial clock state $|0\rangle_{\mathrm{C}}$ above the single ancilla on the tape. As per the constructions in [GI09; CPGW15a], this state can indeed be locally detected.

Finalisation Penalty. The final penalty follows precisely the same pattern: we add a local projector of the form $|T\rangle\langle T|_{\mathrm{C}} \otimes (\mathbb{1} - |1\rangle\langle 1|)_{\mathrm{anc}}$, and ensure that the final clock

state $|T\rangle_C$ can be recognized locally above where $|\text{out}\rangle_{\text{anc}}$ sits. To realise this, we note that the ancilla bit is located at the end of string of qudits encoding the TM tape. The final clock state can then be locally determined by a nearest-neighbour, translationally invariant term that recognises the final clock state by looking at the pair of qudits at the end of the chain. Again, this is done in [GI09; CPGW15a].

Penalty Term Construction. The amplitude of the output ancilla $|\psi\rangle_{\text{anc}}$ depends on correct initialization of the ancillas for the QTM, as well as on the halting amplitude, and is given in eq. (3.13). To penalize the overlap $\langle 1|\psi\rangle_{\text{anc}}$—which corresponds to wrong initialization, or halting—we add the following nearest neighbour term to the Hamiltonian:

$$h_{i,i+1}^{(\text{out})} = \left|[\blacksquare][\textcircled{T},\ldots,\xi]\right\rangle\!\left\langle[\blacksquare][\textcircled{T},\ldots,\xi]\right|_{i,i+1} \otimes (\mathbb{1}_i - |1\rangle\langle 1|_i) \otimes \mathbb{1}_{i+1}.$$

As just mentioned, the input penalty term $h_{i,i+1}^{(\text{in})}$ can similarly be written as a nearest-neighbour projector onto a clock state at $t = 0$. Thus, on the entire chain we have the penalty terms

$$H^{(\text{in})} = \sum_{i=1}^{L-1} h_{i,i+1}^{(\text{in})} \tag{3.16}$$

$$H^{(\text{out})} = \sum_{i=1}^{L-1} h_{i,i+1}^{(\text{out})}. \tag{3.17}$$

Definition 3.5. *We denote the QPE+QTM history state Hamiltonian including the in- and output penalties from eq.* (3.16) *with* $H_{\text{comp}}(L,\varphi) := H_{\text{QTM}}(L,\varphi) + H^{(\text{in})} + H^{(\text{out})}$.

3.6.3 Ground State Energy in Halting and Non-Halting Case

The ground state energy of H_{comp} depends on how much penalty is picked up throughout the computation. Known techniques like Kitaev's geometrical lemma [KSV02; BCO17] for a lower bound and a simple triangle inequality for the upper

bound can be used to show that

$$\Omega\left(\frac{1}{T^3}\right) \leq \lambda_{\min}(H_{\text{comp}}(\varphi(\eta))) \leq \mathrm{O}\left(\frac{1}{T}\right) \tag{3.18}$$

for a non-halting instance $\eta \in \mathbb{N}$. However, both the upper and lower bounds here are not tight enough for our purposes.

In order to obtain tighter bounds, we realize that our history state construction has a linear clock (i.e. one that never branches and simply runs from $t = 0$ to $t = T$); in this case, tight bounds on the overall energy effect of the penalty terms already exist; we refer the reader to [BC18a; CLN18; Wat19] for an extended analysis. For the sake of completeness and brevity, we quote some of the definitions and lemmas from prior literature in the appendix and reference them in the following.

Lemma 3.7. *In case $\eta \in \mathbb{N}$ correspond to a non-halting instance, the lowest eigenvalue of H_{comp} satisfies $\lambda_{\min}(H_{\text{comp}}(\varphi(\eta))) = \Omega(T^{-2})$.*

Proof. It can straightforwardly be checked that H_{comp} is a standard-form Hamiltonian as per definition A.2, and so as per the Clairvoyance Lemma [Wat19, Lem. 5.6] (and also proven in lemma A.3, appendix A.1) we know that H_{comp} breaks down into three subspaces. The subspaces of types 1 and 2 are trivially shown to have ground state energies $\Omega(T^{-2})$.

Within the third subspace, which we label S, there are no illegal terms and only the in- and output penalties $H^{(\text{in})} + H^{(\text{out})}$ from eq. (3.16) have to be considered. Again, using the Clairvoyance Lemma [Wat19, Lem. 5.6] or lemma A.3 the clock evolution within this subspace is linear—meaning there is never any branching—and hence $H_{\text{comp}}|_S$ is equivalent to Kitaev's original circuit-to-Hamiltonian construction. This means that the Hamiltonian therein is of the form

$$H_{\text{comp}}|_S = H_{\text{prop}} + |0\rangle\langle 0|_{\mathrm{C}} \otimes \Pi^{(\text{in})} + |T\rangle\langle T|_{\mathrm{C}} \otimes \Pi^{(\text{out})}$$

where $H_{\text{prop}} \sim \Delta \otimes \mathbb{1}$ for a path graph Laplacian Δ, and $\Pi^{(\text{in})/(\text{out})}$ are the in- and output penalties inflicted at time 0 and T; this Hamiltonian is then explicitly of the family of Hamiltonians studied in [BC18a]. In particular, by [BC18a, Th. 7],

Hamiltonians of this form have ground state energy $\lambda_{\min}(H_{\text{comp}}|_S) = \Omega(T^{-2})$. Thus all three of the subspaces have a minimum eigenvalue of the form $\Omega(T^{-2})$, and since they are invariant subspaces, we see that the overall minimum eigenvalue must be $\lambda_{\min}(H_{\text{comp}}) = \Omega(T^{-2})$. □

Lemma 3.8 (Theorem 6.1 from [Wat19])**.** *Let $H(\varphi) \in \mathcal{B}(\mathbb{C}^d)^{\otimes L}$ be a standard form Hamiltonian encoding a QTM with runtime $T(L)$, with in- and output penalty terms $H^{(\text{in})/(\text{out})}$. Let there exist a computational path with no illegal states such that the final state of the computation is $|\psi_T\rangle$ and such that the output penalty term satisfies*

$$\langle T| \langle \psi_T| H^{(\text{out})} |\psi_T\rangle |T\rangle \leq \epsilon.$$

Then the ground state energy is bounded by

$$0 \leq \lambda_{\min}\left(H(\varphi)\right) \leq \epsilon\left(1 - \cos\left(\frac{\pi}{2(T - T_{\text{init}}) + 1}\right)\right) = \mathrm{O}\left(\frac{\epsilon}{T^2}\right),$$

where $T_{\text{init}} = \mathrm{O}(\log(T))$ is the time frame within which the input penalty term $H^{(\text{in})}$ applies to the history state.

With this machinery developed, we can derive the following lemma for the specific Hamiltonian H_{comp} at hand.

Theorem 3.3. *Take H_{comp} to encode a phase $\varphi' \in [\varphi(\eta), \varphi(\eta) + 2^{-\ell})$, with $\varphi(\eta)$, as per definition 3.4, and let $\delta(L, m)$ be as in lemma 3.4. Then for*

1. *$m < \eta$ we have*

$$\lambda_{\min}(H_{\text{comp}}) = \Omega\left[T^{-2}\right].$$

2. *$m \geq \eta$ and $\varphi(\eta)$ corresponds to a non-halting instance, then*

$$\lambda_{\min}(H_{\text{comp}}) = \Omega\left[T^{-2}\right].$$

3. *$m \geq \eta$ and $\varphi(\eta)$ corresponds to a halting instance, then*

$$\lambda_{\min}(H_{\text{comp}}) = \mathrm{O}\left[\left(2^{-\ell} + \delta(L,m)\right)^2 \frac{1}{T^2}\right].$$

Proof. Combing lemma 3.5 with lemma 3.6 we derive upper and lower bounds on the magnitude of the amplitudes that a given instance has on a non-halting state. Together with lemma 3.7 this gives us the lower bounds for points 1 and 2. To get the upper bound in 3, by lemma 3.5 and eq. (3.13), the output penalty is bounded as

$$\langle \psi_T | \, \Pi^{(\text{out})} \, | \psi_T \rangle \leq \sin^2\left(\frac{\pi}{3}\right) \left(2^{-\ell} + \delta(L,m)\right)^2.$$

Since no other term contributes a positive energy, the ground state of H_{comp} can be upper-bounded with lemma 3.8 as

$$\lambda_{\min}(H_{\text{comp}}) = \mathrm{O}\left(\frac{\left(2^{-\ell} + \delta(L,m)\right)^2}{T^2}\right). \qquad \square$$

3.7 Checkerboard and TM Tiling

3.7.1 Wang Tiles and Hamiltonians

Wang tilings will play a central role in the construction (and the rest of this thesis).

Definition 3.6 (Wang Tiles)**.** Wang tiles *are unit length square tiles with markings on each of the four edges. For a given set of Wang tiles $\{t_i\}_{i=1}^n$, the markings define horizontal matching rules $\mathcal{R}_{Horz}$ (respectively, vertical matching rules $\mathcal{R}_{Vert}$) such that two tiles t_i, t_j can only be placed next to each other horizontally (vertically) if $(t_i, t_j) \in \mathcal{R}_{Horz}$ ($(t_i, t_j) \in \mathcal{R}_{Vert}$).*

There is then a standard technique to map a set of Wang tiles to a set of spins on a lattice such that the ground state of the spins has the same configuration as the tiling pattern when all tiling rules are satisfied (see [GI09, Section 3] or the appendix of [Bau+18a] for a summary). Consider a set of Wang tiles $\mathcal{T}$ rules with horizontal constraints $\mathcal{R}_{Horz} \subseteq \mathcal{T} \times \mathcal{T}$ such that if t_i is placed to the left of t_j, then it must be the case that $(t_i, t_j) \in \mathcal{R}_{Horz}$ and likewise for the vertical tiling rules $\mathcal{R}_{Vert}$.

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

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

$$h^{col}_{k,k+1} = \sum_{(t_i,t_j)\notin\mathcal{R}_{Horz}} |t_i t_j\rangle \langle t_i t_j|_{k,k+1} \tag{3.19}$$

$$h^{row}_{k,k+1} = \sum_{(t_i,t_j)\notin\mathcal{R}_{Vert}} |t_i t_j\rangle \langle t_i t_j|_{k,k+1}\,, \tag{3.20}$$

3.7.2 Tiling to Hamiltonian Mapping

Given a fixed set of Wang tiles on a 2D lattice, we can map the corresponding tiling pattern to a classical translationally invariant, nearest neighbour Hamiltonian over spins on the same lattice, as described in section 3.7.1.

In case we need to have more than one tile set on the same lattice, we can simply introduce lattice layers:

Remark 3.2 (Tiling Layers [GI09])**.** *For multiple tile sets $\mathcal{T}_1, \ldots, \mathcal{T}_\ell$, there exists a meta tileset $\mathcal{T}$ with a set of meta-tiling rules, such that the meta-tiling rules are only satisfied iff the tiling rule for each element of the tuple is satisfied. The corresponding Hamiltonian is defined on the tensor product of the individual Hilbert spaces. Tile constraints may also be placed between layers.*

Proof. Given a lattice, we represent the meta-tile set as an ℓ-tuple associated with each site. Each element represents a layer in the tiling. Tiling rules for the k^{th} layer are enforced between the k^{th} elements of tuples on neighbouring sites. Tiling rules between layers can be prevented from occurring by disallowing certain tuples from appearing. □

3.7.3 Checkerboard Tiling

In this section, we define a tile set that periodically tiles the infinite plane. The underlying pattern we wish to create is that of a square lattice, where each grid

cell within the pattern has the same side length, much like the boundaries on a checkerboard. The tiling will not be unique; in fact, there will be a countably infinite number of variants of the tiling which satisfy the tiling rules, corresponding to the pattern's periodicity. This non-uniquess is intended: the corresponding tiling Hamiltonian will have a degenerate ground state, the interplay of the other Hamiltonians' energy eigenstates that are conditioned on this underlying lattice pattern will then single out a unique ground state.

We constructively define this checkerboard tiling in this section. In order to explain and proof rigorously how the highest net-bonus tiling, we break the proof up into two parts; in the first part, we will create a checkerboard pattern of various square sizes, but such that the offset from the lower left corner in the lattice is left unconstrained. In the second part, we will lift this degeneracy.

Proposition 3.1 (Unconstrained Checkerboard Tiling)**.** *We define the tileset $\mathcal{T}_1$ to contain the following edge-colored tiles:*

corner & edge :

interior :

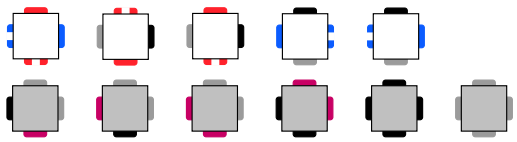

The rules for these tiles—by convention—are such that edges have to match up. Then all valid tilings for a lattice Λ will either:

1. *have no corner tile present, or*

2. *have corner tiles present as shown in fig. 3.7, i.e. such that they are part of a checkerboard pattern of squares, where the squares' side length—and the offset of the left- and bottommost corner tile—is unconstrained.*

Proof. Fig. 3.7 forms a valid tiling by inspection. What is left to prove is that given we demand at least one corner tile to be present this is the only tiling pattern possible.

To this end, we first note that the tiles directly adjacent to the corner tile are necessarily of the configuration in fig. 3.5.

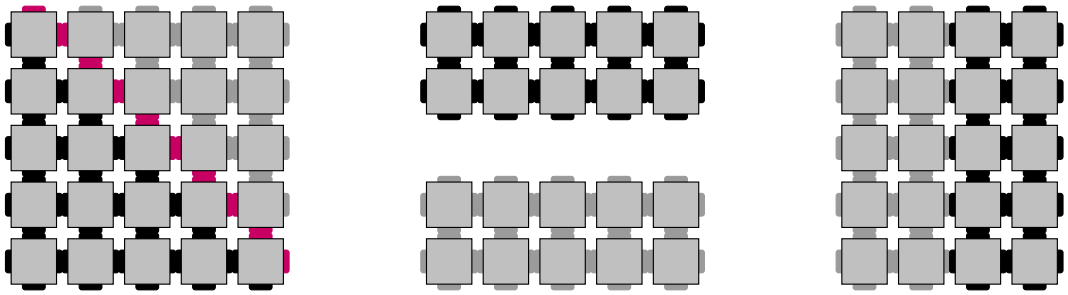

Figure 3.6: Sub-tiling patterns A_1, A_2 and A_3 from left to right, possible with at most four tiles in proposition 3.1.

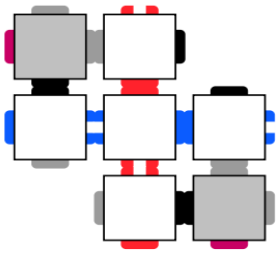

Figure 3.5: A tiling configuration.

We then note that the only way for multiple of these corner tiles to join up is via blue horizontal links (called configuration A_2), red vertical ones (configuration A_3), or diagonal purple ones (configuration A_1); we show sections of these links A_1, A_2 and A_3 in fig. 3.6.

This reduces the problem to finding valid geometric patterns of horizontal blue, vertical red and diagonal purple lines, which are only ever allowed to intersect jointly together; the resulting pattern is a grid of squares laid out by the red and blue edge tiles, where the fact that each enclosed area is a square is enforced by the purple diagonals. If the square size is bigger than the lattice Λ, this means that only a single corner tile is present; otherwise there is multiple ones, as shown in fig. 3.7. Naturally, offset and square sizes remain unconstrained; the claim follows. □

We emphasize that the tileset $\mathcal{T}$ in proposition 3.1 does have valid tilings that are e.g. all-grey- or all-black-edged areas, or those where only a purple diagonal with grey on one side, and black on the other side is present, as shown in fig. 3.6. For this reason, and in order to lift the offset degeneracy still present, we add extra constraints

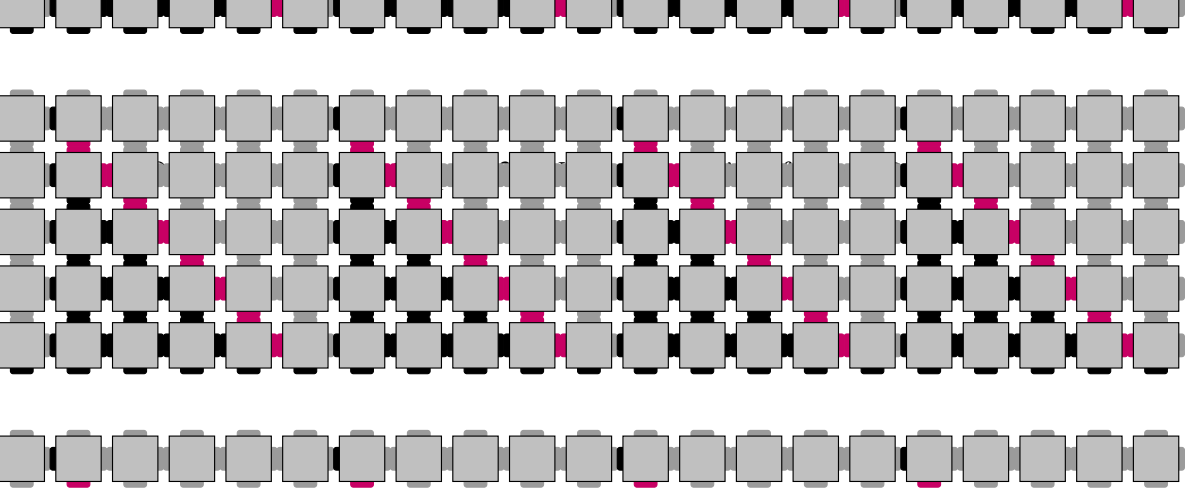

Figure 3.7: Section of the checkerboard tiling Hamiltonian's ground state.

to the tileset.

In order to single out all those patterns that commence with a full square in the lower left corner of the lattice region, we employ Gottesman and Irani's boundary trick which exploits the fact that on any hyperlattice there is always a very specific mismatch between the number of vertices and the number of edges. In our case it reads as follows.

Proposition 3.2 (Constrained Checkerboard Tiling)**.** *Take the tileset $\mathcal{T}$ from proposition 3.1 with the same edge-matching tiling rules, and define a new tileset $\mathcal{T}'$ with the following additional bonuses and penalties:*

1. *any interior tile gets a bonus of -1 if it appears to the top, and a bonus of -1 if it appears to the right of another tile, and*

2. *any interior tile gets an unconditional penalty of 2.*

Then the highest score tilings possible with $\mathcal{T}'$ on a square lattice Λ are the checkerboard patterns shown in fig. 3.7, but such that a corner tile lies in the lower left of the lattice. All other tilings have net score ≥ 1.

Proof. The only effect of the extra bonus and penalty terms are that the grey interior tiles can no longer appear on the left or bottom lattice boundaries; edge tiles have to be placed there. This, in turn, means that the only zero penalty configuration for the lower left corner is to place a corner tile there, meaning that the only net zero penalty configurations have at least one corner tile present. The rest of the claim then follows from proposition 3.1. □

With the tileset $\mathcal{T}'$ defined such that the highest net-score tilings are checkerboard patterns with unconstrained square sizes and offset $(0,0)$ from the lower left corner in the spin lattice, we can formalize the tiling Hamiltonian in the following lemma.

Lemma 3.9 (Checkerboard Tiling Hamiltonian)**.** *There exists a diagonal Hermitian operator $h \in \mathcal{B}(\mathbb{C}^d \otimes \mathbb{C}^d)$ for $d = 11$ with matrix entries in $\mathbb{Z}$ as in eq. (3.19) such that the corresponding tiling Hamiltonian $H_{\text{cb}} = \sum_{i \sim j} h^{(i,j)}$ on a square lattice Λ has a degenerate zero energy ground space S_{cb} spanned by checkerboard tilings as in fig. 3.7, of all possible square sizes, where the pattern starts with a corner tile at the origin (i.e. in the lower left corner of the lattice), as laid out in proposition 3.2. Any other eigenstate not contained in this family of zero energy states has eigenvalue ≥ 1.*

Proof. Translating the tileset $\mathcal{T}'$ from proposition 3.2 into local terms as in eq. (3.19) via [Bau+18a, Cor. 2] yields local Hamiltonian terms $h \in \mathcal{B}(\mathbb{C}^d \otimes \mathbb{C}^d)$, where d is the number of tiles in the tileset—here 11; the local terms have entries in $\mathbb{Z}$ because all the weights (bonuses and penalties) in the tileset are integers. H_{cb} will have a ground space spanned by tilings with net score 0, which we proved in proposition 3.2 to look as claimed.

Furthermore, since all other tilings must have integer net penalty, all other tiling eigenstates will have energy ≥ 1. The claim follows. □

3.7.4 Classical Turing Machine Tiling

It is well know that a classical TM which runs for time N and uses a tape of length N can be encoded in an $N \times N$ grid of tiles [Ber66]. A brief overview of how this is done is given in the following. We first recall that a TM is specified by a tuple (Σ, Q, δ) where Q is the TM state, Σ is the TM alphabet, and δ is a transition function

$$\delta : Q \times \Sigma \to Q \times \Sigma \times \{L, R\}. \tag{3.21}$$

as well as an initial state q_0, an accepting state q_a, and a blank symbol $\# \in \Sigma$. Here L, R in the transition function output tell the TM head whether to move left or right respectively.

We now take the $N \times N$ grid which we can place tiles on. We will identify the rows of the grid with the tape of the TM, where successive rows will be successive time steps. Each tile now represents a cell of the TM's tape at a given time step. We now introduce a set of tiles which encode the evolution of the TM. We will need tiles which represents every possible configuration that a cell can take (what is written in the cell, whether the TM head is there, etc.).

To encode the evolution of such a TM into a set of tiles, we introduce three types of tiles: variety 1 which is specified only by an element of Σ, variety 2 specified by $\Sigma \times Q \times \{r, l\}$, and variety 3 specified by $\Sigma \times Q \times \{R, L\}$. At position P offset from the left within a row, these tiles have the following function:

Variety 1 With marking (c), $c \in \Sigma$, the corresponding cell on the TM's tape contains c, and the TM head is *not* at position P at the corresponding time step.

Variety 2 With marking (c, q, d), $c \in \Sigma$, $q \in Q$, $d \in \{r, l\}$, the corresponding cell on the TM's tape contains c, the TM head is at position P at this time step, but has not yet overwritten the tape symbol. The TM is in state q and the TM head has just moved from the right/left of P.

Variety 3 With entry (c, q, D), $c \in \Sigma$, $q \in Q$, $D \in \{R, L\}$, the corresponding cell on the TM's tape contains c, the TM head has just moved right/left from position P where it has just overwritten the previous symbol. The TM is in state q at this time step.

As a last remark, we note that one can always dovetail multiple TM tilings, as shown by Gottesman and Irani.

Lemma 3.10 (Tiling-Layer Dovetailing [GI09])**.** *Let $\mathcal{M}_1$ and $\mathcal{M}_2$ be classical Turing machines with the same alphabet Σ such that their evolution is encoded in a tiling pattern on different tiling layers (see remark 3.2) of a rectangular grid with a border as in fig. 3.7. Then—by potentially altering the tile sets—it is possible to constrain the tiling layers at the border such that $\mathcal{M}_2$ takes the output of $\mathcal{M}_1$ as its input and continues the computation.*

Proof. If $\mathcal{M}_1$ and $\mathcal{M}_2$ are TMs, then there exists a TM M which carries out $\mathcal{M}_1$ followed by $\mathcal{M}_2$ [Tur37]. Define a tileset on each layer that corresponds to said Turing machine, such that $\mathcal{M}_1$ runs from bottom-to-top and $\mathcal{M}_2$ runs top-to-bottom on each respective layer. We now need to show that there is a way of enforcing equality of the tapes of the two tiling layers next to the boundary; then the claim follows.

Similar to remark 3.2, let the meta tile at position k be specified by a 2-tuple $T_k = (t_i, t_j)_k$. Let the set of tiles making up the border be B. Then we enforce the 2-local tiling rule that the only valid tiles that can appear next to the upper border tiles have the form $((t_i, t_i), b)$, where $b \in B$ (i.e. the tiles must have the same markings in both layers). Thus the output of $\mathcal{M}_1$ is the input of $\mathcal{M}_2$. $\mathcal{M}_2$ then continues the computation on the top layer of the grid. □

In this fashion, any Turing machine (e.g. a universal one) can be encoded in a grid of tiles, which in turn can be used to define a local Hamiltonian with a ground state that corresponds to the TM's valid evolution; given a TM tiling, this can be achieved by using the tiling-to-Hamiltonian mapping already explained. Giving due credit, we capture this mapping for TM tilings in the following lemma.

Lemma 3.11 (Berger's Turing machine Tiling Hamiltonian [Ber66])**.** *For any classical Turing machine* (Σ, Q, δ) *there exists a diagonal Hermitian operator* $h \in \mathcal{B}(\mathbb{C}^d \otimes \mathbb{C}^d)$ *for* $d = \mathrm{poly}(|\Sigma|, |Q|)$ *with matrix entries in* $\mathbb{Z}$ *as in eq.* (3.19) *such that the corresponding tiling Hamiltonian* $\sum_{i\sim j} h^{(i,j)}$ *on a square lattice* Λ *has a degenerate ground space* $S_{\mathrm{TM,tiling}}$ *containing*

1. *any tape configuration without TM head tiling the plane forward indefinitely,*

2. *a tiling pattern corresponding to valid Turing machine evolutions where the initial head is aligned on one side of the lattice and where the TM does not halt on the initial tape and space provided, and*

3. *any valid Turing machine evolution starting mid-way that does not halt within the space provided.*

Proof. See [Ber66]; the fact that the tape without head tiles the infinite plane is obvious since the tape can be initialized arbitrarily and will consistently cover the lattice, i.e. by being copied forward. If the TM's head is present in a tile, and since there is no transition *into* the initial state $q_i \in Q$ of the TM, if the initial state is present it has to reside on one side of the lattice. Similarly, if the TM halts within the space provided there is no forward transition, meaning that tiling cannot have zero energy. Finally, if neither initial nor final state are present the tiling can show a consisten Turing machine evolution starting mid-way, with the tape being copied forward, or potentially altered if the TM head passes by. □

We emphasize that the TM in lemma 3.11 does not have to be reversible. We will later lift the large degeneracy of the so-defined ground space by forcing an initial tape and head configuration; with such an initial setup, the tiling becomes unique since the forward evolution of a TM head and tape is always unambiguous.

3.7.5 Combining Checkerboard and Turing Machine Tiling

As seen in remark 3.2 and lemma 3.10, we can combine two tilesets into one, by defining the new tileset as the Cartesian product of the two. In this fashion we couple the TM tile set to apper above grey-shaded interior of the squares in the underlying checkerboard pattern from proposition 3.2; the area above the edge tiles we fill with a dummy border tile. We use this dummy border to enforce initialization of the TM's tape and head: for a tape cell *above but not to the right* of the border, the tape cell is blank. For a cell *above and to the right* of a border, we put the TM into its initial configuration q_0.

In case we need our Turing machine to run for more steps than are available on a single $L \times L$ grid, we can do so as well by introducing multiple layers as per remark 3.2.

Lemma 3.12. *Let $n_1, n_2 \in \mathbb{N}$ be constant, and take a TM tileset $\mathcal{S}$ such that the TM tiles appear over the grey-shaded interior of the checkerboard pattern in fig. 3.7. We can define a new tileset $\mathcal{S}'$ such that the TM head will start in the lower left corner on an empty tape; on a grey square of side length L it will have a tape of length $n_1 L$*

and runtime $n_2 L$ available.

Proof. Initializing the head and tape on one edge of the grey square is achieved by penalizing any other tiles from appearing there, which we can do using inter-layer constraints as in remark 3.2. Once the TM tiling reaches one end of the grey square, we can similarly copy its state to another layer with a separate TM tileset that makes it evolve in the opposite direction. This shows that one can increase the available number of time steps by another constant n_2. An even simpler argument shows that on finitely many grid cells L one can always increase the number of tape cells by a constant factor n_1, by redefining n_1 sets of separate symbols. The claims follow. □

3.7.6 Cordonning off an Edge Subsection

In this section, we show that one can define a classical TM tiling that puts a single marker on an extra layer within each checkerboard square in fig. 3.7, namely on the lower edge, and at position $x = \lceil L^{1/c} \rceil$ for any $c \in \mathbb{N}$, $c > 0$. Since we have already shown how to define a classical TM tiling to appear only within the grey shaded interior of each checkerboard sqaure (remark 3.2), how to allow constant tape and runtime overhead (lemma 3.12), and how to dovetail TM tilings (lemma 3.10), the claim is immediate from the following two lemmas.

Lemma 3.13. *Let $f : \mathbb{N} \to \mathbb{N}$ such that $1 \leq f(N) \leq N$ be computable within time $O(2^{|x|})$ where $|x|$ is the binary length of its input. Consider the checkerboard tiling constructed in proposition 3.2 such that each square has side length $N+2$ (which is measured between corner tiles). Then there exists a set of tiles which has this same checkerboard pattern, but for every corner tile, except those along the bottom edge, there is a special symbol $\bullet$ at distance $f(N)$ from the left border along the top edge.*

Proof. The proof is a variant of a construction from [GI09]. First, we add an extra tiling layer above the grey interior of the checkerboard tilings which translates the square's side length N into binary; this can be done with a counter tiling, see [GI09; Pat14] and [BP17a, sec. F2.3].

Using lemma 3.10, we then dovetail this output with a TM that computes the function $f(N)$ by taking input from the previous layer. Since this is promised to be

computable in time $O(2^{|N|}) = O(N)$, this can be done via lemma 3.12. The output of this computation is then $f(N)$ in binary.

Finally we run a binary-to-unary converting TM on the binary output of $f(N)$ by reversing the binary counter tiling in [GI09; Pat14]; this requires N steps. This leaves a marker at distance $f(N)$ along the square interior. We can then introduce a tiling rule which forces a • marker onto the edge above it. The configuration on the upper white edge of each complete square of the tiling is then as per fig. 3.8

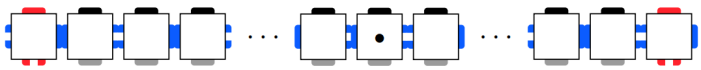

Figure 3.8: A tiling configuration.

where the black dot • marks distance $f(N)$ away from the left border. □

With lemma 3.13 in place, all that is left is to show existence of a TM that calculates the 8th root of a number given in binary, and obeys the required constraint on the number of steps—i.e. at most linear in the square's side length L.

Lemma 3.14. *Let $c \in \mathbb{N}$, $c > 0$. There exists a classical TM which, on binary input L, computes $\lceil L^{1/8} \rceil$ in binary, and requires at most $O(\log_2^8(L))$ steps.*

Proof. It is known that taking the square root of a number has the same time complexity as multiplication (see [Alt79]). For a number of $\ell = \log_2 L$ digits, long multiplication has time- and space complexity $\sim \log_2^2 L$. Taking the 8th root can thus be done in $\sim \log_2^8 L$ steps by calculating

$$\sqrt[8]{\cdot} = \sqrt{\sqrt{\sqrt{\cdot}}}$$

□

Now we have all the ingredients together to define the following augmented checkerboard Hamiltonian, which is in essence the checkerboard tiling Hamiltonian defined in lemma 3.9, but with a classical TM acting within its grey squares to place an additional marker onto the horizontal edges.

Lemma 3.15 (Augmented Checkerboard Tiling Hamiltonian)**.** *Let H_{cb} be the Hamiltonian defined in lemma 3.9. Then we can increase the local Hilbert space dimension to accommodate for the extra tileset necessary in lemma 3.14, and define a new Hamiltonian H'_{cb} as per eq. (3.19) where*

1. *the zero energy ground state is spanned by the same checkerboard patterns as in lemma 3.9, but such that the horizontal edges* above *a grey square carry a special marker • at offset $L^{1/8}$ from the left cornerstone,*

2. *any other eigenstate has eigenvalue ≥ 1.*

Proof. The first claim follows by lemmas 3.13 and 3.14, and lemma 3.11. The second claim follows since the grey TM interiors feature unique tilings, enforced by penalties only. □

For later reference, we further prove the following two tiling robustness facts.

Remark 3.3 (Checkerboard Tiling Robustness)**.** *We single out the pair of tiles in fig. 3.9,*

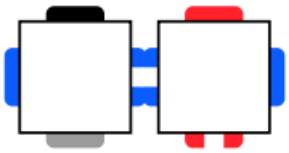

Figure 3.9: A tiling configuration.

in the tileset $\mathcal{T}'$ used in proposition 3.2. Then either

1. *the pair of tiles is part of an edge of some length L as shown in the proof of lemma 3.13—i.e. fig. 3.8—with a grey square of size $L \times L$* below *it, and a valid TM tiling enforcing the position of the extra edge marker • at position $\lceil L^{1/8} \rceil$, or*

2. *there exists a unique penalty ≥ 1 at another location in the lattice that can be associated to the tile pair.*

Proof. Since the corner tile in the pair cannot be one on the left lattice boundary, we follow the tiling to its left; it has to be a blue edge A_2 pattern as in fig. 3.6, and necessarily end in another corner tile—if not, take the mismatching tile and resulting penalty of size 1 as the unique associated one.

Given the blue horizontal edge is intact, this defines a distance between the two corner tiles, L. The subsquare $L \times L$ *below* this defined edge then has to be a valid checkerboard square with augmenting TM as in lemma 3.15, which in turn enforces the position of the • marker between the two upper corner tiles at the specified offset. If the square is not intact—which includes it being cut off—take the closest penalty in Manhattan distance from the tile pair as the associated penalty of size ≥ 1 (or one of the closest one in case of ambiguities). □

Remark 3.4 (Augmented Checkerboard Tiling Robustness)**.** *In any given ground state of the checkerboard tiling, there can be at most one • between two cornerstone markers; this marker is only ever present on blue horizontal edges that have a full grey interior square below them, meaning the • is offset at $\lceil L^{1/8} \rceil$ from its left, as in lemma 3.14. Any other configuration introduces a penalty ≥ 1.*

Proof. A bullet can only appear above the appropriate marker in the classical TM. We design the TM such that it produces exactly one such marker and such a marker gets a penalty if it is not above the point at which the TM places it. Thus, if there exists more than one • per edge joining two cornerstones, at most one of them can be above the marker left by the classical TM, and hence the other will receive an energy penalty.

Furthermore, since • can only occur at the output of a valid TM tiling, it can only occur on edges that lie above a full TM tiling. Since the lattice boundaries are white edges by proposition 3.2, the claim follows. □

3.8 A 2D Marker Checkerboard

In this section we will introduce a Hamiltonian on a one-dimensional spin chain which has a fine-tuned negative energy. More specifically, our goal in this section is to take the tiling pattern given in fig. 3.7 used to define H'_{cb} in lemma 3.15, and on a

separate layer add the Marker Hamiltonian $H^{(f)}$ from [Bau+18b, Thm. 11].

In slight extension from the construction therein, we only allow the boundary markers $|\blacksquare\rangle$ to coincide with the cornerstones of the checkerboard tiling (see fig. 3.10)

Figure 3.10: A tile.

and condition the transition terms h_1 and h_2 from [Bau+18b, Lem. 2] to only occur in between two cornerstones and if and only if the marker $\bullet$ is present there,[1] i.e. on the blue horizontal edge as per fig. 3.11.

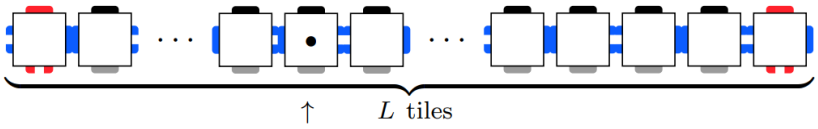

Figure 3.11: A tiling configuration.

All other configurations are energetically penalised. The negative energy contribution of one such edge—and thus by remark 3.4 also of one square *below* said edge in the checkerboard pattern—is

$$E_{\text{edge}}(L) := \lambda_{\min}(H^{(f)}|_{S(L)}), \tag{3.22}$$

where $S(L)$ denotes a single $\blacksquare$-bounded segment of the original marker construction of length L. The arrow $\uparrow$ denotes the position of the special symbol that indicates position $L^{1/8}$, as explained in lemmas 3.14 and 3.15.

As the ground state energy of $H^{(f)}$ depends on the choice of the falloff f we carefully pick this function to be able to discriminate between the halting and non-halting cases in theorem 3.3. In particular, we will choose f such that if a universal TM halts on input $\varphi(\eta)$, then $\min_L(E_{\text{edge}}(L) + \lambda_{\min}(H_{\text{comp}})) < 0$, if it

[1] This can easily be enforced with a regular expression.

does not halt then $\min_L(E_{\text{edge}}(L) + \lambda_{\min}(H_{\text{comp}})) \geq 0$, where we assumed the Turing machine's tape length is L as well. For inputs $\varphi' \in [\varphi(\eta), \varphi(\eta) + 2^{-\eta-\ell})$ for some $\ell \geq 1$ a similar condition will be true depending on the amplitude that the output state has on halting and non-halting.

One obstacle is that the bounds on the energy contribution in [Bau+18b, Lem. 7] is too loose for our purposes, i.e. it was asymptotically bounded as lying in the interval $\lambda_{\min}(H^{(f)}) \in (-2^{-f(L)}, -4^{-f(L)})$. In the following section, we prove that the scaling of the upper bound is in fact tight.

3.8.1 A Tight Marker Hamiltonian Bound

In this section we improve on the bounds set out in [Bau+18b, Lem. 7] for the ground state energy of the Marker Hamiltonian. To do this, we consider the following $w \times w$ matrix:

$$\Delta'_w = \Delta^{(w)} - |w\rangle\langle w| . \tag{3.23}$$

We now adapt [Bau+18b, Lem. 7] to prove a better lower bound on the lowest eigenvalue.

Lemma 3.16. *The minimum eigenvalue of Δ'_w satisfies*

$$\lambda_{\min}(\Delta'_w) \geq -\frac{1}{2} - \frac{3}{4^w}. \tag{3.24}$$

Proof. Our proof is essentially the same as in Bausch et al. except we use a better ansatz for the lower bound on the ground state energy. We begin by noting that, as in the proof of [Bau+18b, Lem. 7], the characteristic polynomial of Δ'_w is

$$p_w(\lambda) = -\frac{2^{-w-1}}{\sqrt{\lambda - 4}} \left(3\sqrt{\lambda}(x_w(\lambda) - y_w(\lambda)) + \sqrt{\lambda - 4}(x_w(\lambda) + y_w(\lambda)) \right) \tag{3.25}$$

where

$$x_w(\lambda) = \left(\lambda - \sqrt{\lambda - 4}\sqrt{\lambda} - 2\right)^w$$
$$y_w(\lambda) = \left(\lambda + \sqrt{\lambda - 4}\sqrt{\lambda} - 2\right)^w.$$

Since it is not clear if $p_w(\lambda) = 0$ has any closed form solutions in expressible in λ directly, we instead try to bound where the solutions can be.

First we calculate $p_w(-1/2) = (-1)^{1+w}2^{-w}$, and thus know that sgn $p_w(-1/2) = 1$ for w odd, and -1 for w even. If we can show that $p_w(-1/2 - f(w))$ has the opposite sign for some function $f(w) \geq 0$, then by the intermediate value theorem we know there has to exist a root in the interval $[-1/2 - f(w), -1/2]$. Since we are trying to prove a tighter bound than [Bau+18b, Lem. 7], we will assume $0 \leq f(w) \leq 2^{-w}$.

Let $p_w(-1/2 - f(w)) =: A_w/B_w$, where we use the notation of [Bau+18b, Lem. 7]:

$$\begin{aligned}
B_w &= 2^{w+1}\sqrt{f(w) + \frac{9}{2}}, \\
A_w &= -a_{1,w}(x'_w - y'_w) - a_{2,w}(x'_w + y'_w), \\
a_{1,w} &= 3\sqrt{f(w) + \frac{1}{2}}, \\
a_{2,w} &= \sqrt{f(w) + \frac{9}{2}}, \\
x'_w &= \left(\sqrt{f(w) + \frac{9}{2}}\sqrt{f(w) + \frac{1}{2}} - f(w) - \frac{5}{2}\right)^w, \\
y'_w &= \left(-\sqrt{f(w) + \frac{9}{2}}\sqrt{f(w) + \frac{1}{2}} - f(w) - \frac{5}{2}\right)^w.
\end{aligned}$$

Then B_w, $a_{1,w}$ and $a_{2,w}$ are real positive for all w. We distinguish two cases.

w **even.** If w is even, we need to show $p_w(-1/2-1/2^w) \geq 0$, which is equivalent to

$$\begin{aligned}
& 0 \leq \frac{A_w}{B_w} \\
\iff \quad & 0 \leq A_w = -a_{1,w}(x'_w - y'_w) - a_{2,w}(x'_w + y'_w) \\
\iff \quad & 0 \geq a(x'_w - y'_w) + (x'_w + y'_w) \qquad \text{where } a := \frac{a_{1,w}}{a_{2,w}} \in [1,2] \\
\iff \quad & \frac{a-1}{a+1} y'_w \geq x'_w.
\end{aligned}$$

For w even, $y'_w \geq x'_w$, and furthermore we find that $x'^{1/w}_w / y'^{1/w}_w$ is monotonically decreasing (assuming that $f(w) \geq 0$ and is itself monotonically decreasing), so it suffices to find a $f(w)$ which satisfies

$$\frac{a-1}{a+1} \geq \left(\frac{5}{2} - \frac{3}{2}\right)^w \left(\frac{5}{2} + \frac{3}{2}\right)^{-w} = \frac{1}{4^w}. \tag{3.26}$$

Expanding out a as

$$a = 3\sqrt{\frac{f(w)+1/2}{f(w)+9/2}}, \tag{3.27}$$

and substituting this into the above, we find

$$f(w) \geq \frac{9}{4(4^w) - 10 + 5(4^{-w})}. \tag{3.28}$$

Hence we can choose $f(w) = 3/4^w$, which works for all $w \geq 2$.

w **odd.** Now $y'_w \leq x'_w$, and it suffices to show

$$\frac{a-1}{a+1} y'_w \leq x'_w$$

which is true provided $\frac{a-1}{a+1} \leq 1$. This also holds true for all $w \geq 0$ for $f(w) = 3/4^w$. This finishes the proof. □

Theorem 3.4. *The minimum eigenvalue of Δ'_w satisfies*

$$-\frac{1}{2}-\frac{3}{4^w} \leq \lambda_{\min}(\Delta'_w) \leq -\frac{1}{2}-\frac{1}{4^w}. \tag{3.29}$$

Proof. Lemma 3.16 gives the lower bound, and [Bau+18b, Lem. 8] gives the upper bound. □

3.8.2 Balancing QPE Error and True Halting Penalty

With this tighter bound derived in theorem 3.4, we can calculate the necessary magnitude and scaling of $E_{\text{edge}}(L)$ as explained at the start of section 3.8 as follows. As a first step, we notice that the clock runtime $T = T(L)$ of the QTM is bounded by eq. (3.15), which holds both in the halting and non-halting case, since the clock idles after the computation is done. That is, the clock runtime *does not* depend on the input to the computation.

Let $E_{\text{pen,halt}}(L)$ and $E_{\text{pen,too short}}(L)$ be the ground state energies of $H_{\text{comp}}(L)$ in the case where the encoded computation does not halt with high probability, and when the binary expansion of the encoded phase is too long, respectively, i.e. when $|\varphi'| > m$. Then from theorem 3.3 we get:

$$E_{\text{pen,non-halt}}(L) \geq E_{\text{pen,too short}}(L) = \Omega\left[\frac{1}{T^2}\right] \overset{*}{\geq} \frac{K_1}{L^2 \xi^{2L} \log^2 L}, \tag{3.30}$$

where we made use of remark 3.1 at step $(*)$. Similarly, let $E_{\text{pen,halt}}(L)$ be the minimum eigenvalue when the QTM halts on input $\varphi' \in [\varphi(\eta), \varphi(\eta) + 2^{-\eta-\ell})$, as given in definition 3.4. Then again from theorem 3.3 and for sufficiently large ℓ we get:

$$\begin{aligned} E_{\text{pen,halt}}(L) &= \mathrm{O}\left[\left(2^{-\ell} + \delta(L,m)\right)\frac{1}{T^2}\right] \\ &\overset{**}{=} \mathrm{O}\left[\left(2^{-\ell} + L^2 2^{-L^{1/4}}\right)\frac{1}{T^2}\right] \leq \frac{K_2 2^{-L^{1/4}}}{\xi^{2L}}. \end{aligned} \tag{3.31}$$

where in step $(**)$ we have used the fact that $m \leq L$, $c_1 < 4$ and $c_2 \geq 1$. Both K_1 and K_2 in eqs. (3.30) and (3.31) are positive constants, chosen sufficiently small and

large to satisfy the two bounds. How large does ℓ have to be—or in other words, how small does the interval around $\varphi(\eta)$ have to be that φ' is chosen from—for eq. (3.31) to hold?

$$2^{-\ell} \leq L^2 2^{-L^{1/4}} \quad \Leftrightarrow \quad \ell \geq \log_2\left(L^{-2} 2^{L^{1/4}}\right). \tag{3.32}$$

In order to discriminate between the two asymptotic history state penalties in eqs. (3.30) and (3.31), $E_{\text{edge}}(L)$ thus has to lie asymptotically between these two bounds, i.e. we need

$$E_{\text{edge}}(L) = \mathrm{o}\left(\frac{1}{L^2 \xi^{2L} \log^2 L}\right) \quad \text{and} \quad E_{\text{edge}}(L) = \omega\left(\frac{1}{\xi^{2L} 2^{L^{1/4}}}\right).$$

Now we know by theorem 3.4 that $E_{\text{edge}}(L) \sim 4^{-f(L)}$ for some $f : \mathbb{N} \longrightarrow \mathbb{N}$ marker falloff, which itself has to be computable by a history state construction on the segment of length L. We therefore require

$$\begin{aligned} \mathrm{o}\left(\frac{1}{L^2 \xi^{2L} \log^2 L}\right) = \frac{1}{4^{f(L)}} = \omega\left(\frac{1}{\xi^{2L} 2^{L^{1/4}}}\right), \quad \text{or} \\ \omega\left(L + \log L + \log\log L\right) = f(L) = \mathrm{o}\left(L + L^{1/4}\right). \end{aligned} \tag{3.33}$$

This lets us formulate the following conclusion.

Corollary 3.5. *There exists a constant C such that $f(L) = C(L + L^{1/8})$ asymptotically satisfies eq.* (3.33).

3.8.3 Marker Hamiltonian with L + L$^{1/8}$ Falloff

The crucial question is: can we create a Marker Hamiltonian with a falloff exponent like $f(L) = C(L + L^{1/8})$, which would satisfy corollary 3.5? As discussed in [Bau+18b], this is certainly possible for any polynomial of L, or even an exponential—in essence it is a question of creating another history state clock for which the runtime of the segment of length L equals $f(L)$. Herein lies the problem: while a runtime L is easy—just have a superposition of a particle sweeping from one side to the other—how do we perform $L^{1/8}$ additional steps?

While there might be a clever way of doing this purely within the scope of a

history state construction, we take the easy way out.[2] In section 3.7.6, we discussed how we can place a special symbol on the lower edge, which by lemma 3.14 can be at distance $L^{1/8}$ from the left corner. With this in mind and with the tighter marker Hamiltonian spectral bound from lemma 3.16 to define the following variant of a marker Hamiltonian:

Lemma 3.17. *Let $C \in \mathbb{N}$ be constant. Take the standard marker Hamiltonian $H_0^{(f)}$ from [Bau+18b] defined on a local Hilbert space $\mathcal{H}_0 = \mathbb{C}^{d'}$, where d' depends on the decay function f to be implemented. Then there exists a variant $H^{(f)}$ with local Hilbert space $\mathcal{H} = \mathcal{H}_0 \otimes \mathbb{C}^2$, where $|\bigstar\rangle$ is one of the basis states of the second subspace, such that $H^{(f)}$ has the following additional properties:*

1. *$H^{(f)} = \sum_i h_i$, with $h_i \in \mathcal{B}(\mathbb{C}^d \otimes \mathbb{C}^d)$, and $d = \mathrm{O}(C)$.*

2. *$[h, |\bigstar\rangle\langle\bigstar|] = 0$.*

3. *If $S(r)$ is the subspace of a single ■-bounded segment of length L, containing a single ★ offset at position r, then*

$$-\frac{3}{4^{f(L)}} \leq \lambda_{\min}\left(H^{(f)}|_{S(r)}\right) \leq -\frac{1}{4^{f(L)}}, \tag{3.34}$$

where $f(L) = C(L+r)$.

Proof. We design the marker Hamiltonian variant to perform the following procedure before stopping:

1. Sweep the length of the edge L,

2. Sweep back to the $|\bigstar\rangle$ symbol sitting at offset r.

3. If the number of rounds is not yet C, switch to another head state and repeat, where even iterations run in reverse.

[2] We note that if this task is possible within the history state framework, then it may be possible to prove the main result of this paper for 1D. Indeed, the 2D tiling construction is only used to allow the 1D Marker Hamiltonian to have the correct drop off.

Finally, employ Gottesman and Irani's boundary trick, used as in [Bau+18b, Rem. 3], which exploits the mismatch in number of one- and two-local interaction terms to remove the constant $-1/2$ offset present in theorem 3.4 by only adding translationally-invariant nearest neighbour terms to the Hamiltonian. The energy scaling then follows directly from theorem 3.4, and the dimension and $[h_M, |\bigstar\rangle\langle\bigstar|] = 0$ follow by construction. □

This marker Hamiltonian we will now combine with the Hilbert space of the checkerboard Hamiltonian H'_{cb} from lemma 3.15, to obtain a 1D marker Hamiltonian where the location of the boundary symbols ■ and offset marker ★ align with the checkerboard tiles as

■ ⟷  and ★ ⟷  (3.35)

and such that the marker Hamiltonian terms do not occur above any other but the blue edge tiles.

Corollary 3.6 (1D Marker Hamiltonian)**.** *Let H'_{cb} be the checkerboard Hamiltonian from lemma 3.15, with local Hilbert space $\mathcal{H}_{\text{cb}}$. Take $H^{(f)}$ from lemma 3.17, with local Hilbert space $\mathcal{H}$, and let $C \in \mathbb{N}$, $C \geq 1$. Then there exists a marker Hamiltonian $H_1^{(f)}$ with one- and two-local interactions $h_1 \in \mathcal{B}(\mathcal{H}')$, $h_2 \in \mathcal{B}(\mathcal{H}' \otimes \mathcal{H}')$ where $\mathcal{H}' := (\mathcal{H} \oplus \mathbb{C}) \otimes \mathcal{H}_{\text{cb}}$, and such that $H_1^{(f)}$ has the following properties.*

1. *If $S(r)$ denotes the subspace of a good tiling edge segment fig. 3.8 of length L, where the marker $\bullet$ is offset at position r from the left, then*

$$-\frac{3}{4f(L)} \leq \lambda_{\min}\left(H_1^{(f)}|_{S(r)}\right) \leq -\frac{1}{4f(L)},$$

   with $f(L) = C(L+r)$.

2. *Restricted to any other tiling subspace S' which does not contain the pair of tiles in fig. 3.12,*

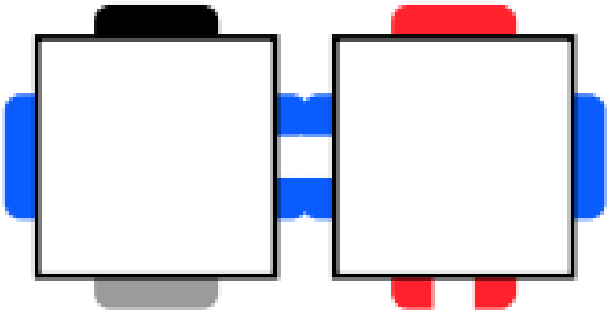

Figure 3.12: A tiling configuration.

we have $\lambda_{\min}(H_1^{(f)}|_{S'}) \geq 0$.

Proof. Let h'_1 and h'_2 denote the one- and two-local terms of $H^{(f)}$, trivially extended to the larger Hilbert space $\mathcal{H} \oplus \mathbb{C}$. Let $|0\rangle$ denote the extra basis state in $\mathcal{H} \oplus \mathbb{C}$. Denote with Π a projector onto the tiling subspace spanned by the corner and blue edge tiles given in proposition 3.1. We explicitly construct the local interactions h_1 and h_2 of $H_1^{(f)}$ by setting

$$h_1 := h'_1 \otimes \Pi + |0\rangle\langle 0| \otimes \Pi + (\mathbb{1} - |0\rangle\langle 0|) \otimes \Pi^{\perp} +$$
$$(\mathbb{1} - |\bigstar\rangle\langle\bigstar|) \otimes \left|\begin{array}{c}\square\\[-4pt]\bullet\end{array}\right\rangle\left\langle\begin{array}{c}\square\\[-4pt]\bullet\end{array}\right| + (\mathbb{1} - |\blacksquare\rangle\langle\blacksquare|) \otimes \left|\square\right\rangle\left\langle\square\right|$$

and

$$h_2 := h'_2 \otimes \Pi^{\otimes 2}.$$

The marker bonus is only ever picked up by the (final state) marker head running into the right boundary in a configuration $|\cdots \triangleright\triangleright\blacktriangleright\blacksquare\rangle$, which by the one-local Hamiltonian constraints newly imposed can only occur above the tile pair blue edge–corner given; any other configuration will have a net penalty ≥ 0. By construction, the ground space of $H_1^{(f)}$ features the required alignment from eq. (3.35). The claim then follows from lemma 3.17. □

This is the last ingredient we require to formulate a two-dimensional variant of the Marker Hamiltonian, with the required falloff from corollary 3.5.

Theorem 3.5 (2D Marker Hamiltonian)**.** *We denote with Λ the given lattice. Let h_1 and h_2 be the local terms defining the 1D marker Hamiltonian from corollary 3.6 with constant $C \in \mathbb{N}$, $C \geq 1$. Further let H'_{cb} be the augmented checkerboard lattice with symbol $\bullet$ offset by $L^{1/8}$ on each of the horizontal edges, as defined in lemma 3.15. On the joint Hilbert space we set*

$$H^{(\boxplus,f)} := \mathbb{1} \otimes H'_{\mathrm{cb}} + \sum_{i \in \Lambda} h_1^{(i)} + \sum_{i \in \Lambda} h_2^{(i)}$$

where the second sum runs over any grid index where the 2×1-sized interaction can be placed. Then the following hold:

1. *$H^{(\boxplus,f)}$ block-decomposes as $H^{(\boxplus,f)} = \bigoplus_{s=1}^{L} H_s^{(\boxplus,f)} \oplus B$; the family $H_s^{(\boxplus,f)}$ corresponds to all those tiling patterns compatible with the augmented checkerboard pattern in lemma 3.15 with square size s. B collects all other tiling configurations.*

2. *The ground state of $H_s^{(\boxplus,f)}$, labelled $|\psi_s\rangle$ is product across squares $|\psi_s\rangle = \bigotimes_i |\phi_i\rangle$, where i runs over all squares in the tiling.*

3. *$B \geq 0$.*

4. *Denote with A a single square of the ground state $|\boxplus_s\rangle$ (i.e. a square making up the grid), denoted $|\boxplus_s\rangle_A$. Then its energy contribution to the ground state of $H_s^{(\boxplus,f)}$ is*

$$-\frac{3}{4^{C(s+s^{1/8})}} \leq \langle \boxplus_s |_A \, H^{(\boxplus,f)}(s)|_A \, |\boxplus_s\rangle_A \leq -\frac{1}{4^{C(s+s^{1/8})}}.$$

   where C is the constant from corollary 3.6.

5. *Denote with $\Pi = |\boxplus_s\rangle\langle\boxplus_s|_A$ the projector onto the orthogonal complement of*

the ground state of $H_s^{(\boxplus,f)}|_A$. *Then*

$$\Pi H_s^{(\boxplus,f)}|_A \Pi \geq 0.$$

Proof. We prove the claims step by step.

Claim 1 & 2 The classical tiling Hamiltonian H'_{cb} is diagonal in the computational basis. Furthermore, by construction, h_1 and h_2 defined in corollary 3.6 commute with the tiling terms.

Claim 3 The bonus of $-1/2$ introduced in the marker Hamiltonian can only ever act across a pair of tiles

Since we have proven the checkerboard tiling to be robust with respect to the occurence of this tile pair in remark 3.3, we know that the combination carries at least a penalty ≥ 1 if it occurs in any non-checkerboard configuration; this means that any tiling in B can never have a sub-configuration such that the marker bonus offsets penalties inflicted by the tiling constraints; $B \geq 0$ follows.

Claim 4 The "good" subspace in the fourth claim we know by remark 3.4 to necessarily look as the blue edge segment fig. 3.8. This, in turn, means that $r = \lceil L^{1/8} \rceil$ in lemma 3.17, and the claim follows from the first energy bound proven therein.

Claim 5 Follows in a similar fashion as the fourth claim, from remark 3.4 and from the second claim in lemma 3.17. □

3.9 Spectral Gap Undecidability of a Continuous Family of Hamiltonians

In this section we combine the 2D Marker Hamiltonian with the QPE History State construction.

3.9.1 Uncomputability of the Ground State Energy Density

Lemma 3.18. *Let $h_1, h_2^{\text{row}}, h_2^{\text{col}}$ be the one- and two-local terms of $H^{(\boxplus,f)}$ with local Hilbert space $\mathcal{H}_{\text{m}}$, and similarly denote with q_1, q_2 be the one- and two-local terms of H_{comp} from definition 3.5 with local Hilbert space $\mathcal{H}_{\text{q}}$, respectively. Let Π_{edge} be a projector onto the edge tiles in proposition 3.1. Define the combined Hilbert space $\mathcal{H} := \mathcal{H}_{\text{m}} \otimes (\mathcal{H}_{\text{q}} \oplus \mathbb{C})$, where $|0\rangle$ denotes the basis state for the extension of $\mathcal{H}_{\text{q}}$.*

We define the following one- and two-local interactions:

$$h_1^{\text{tot}} := h_1 \otimes \mathbb{1} + \Pi_{\text{edge}} \otimes q_1 + \Pi_{\text{edge}} \otimes |0\rangle\langle 0| + \Pi_{\text{edge}}^{\perp} \otimes (1 - |0\rangle\langle 0|)$$

$$h_2^{\text{tot,row}} := h_2^{\text{row}} \otimes \mathbb{1} + \Pi_{\text{edge}}^{\otimes 2} \otimes q_2$$

$$h_2^{\text{tot,col}} := h_2^{\text{col}} \otimes \mathbb{1}$$

$p_2^{\text{tot,row}} :=$ 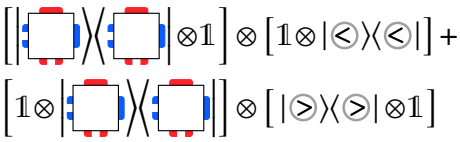

On a lattice Λ define the overall Hamiltonian

$$H := \sum_{i \in \Lambda} h_{1,(i)}^{\text{tot}} + \sum_{i \in \Lambda} \left(h_{2,(i)}^{\text{tot,row}} + p_{2,(i)}^{\text{tot,row}} \right) + \sum_{i \in \Lambda} h_{2,(i)}^{\text{tot,col}},$$

where each sum index runs over the lattice Λ where the corresponding Hamiltonian term can be placed. Then H has the following properties:

1. *$H = \bigoplus_s H_s \oplus B'$ block-decomposes as $H^{(\boxplus,f)}$ in theorem 3.5, where $B' = B \otimes \mathbb{1}$.*
2. *$B' \geq 0$.*
3. *All eigenstates of H_s are product states across squares in the tiling with square size s, product across rows within each square, and product across the local Hilbert space $\mathcal{H}_{\text{m}} \otimes (\mathcal{H}_{\text{q}} \oplus \mathbb{C})$.*
4. *Within a single square A of side length s within a block H_s, all eigenstates are of the form $|\boxplus_s\rangle|_A \otimes |r_0\rangle \otimes |r\rangle$, where*
   (a) *$|\boxplus_s\rangle$ is the ground state of the 2D marker Hamiltonian block $H_s^{(\boxplus,f)}$,*

(b) $|r_0\rangle$ *is an eigenstate of* $H_{\text{comp}} \oplus 0$*, i.e. the history state Hamiltonian with local padded Hilbert space* $\mathcal{H}_{\text{q}} \oplus \mathbb{C}$*, and*

(c) $|r\rangle \in (\mathcal{H}_{\text{q}} \oplus \mathbb{C})^{\otimes(s\times(s-1))}$ *defines the state elsewhere.*

5. *The ground state of* $H_s|_A$ *is unique and given by* $|r\rangle = |0\rangle^{\otimes(s\times(s-1))}$ *and* $|r_0\rangle = |\Psi\rangle$*, where*

$$|\Psi\rangle = \sum_{t=0}^{T-1} |t\rangle\,|\psi_t\rangle$$

is the history state of H_{comp} *as per theorem 7.9, and such that* $|\psi_0\rangle$ *is correctly initialized.*

Proof. We already have all the machinery in place to swiftly prove this lemma. First note that, by construction, all of $\{h_1^{\text{tot}}, h_2^{\text{tot,row}}, h_2^{\text{tot,col}}, h_2^{\text{col}}, p_2^{\text{tot,row}}\}$ pairwise commute with the respective tiling Hamiltonian terms $\{h_1, h_2^{\text{row}}, h_2^{\text{col}}\}$. Furthermore, the local terms from H_{comp}—q_1 and q_2—are positive semi-definite; together with theorem 3.5 this proves the first three claims. As shown in theorem 7.9 and since the Hamiltonian constraints in $p_2^{\text{tot,row}}$ enforce the ground state of the top row within the square A to be bracketed, the first and third claim imply the fourth and fifth. □

Lemma 3.19. *Take the same setup as in lemma 3.18, and let* $H_{\text{comp}} = H_{\text{comp}}(\varphi')$ *for* $\varphi' \in [\varphi(\eta), \varphi(\eta) + 2^{-\eta-\ell})$*, where* $\varphi(\eta)$ *is the unary encoding of* $\eta \in \mathbb{N}$ *from definition 3.4. As usual* $\ell \geq 1$*. Then for a block* H_s *we have*

1. *If* $s < \eta$*,* $H_s \geq 0$*.*

2. *If* $s \geq \eta$ *and* $\mathcal{M}$ *does not halt on input* η *within space* s*, then* $H_s \geq 0$*.*

3. *If* $s \geq \eta$ *and* $\mathcal{M}$ *halting on input* η*, and* $\ell \geq \log_2(s^{-2}2^{s^{1/4}})$ *as per eq. (3.32), then* $\lambda_{\min}(H_s) < 0$*.*

Proof. We start with the first claim. By lemma 3.18, it suffices to analyse a single square A of side length s; the proof then essentially follows that of [Bau+18b, Thm. 20]. We first assume $s < \eta$. Using the same notation as in theorem 3.5, and denoting

with Π_{edge} the projector onto the white horizontal edge within A, we have

$$\begin{aligned}\lambda_{\min}(H_s|_A) &= \lambda_{\min}\left[H^{(\boxplus,f)}(s)|_A \otimes \mathbb{1} + \Pi_{\text{edge}} \otimes H_{\text{comp}}(\varphi')\right] \\ &= E_{\text{edge}}(s) + E_{\text{pen,tooshort}}(s) \geq 0,\end{aligned}$$

where we used corollary 3.5 and theorem 3.5 and the fact that the two Hamiltonian terms in the sum commute.

The other claims follow equivalently: in each case by corollary 3.5, the sum of the edge bonus and TM penalties satisfy eq. (3.33). For the second claim, by the same process we thus get

$$\lambda_{\min}(H_s|_A) = E_{\text{edge}}(s) + E_{\text{pen,non-halt}}(s) \geq 0.$$

Then for the third claim,

$$\lambda_{\min}(H_s|_A) = E_{\text{edge}}(s) + E_{\text{pen,halt}}(s) < 0. \qquad \square$$

Corollary 3.7. *Take the same setup as in lemma 3.19, and let $\varphi(\eta)$ encode a halting instance. Set $w = \operatorname{argmin}_s\{\lambda_{\min}(H_s) < 0\}$, and W a single tile of size $w \times w$. Then the ground state energy of $H(\varphi')$ on a grid Λ of size $L \times H$ is bounded as*

$$\lambda_{\min}(H(\varphi')) = \left\lfloor \frac{L}{w} \right\rfloor \left\lfloor \frac{H}{w} \right\rfloor \lambda_{\min}(H(\varphi')|_W). \tag{3.36}$$

Proof. From lemma 3.18, we know the ground state of $H(\varphi')$ is a grid with offset $(0,0)$ from the lattice's origin in the lower left. Each square of the grid contributes energy $\lambda_{\min}(H(\varphi')|_W) < 0$; the prefactor in eq. (3.36) is simply the number of complete squares within the lattice.

For all truncated squares on the right hand side, H_{comp} from definition 3.5 with either the left or right ends truncated has zero ground state energy, since it is either free of the in- or output penalty terms. Furthermore, we see that if we truncate the right end of the 1D Marker Hamiltonian $H_1^{(f)}$ in lemma 3.17, it has a zero energy

ground state since it never encounters the tile pair

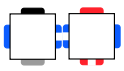

from theorem 3.5 necessary for a bonus. Truncating squares at the top does not yield any positive or negative energy contribution. The total lattice energy is therefore simply the number of complete squares on the lattice, multiplied by the energy contribution of each square. □

Theorem 3.6 (Undecidability of Ground State Energy Density)**.** *Discriminating between a negative or nonnegative ground state energy density of $H(\varphi')$ is undecidable.*

Proof. Immediate from lemma 3.19 and corollary 3.7; the energy of a single square is either a small negative constant, or nonnegative. Determining which is at least as hard as solving the halting problem. □

With this result we can almost lift the undecidability of ground state energy density to the spectral gap problem. In order to make the result slightly stronger, for this we first shift the energy of $H(\varphi')$ by a constant.

Lemma 3.20 ([Bau+18b, Lem. 23])**.** *By adding at most two-local identity terms, we can shift the energy of H from lemma 3.18 such that*

$$\lambda_{\min}(H) \begin{cases} \geq 1 & \text{in the non-halting case, and} \\ \longrightarrow -\infty & \text{otherwise.} \end{cases}$$

3.9.2 Undecidability of the Spectral Gap

With the proven uncomputability of the ground state energy density, we can lift the result using the usual ingredients—a Hamiltonian with a trivial ground state, as well as a dense spectrum Hamiltonian that will be pulled down alongside the spectrum of the QPE Hamiltonian, if the encoded universal Turing machine halts on the input encoded in the phase parameter—to prove that the existence of a spectral gap for our constructed one-parameter family of Hamiltonians is undecidable as well.

Theorem 3.7 (Undecidability of the Spectral Gap)**.** *For a continuous-parameter family of Hamiltonians, discriminating between gapped with trivial ground state* $|0\rangle^{\otimes\Lambda}$*, and gapless as defined in definitions 3.1 and 3.2, is undecidable.*

Proof. So far we have constructed a Hamiltonian $H(\varphi')$ with undecidable ground state energy asymptotics given in lemma 3.20; we denote its Hilbert space with $\mathcal{H}_1$. We add the usual Hamiltonian ingredients as in [CPGW15a] or [Bau+18b, Thm. 25]:

H_{dense} Asymptotically dense spectrum in $[0,\infty)$ on Hilbert space $\mathcal{H}_2$.

H_{trivial} Diagonal in the computational basis, with a single 0 energy product ground state $|0\rangle^{\otimes\Lambda}$, and a spectral gap of 1 (i.e. all other eigenstates have nonnegative energy ≥ 0); its Hilbert space we denote with $\mathcal{H}_3$.

H_{guard} A 2-local Ising type interaction on $\mathcal{H} := \mathcal{H}_1 \otimes \mathcal{H}_2 \oplus \mathcal{H}_3$ defined as

$$H_{\text{guard}} := \sum_{i\sim j} \left(\mathbb{1}_{1,2}^{(i)} \otimes \mathbb{1}_{3}^{(j)} + \mathbb{1}_{3}^{(i)} \otimes \mathbb{1}_{1,2}^{(j)} \right),$$

where the summation runs over all neighbouring spin sites of the underlying lattice Λ (horizontal and vertical).

We then define

$$H^{\Lambda(L)}(\varphi') := H(\varphi') \otimes \mathbb{1}_2 \oplus 0_3 + \mathbb{1}_1 \otimes H_{\text{dense}} \oplus 0_3 + 0_{1,2} \oplus H_{\text{trivial}} + H_{\text{guard}}.$$

The guard Hamiltonian ensures that any state with overlap both with $\mathcal{H}_1 \otimes \mathcal{H}_2$ and $\mathcal{H}_3$ will incur a penalty ≥ 1. It is then straightforward to check that the spectrum of H_{tot} is given by

$$\text{spec}(H^{\Lambda}) = \{0\} \cup (\text{spec}(H(\varphi')) + \text{spec}(H_{\text{dense}})) \cup G$$

for some $G \subset [1,\infty)$, where the single zero energy eigenstate stems from H_{trivial}.

In case that $\lambda_{\min}(H(\varphi')) \geq 1$, $\text{spec}(H(\varphi')) + \text{spec}(H_{\text{dense}}) \subset [1,\infty)$ and hence the ground state of H^{Λ} is the ground state of H_{trivial} with a spectral gap of size one.

For $\lambda_{\min}(H(\varphi')) \longrightarrow -\infty$, H_{dense} is asymptotically gapless and dense; this means that H^{Λ} becomes asymptotically gapless as well. □

Since the spectral properties of $H(\varphi')$ are—by lemma 3.19—robust to a choice of φ' within an interval around an encoded instance $\varphi(\eta)$ as per definition 3.4—i.e. for large enough ℓ we can vary $\varphi' \in [\varphi(\eta), \varphi(\eta) + 2^{-\eta-\ell})$—theorem 7.5 immediately proves theorem 3.1 and corollaries 3.1 and 3.2.

The Hamiltonian we have used to prove this result had phases distinguishable by an order $O_{A/B}$. However, we can adapt the Hamiltonian to change between any two phases which have different properties. For example, let H_X and $H_{\neg}$ be Hamiltonians with zero ground state energy such that their grounds states have property X and not X respectively in the thermodynamic limit (and assume $H(\varphi')$ has property X). Then determining whether the ground state of the Hamiltonian:

$$H^{\Lambda(L)}(\varphi') := H(\varphi') \otimes \mathbb{1}_2 \oplus 0_3 + \mathbb{1}_1 \otimes H_X \oplus 0_3 + 0_{1,2} \oplus H_{\neg X} + H_{\text{guard}}.$$

has property X in the thermodynamic limit is undecidable. Thus the phase diagram between these two phases is uncomputable.

3.10 Discussion

One of the main aims of this work was as a first foray into the study of the complexity of phase transitions. Quantum phase transitions are one of the best studied, but poorly understood, physical phenomena. We envision this work can be extended in several directions:

Uncomputablity in 1D. Here we have only studied phase diagrams in 2D. As described in 3.8, our construction has relied on the fact we can encode a classical Turing Machine into 2D tilings. This is not possible in 1D. However, since 1D systems tend to be fundamentally easier to solve than 2D systems, it may still be the case that the phase diagram of a 1D system is computable. However, given the undecidability of the spectral gap in 1D [Bau+18b], it would not be unexpected that computing the phase diagram in 1D is also uncomputable.

More Realistic Systems. The complexity of determining the critical value of φ at which the quantum phase transition occurs. This work has shown it is in general undecidable, but for more physically realistic systems—for example those with smaller Hilbert space dimension—does this remain the case?

Finite Systems. In this work we have only studied phase diagrams in the thermodynamic limit; naturally, those cannot occur in reality. Yet for any finite-sized system, determining any property is necessarily decidable (as we can simply diagonalise the Hamiltonian). A natural question is thus what we can say about the complexity of determining phases and phase parameters for finite system sizes, for a suitable notion of phase transitions in this context. We do not know the limits for which the properties of condensed matter systems become decidable. However, the study of these limits has potentially far-reaching consequences for high-energy physics and quantum chemistry, among other areas. We explore an element of this later in chapter 7.

Chapter 4

The Computational Complexity of the Ground State Energy Density Problem

4.1 Introduction

In section 2.3 of the introduction of this thesis we reviewed a number of Hamiltonian complexity problems which characterised how difficult it is to determine properties of Hamiltonians beyond the ground state energy. The input to all of these problems is (usually) a description of a local Hamiltonian on a finite number of particles, and the complexity-theoretic hardness is a function of varying the Hamiltonian.

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

Furthermore, they are typically interested in computing the physical properties of a *single* fixed Hamiltonian. Often, the local interactions have some regular structure, such as translational invariance where all the local interactions take the same form.

The standard formulation of the Local Hamiltonian problem does not capture this type of question — every different instance corresponds to a different Hamiltonian – rather than a fixed Hamiltonian.

Examples of these cases are the Fermi-Hubbard model (believed to be a model of high temperature superconductivity) or Lattice Quantum Chromodynamics (a model of the gluon-quark interactions). We wish to extract the properties both of these *fixed* Hamiltonians in the thermodynamic limit. The Local Hamiltonian problem and other problems where the Hamiltonian varies in the input do not capture the complexity of this.

Related Work There are a small number of results proving hardness of estimating the ground state energy for a translationally invariant Hamiltonian where the local interaction is fixed, and the only input to the problem is the lattice size. Here, since a lattice of size 2^n can be specified in n bits, the natural complexity class is NEXP (or $\mathsf{QMA}_{\mathsf{EXP}}$ in the quantum case), rather than NP. The Wang tiling completion problem is known to be NEXP-complete [Pap94; GI09], which can trivially be translated to the ground state energy problem for a single, fixed, translationally invariant, nearest-neighbour, classical Hamiltonian on a 2D square lattice, where the state at some of the boundaries is fixed (fixed boundary conditions). As the interaction is fixed, the only remaining problem input is the size of the lattice. Remarkably, this alone suffices for the hardness result. Gottesman and Irani [GI09] also extended these results to more natural types of boundary condition. They went on to prove the analogous $\mathsf{QMA}_{\mathsf{EXP}}$-completeness result for families of quantum Hamiltonians on a 1D chain. However, these results still concern Hamiltonians on finite numbers of particles; indeed, the problem input is the number of particles the Hamiltonian acts on.

In the thermodynamic limit, the ground state energy is no longer a meaningful quantity; it typically has infinite magnitude, and is not physically measurable. In this setting, the more relevant quantity is the ground state energy *density*: the minimum energy *per particle*. Just as the ground state energy is a key starting point for studying the physics of finite many-body systems, the ground state energy density (GSED) is a

key starting point for physics in the thermodynamic limit. Methods of approximating the ground state energy density in condensed matter systems have been the subject of much study in the physics literature [Per+92; HW94].

Less is known about the computational complexity of the ground state energy density problem, than for the ground state energy. Gottesman and Irani [GI09] proved that the ground state energy density problem for translationally invariant, nearest-neighbour, quantum Hamiltonians on a 1D chain with a $\Omega(1/2^n)$ promise-gap is NEXP-complete. Here, the input is a description of the local interaction of the system, and the complexity is a function of varying over the Hamiltonian. Meanwhile, as a stepping stone to their undecidability result for the spectral gap, [CPGW15b; CPGW15a] proved that deciding whether the ground state energy density is 0 or strictly positive, with no promise gap, is undecidable, Their result holds for quantum, translationally invariant, nearest neighbour Hamiltonians on a 2D square lattice with a fixed local dimension. [Bau+18b] later extended this undecidability result to 1D chains (again as a stepping stone to the spectral gap problem) and [BCW21] extends to to 2D systems for which the local interaction are analytic in the input parameter (as we saw in chapter 3).

However, as with most ground state energy complexity results, these results still have as input the description of the Hamiltonian, and the hardness is a result of varying the Hamiltonian. That is, to prove hardness we need to vary the physical systems itself. Indeed, it is perhaps better to think of these results as being about a parameterised family of Hamiltonians rather than a fixed Hamiltonian.

4.1.1 The Ground State Energy Density problem

If we restrict to a single, fixed Hamiltonian in the thermodynamic limit, it may seem that there are no input parameters left, and complexity theory can have nothing to say! However, one can still ask about the complexity of estimating the ground state energy density to a given precision, where the only input is the precision required. (See section 4.2 for precise problem definitions.) Arguably, this is the problem formulation closest to that often encountered in condensed matter physics.

The ground state energy density of the specific Hamiltonian we construct is a

single, real number $\mathcal{E}_\rho$. Our hardness results imply the solutions to *all* instances of NEEXP-complete problem are encoded in the digits of this single number, with successive digits of $\mathcal{E}_\rho$ giving the solution to successive instances of a canonical NEEXP-complete problem. In this sense, the ground state energy density of this Hamiltonian is somewhat reminiscent of Chaitin's constant [Cha75], but encoding solutions to problems in a certain complexity class, rather than the Halting problem.

Aside: at the time that a preprint of this work was released, a similar work by Aharonov and Irani was released with similar results on the estimation of the ground state energy density [AI21] (with the published work now appearing here [AI22]). Both results were achieved independently and publication of the preprints was organised to appear simultaneously.

4.2 Results

Define the energy density of the finite lattice as:

Definition 4.1 (Ground State Energy Density)**.** *Consider a translationally invariant Hamiltonian defined on an $L \times H$ lattice, $H^{\Lambda(L\times H)}$). The* ground state energy density *is defined as*

$$\mathcal{E}_\rho(L,H) := \frac{\lambda_0(H^{\Lambda(L\times H)})}{LH}. \tag{4.1}$$

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

$$\mathcal{E}_\rho := \lim_{L,H\to\infty} \mathcal{E}_\rho(L,H). \tag{4.2}$$

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

This limit is well defined [CPGW15a]. We now consider some useful definitions for the computational problems. For all these definitions we will be referring to the infinite lattice case. We can cast the problem of finding $\mathcal{E}_\rho$ as a computational promise problem similar in spirit to the local Hamiltonian problem:

Definition 4.2 (Ground State Energy Density (GSED) promise problem)**.**
Problem Parameters: *A fixed, translationally invariant, nearest-neighbour Hamiltonian on a* $2D$ *infinite square lattice of* d*-dimensional spins.*
Input: *Two real numbers* β *and* α*, such that* $\beta - \alpha = \Omega(2^{-p(n)})$*, for some integer* n *and polynomial* $p(n)$*.*
Output: *Determine whether* $\mathcal{E}_\rho > \beta$ *(*No *instance) or* $\mathcal{E}_\rho < \alpha$ *(YES instance).*
Promise: *The ground state energy density does not lie between in the interval* $[\alpha, \beta]$*.*

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

Definition 4.3 (Ground State Energy Density (FGSED) function problem)**.**
Problem Parameters: *A fixed, translationally invariant, nearest-neighbour Hamiltonian acting on an* $2D$ *infinite lattice of* d*-level spins.*
Input: *An error bound* ϵ*, specified in binary.*
Output: *An approximation to the ground state energy density,* $\tilde{\mathcal{E}}_\rho$ *such that* $|\mathcal{E}_\rho - \tilde{\mathcal{E}}_\rho| \leq \epsilon$*.*

We will often restrict GSED in definition 4.2 to classical Hamiltonians, rather than general (quantum) Hamiltonians. When we wish to highlight this distinction, we refer to these as *classical* GSED and *quantum* GSED, respectively.

The main results of this work are as follows:

Theorem 4.1. $\mathsf{P}^{\mathsf{NEEXP}} \subseteq \mathsf{EXP}^{\mathsf{GSED}} \subseteq \mathsf{EXP}^{\mathsf{NEXP}}$ *for classical* GSED.

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

Corollary 4.1. GSED *is* NEEXP*-hard under exponential time Turing reductions, for a classical, translationally invariant, nearest-neighbour Hamiltonian.*

We also prove:

Theorem 4.2. *Classical* GSED $\in$ NEXP.

Corollary 4.1 and theorem 4.2 are not in conflict with each other. Allowing exponential-time Turing reductions (as opposed to the polytime Turing reductions usually considered) allows exponentially harder problems to be solved.

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

We can also consider the case of quantum Hamiltonians:

Theorem 4.3. $\mathsf{P}^{\mathsf{NEEXP}} \subseteq \mathsf{EXP}^{\mathsf{GSED}} \subseteq \mathsf{EXP}^{\mathsf{QMA}_{\mathsf{EXP}}}$ *for quantum* GSED.

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

Theorem 4.4. FGSED $\in \mathsf{FP}^{\mathsf{NEXP}}$ *for classical* FGSED.

We also get the bound

Lemma 4.1. $\mathsf{FP}^{\mathsf{NEEXP}} \subseteq \mathsf{FEXP}^{\mathsf{FGSED}} \subseteq \mathsf{FEXP}^{\mathsf{NEXP}}$*, for* FGSED *for a fixed classical, translationally invariant, nearest neighbour Hamiltonian.*

4.3 Preliminaries

Definition 4.4. NEEXP or N2EXP

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

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

We also define $\mathsf{QMA}_{\mathsf{EEXP}}$ the same way as QMA, but allowing for a doubly-exponentially long witness and circuit runtime.

We also refer the reader to section section 2.2.3 for definitions of oracle machines and their associated complexity classes. For the particular case of PSPACE^O machines, the PSPACE machine can execute exponentially many computational steps. Thus there is a subtlety as to whether the space bound also applies to the oracle tape or not. Multiple possible definitions for what the PSPACE machine has access to with regards to the oracle tape have been considered in the literature [For94]. We discuss the different results we get depending on the choice of definition in section 4.6.2.1.

4.4 Tiling Preliminaries

4.4.1 Robinson Tiles

Robinson's tiling [Rob71] is based on a set of 5 basic Wang tiles, shown in figure fig. 4.1, with the rule that one tile can be placed next to another only if the arrow heads on the first tile correctly join with the arrow tails on the adjacent tile. I.e. the tiling rules enforce the condition that *arrow heads on one tile must meet arrow tails of the same type on its neighbour in the appropriate direction*.

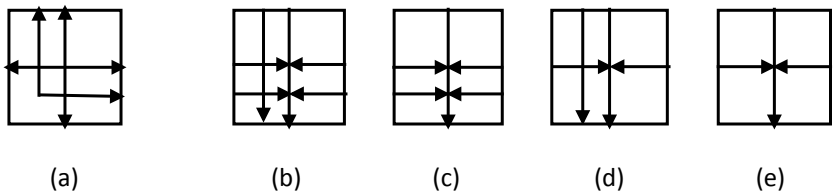

Figure 4.1: The five Robinson tiles we will use. Image taken from [CPGW15a].

Tile (a) in fig. 4.1 has arrows on all sides of the tile and is known as a *cross* and in this depiction is said to face up and to the right. The other 4 tiles are known as *arms*. Each of the arms has a principle arrow across the centre of the tile and which indicates its direction (all the tiles depicted in fig. 4.1 are facing downwards). Arrow markings can be either red or green. On a given arm the horizontal and vertical arrows must have different colours and on cross tiles we force all arrow markings to have the same colour. The Robinson tile set includes all rotations and reflections of these basic tiles.

When these tiles are augmented with certain additional markings, described in [Rob71; CPGW15a], the tiling rules force a pattern of interlocking, nested squares to form in any valid tiling of the plane (see fig. 4.2(c)). The series of squares have side lengths $3, 7, 15, 31, \ldots, 2^n - 1$, for $n \in \mathbb{N}$ (see fig. 4.3). Robinson adds additional coloured markings to the tiles, such that for odd n the borders formed by the double-arrow tile markings running along the edges of the squares are green, and for even n they are red. We direct the reader to [Rob71] and [CPGW15a] for more detailed discussions of tiling pattern and how it is formed.

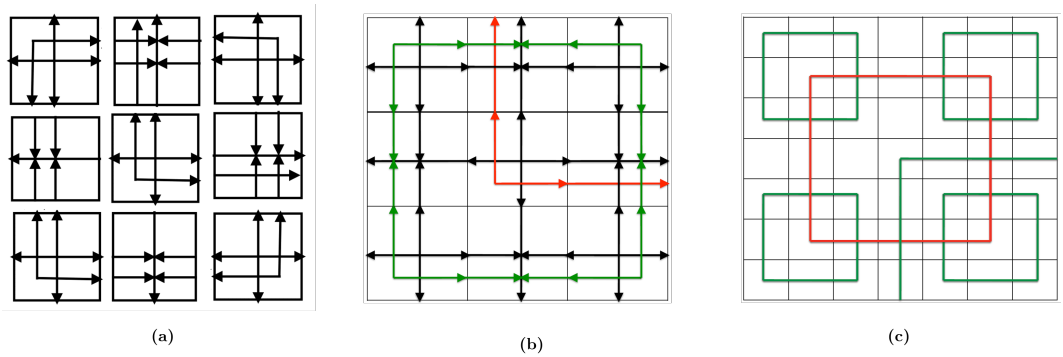

Figure 4.2: (a) A possible tiling arrangement to create a 3-square. (b) shows the same square once the coloured arrows have been introduced. (c) shows a 7-square having combined several 3-squares. Images (b) and (c) taken from [CPGW15a].

For our purposes we will mostly focus on red borders, and refer to these as just *borders*. The combination of the borders and the interior of the border is referred to as a *square*. We refer to a red border of side $4^n - 1$ as an *n-border*. When green borders are referenced, this will always be made explicit.

Consider fig. 4.4. Let $R_v^r, R_v^l, R_h^u, R_h^d$ be the sets of Robinson tiles which contain tiles of type (b), (c), and (d) markings, where the double-arrow markings going across the entire tile are red, and where the arrow markings going across the entire tile are respectively facing right, left, up or down. Let R_X be the set of red crosses, and let $R_X^{UR}, R_X^{UL}, R_X^{DL}, R_X^{DR}$ be the cross tiles that have double arrow markings facing up-right, up-left, down-left and down-right respectively.

Definition 4.5 (n-borders)**.** *Consider a* $(4^n - 1) \times (4^n - 1)$ *subset of a tiling grid, not including its interior. Then the region forms* 2*-border if for every point along the left vertical edge, right vertical edge, bottom horizontal edge, and along the*

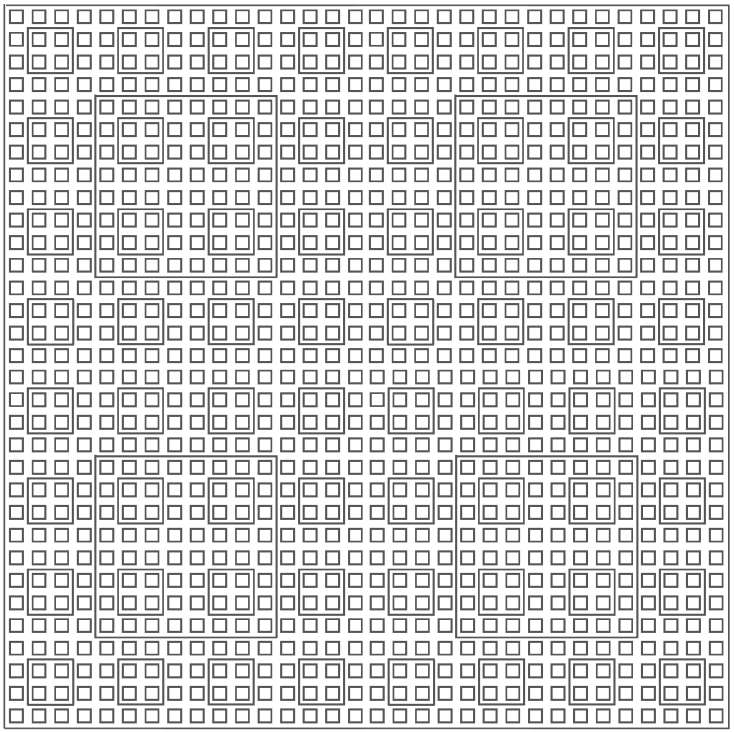

Figure 4.3: A Robinson tiling pattern showing only red borders. Image modified from [CPGW15a].

top horizontal edge is in $R_h^l, R_h^r, R_h^d, R_h^u$*, respectively. Furthermore, the tile in the top-right corner* R_X^{DL}*, top-left corner is* R_X^{DR}*, bottom-left corner is* R_X^{UR}*, and bottom right corner is* R_X^{UL}*.*

Finally note that Robinson tiles allows for two half-planes to be translated relative to each other without violating any of the tiling rules. We wish to avoid this and hence use the modified set of Robinson tilings introduced in [CPGW15a], such that the final set of tiles is all rotations and reflections of those shown in fig. 4.4. It is shown in [CPGW15a] that these tiles produce the same pattern of nested squares, but prevent any two half-planes from be translated relative to each other.

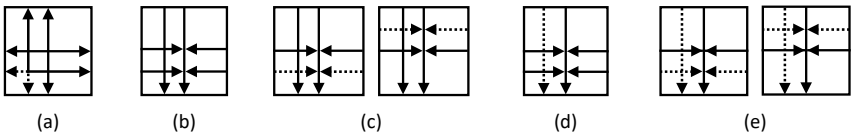

Figure 4.4: The standard Robinson tiles with additional dashed markings added in to prevent slippage between planes. Image modified from [CPGW15a].

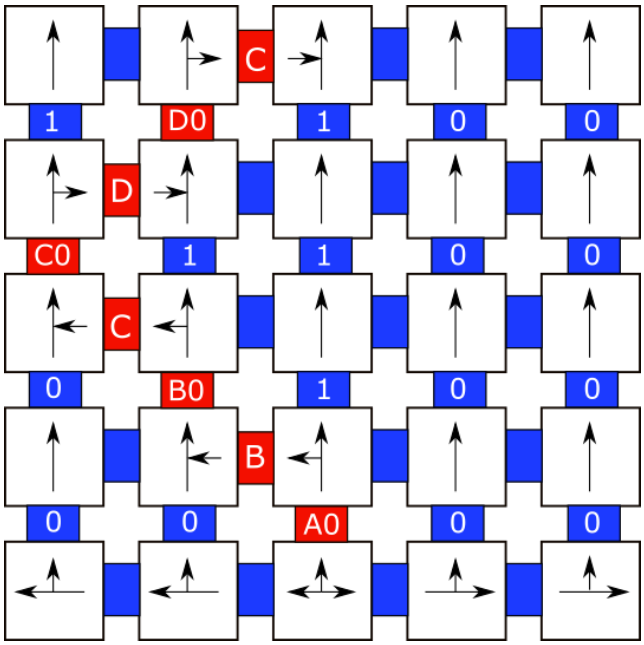

Figure 4.5: The evolution of a classical TM can be represented by Wang tiles, where colours of adjacent tiles have to match, and arrow heads have to meet arrow tails. Here the evolution runs from the bottom of the square to the top. The red labels between adjacent rows represent the position and state of the TM head, and the red labels between adjacent columns represent movement of the TM head after it has acted on the cell.

4.4.2 Encoding Turing Machines with Tiles

It is well known that the evolution of a classical Turing Machine can be encoded as a set of Wang tiles as explained in section 3.7.4 (see fig. 4.5 for an example of such an encoding, and see [Ber66; Rob71; GI09; CPGW15a; Bau+18a] for some further detailed discussions on this topic).

4.4.3 Encoding Turing Machines in the Robinson Tiling

In this section we review how the tiling-encoding of TMs can be combined with the Robinson tiling to create a new set of tiles which, when the plane is tiled according to the tiling rules, encodes the evolution of a separate TM within each n-square in the Robinson tiling pattern. This construction was introduced in [Rob71] to prove undecidability of the tiling of a 2D plane.

Encoding the evolution of a TM directly within the interior of a n-border is not possible as the Robinson tiling pattern is composed of m-squares nested within other n-squares, $m < n$. Thus TMs would overlaps with each other. [Rob71] circumvents

this problem by identifying a sub-grid within each Robinson n-border which allows a TM to be encoded without overlapping with the smaller m-squares, $m < n$, nested within.

Definition 4.6 (Free Rows/Columns and Free Squares, [Rob71])**.** *A* free row/column *of square is a row/column in a Robinson n-border that stretches across the border's interior uninterrupted by any of the m-borders with $m < n$.*

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

Lemma 4.2 (Encoding TM in Robinson Tiling, [Rob71])**.** *Consider any classical Turing Machine which can have its evolution be encoded in a $(2^n + 1) \times (2^n + 1)$ grid of Wang tiles. Then the evolution of this TM can be encoded in the free rows and columns of an n-square in a Robinson Tiling.*

We will use the details of Robinson's construction of lemma 4.2 later, hence we provide some exposition here.

Consider a Robinson n-border. Following [Rob71], to demarcate where the free tiles are so that we can encode a Turing Machine in them, introduce a new kind of marking on the tiles called an 'obstruction signal'. These signals are designed so they are emitted and absorbed from the outside of a red border and while also being absorbed by the inside of a border, as seen in fig. 4.6. In terms of tiles, these markings are formed by adding an additional set of markings such that tiles of type (b) in fig. 4.4 with a red double-arrow "emit" the obstruction signals from one side and "absorb" them on both sides. Tiles that do not emit or absorb obstruction signals force them to propagate in the same direction. The obstruction signals are only emitted from the outer edges of a red Robinson border. A *free tile* is one which does not have an obstruction signal going across it in either direction. In our new tile set, we only encode the Turing Machine tape, head and state symbols in the free tiles.

Transmitting Signals between Free Tiles Thus we are able to encode the evolution of a Turing Machine in these free tiles, effectively creating a $(2^n + 1) \times (2^n + 1)$ square for it to run in. There is a problem in that the free tiles are not spatially close to

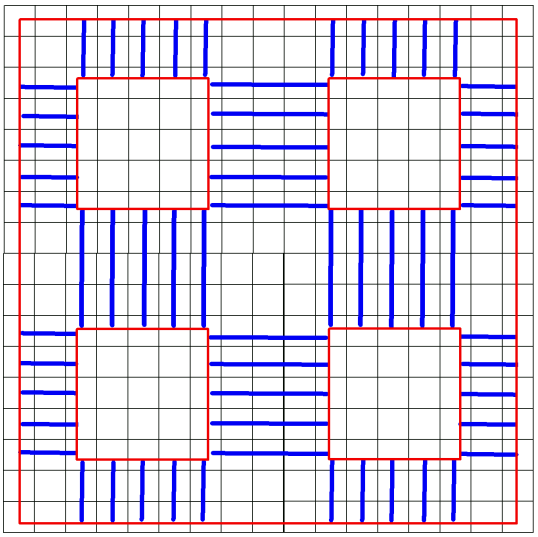

Figure 4.6: The obstruction signals for a red 2^4-square are shown in blue. Each of the tiles within the 2^2-squares emits a signal outwards. The free rows are the rows in which there are no obstruction signals running horizontally (for example the central row). The free columns are the columns in which there are no obstruction signals running vertically (for example the central column).

each other. To solve this, [Rob71] implicitly introduces a new set of tile markings: Turing Machine signals. These signals can be emitted and absorbed by free tiles and run along free rows and columns. Otherwise they are absorbed by tiles with double arrowed red markings: tiles of types (a), (b), or (c), shown in fig. 4.4, on the sides of tiles parallel to the red double arrow lines. Tiles which are not free tiles, and do not absorb the TM signals, force the TM signals to propagate across them. These signal markings allow the tiling to transmit the necessary conditions between spatially distant free tiles.

Initialising the TM Finally, boundary conditions are needed to force the correct initial configuration of the Turing Machine. To ensure this, [Rob71] introduces a further set of tile markings that interact with the Turing Machine markings. The markings are chosen so that every arm tile which is both horizontally facing and forms the bottom border of an n-border, and which *does not* absorb an obstruction signal, must emit a Turing Machine signal upwards. Choose this signal to be s_0 which will force the tiles in the initial layer at free positions to be blank, so that the initial tape configuration is entirely blank.

The exception to this is in the centre of the edge where the tile will emit a Turing

Machine signal $s_0 q_0$ indicating the Turing Machine head starts there. Similarly, choose the tiling markings so any arms in the top, left and right parts of the square's border will absorb any stray Turning Machine signals along their inner edges.

4.5 Robinson Tiling Robustness

In this section we prove a series of results demonstrating that if a region $R \subseteq \mathbb{Z}^2$ is tiled with Robinson tiles, but tiling defects are allowed to occur (i.e. points between adjacent tiles at which the matching rules are not satisfied), then only a finite number of Robinson squares can be destroyed per defect. Similarly bounds were proven in [Mie97; CPGW15a], but are not strong enough for our purposes.

4.5.1 Robinson Border Deficit Bound

Definition 4.7 (Tile configuration)**.** *A* tile configuration *is an assignment of a Robinson tile to each point in the lattice $\mathbb{Z}_2$. The defect set of a tile configuration T is the set of all points in Z_2^* between non-matching tiles in T.*

Definition 4.8 (Border deficit)**.** *The* n-border deficit*,* $\text{deficit}_n(T, S, R)$*, in a region S of a tile configuration T with respect to Robinson tiling R, is the (magnitude of the) difference between the total number of complete n-borders of T within S, and the number of complete n-borders of R within S.*

The total border deficit*,* $\text{deficit}(T)$*, of a tile configuration T is the difference between the total number of complete borders in T and the number of complete borders in a Robinson tiling of the same region, maximised over Robinson tilings.*

The bulk of the work proving a bound on the border deficit was done by Toby Cubitt, and hence we exclude the proofs from this thesis, and instead state the relevant theorem:

Theorem 4.5. *Let T be a tile configuration of a finite subregion of $\mathbb{Z}_2$ with perimeter of length L. Let D denote its defect set. The border deficit of T is bounded by*

$$\text{deficit}(T) \leq 399|D| + L. \tag{4.3}$$

Proof. See the arXiv preprint [WC21]. □

We now use theorem 4.5 to prove a bound on the deficit on the number of Robinson squares.

Definition 4.9. *An* inner border *of an n-border in a Robinson tiling is an m-border (necessarily with $m \leq n$) located in the interior of the n-border, and not contained in the interior of any other border.*

Definition 4.10. *An n-*square *is a tile configuration of a $(4^n - 1) \times (4^n - 1)$ region of the lattice containing an n-border around the perimeter, inner m-borders in the same locations as in a Robinson tiling, and no other borders and no defects in the region between the n-border and the inner m-borders.*

*We call the region of the lattice between an n-border and the locations where the inner m-borders would be in a Robinson tiling, including the n-border and the inner m-borders themselves, the n-*square region*. The* interior *of the n-square region is the region excluding the n-border and the inner m-border locations.*

Definition 4.11 (Square deficit)**.** *The n-*square deficit*,* $\text{square\_deficit}_n(T, S, R)$*, in a region S of a tile configuration T with respect to Robinson tiling R, is the (magnitude of the) difference between the total number of complete n-squares of T within S, and the number of complete n-squares of R within S.*

The total square deficit*,* $\text{square\_deficit}(T)$*, of a tile configuration T is the difference between the total number of complete squares in T and the number of complete squares in a Robinson tiling of the same region, maximised over Robinson tilings.*

Theorem 4.6. *Let T be a tile configuration of a finite subregion of $\mathbb{Z}_2$ with perimeter of length L. Let D denote its defect set. The square deficit of T is bounded by*

$$\text{square\_deficit}(T) \leq 799|D| + 2L. \tag{4.4}$$

Proof. See the arXiv preprint [WC21]. □

4.5.2 Obstruction Signal Bound

We now bound the deficit of squares which have correct obstruction tilings. As discussed in section 4.4.3, obstruction signals are used to demarcate free tiles. We now add these markings to the modified Robinson tiles: we make a small change relative to the obstruction markings in [Rob71] and choose all obstruction signals to have a direction: the horizontal left-to-right or downwards. This new set of tiles modified Robinson tiles + obstruction markings are labelled *obstruction tiles*.

Definition 4.12 (Correct Obstruction Tiling)**.** *A complete Robinson square has a* correct obstruction tiling *if:*

1. *a tile has no obstruction signals iff it is in both a free row and free column.*
2. *a tile has horizontal (vertical) obstruction signals run across it iff it is in a free column (row).*
3. *all tiles not contained a free row or column of the n-border, and not contained within another m-border, have obstruction signals running across them both horizontally and vertically.*

Consider the deficit in the number of squares present with correct obstruction tilings in the case of defects versus defect-free tilings:

Lemma 4.3. *Let T be a tile configuration of a finite subregion of $\mathbb{Z}_2$ with perimeter of length L, and D its defect set.*

Define the obstruction deficit *of T,* obstruction_deficit(T)*, to be the difference between the total number of complete Robinson squares in T with a correct internal obstruction tiling, and the number of these in a Robinson tiling of the same region, maximised over Robinson tilings.*

The obstruction deficit of T is bounded by

$$\text{obstruction\_deficit}(T) \leq 800|D| + 2L. \tag{4.5}$$

Proof. Definition 4.11 and theorem 4.6, bound the square deficit. Given any complete square, obstruction signals which are emitted by the smaller interior borders can

terminate if (a) there is a defect in their path, or (b) they end on a tile with double red arrows. Case (a) implies there is a defect in the interior of the square. Case (b) implies there a tile with a double red arrow marking which is horizontal (vertical) in a free column (row), which must immediately result in a defect if contained in a square. Conversely, a tile without obstruction signals which is not placed in a free row must be adjacent to a tile with obstruction signals, and thus must cause a defect. It follows that if any of the conditions from definition 4.12 are not met, then there must be an interior defect. Furthermore, if there is a complete n-square with a defect, the tiling outside the n-border is unaffected as its n-border is identical to the case without a defect. Thus

$$\text{obstruction\_deficit}(T) \leq \text{square\_deficit}(T) + |D| \tag{4.6}$$

□

4.5.3 Robinson + TM Tiling Bound

Finally, the obstruction tiles need to be combined with the obstructions tiles, which are themselves a combination of the Robinson and obstruction tiles. We use the full set of tiles described in section 4.4.3 which include the Robinson markings, obstruction markings, and Turing Machine signals.

Definition 4.13. *An n-square has a correct TM encoding if its* $(2^n+1)\times(2^n+1)$ *free tiles encode the correct evolution of a TM from some fixed initial state according to the TM's transition rules.*

Lemma 4.4. *Let* T *be a tile configuration of a finite subregion of* $\mathbb{Z}_2$ *with perimeter of length* L*,* D *its defect set.*

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

The total deficit of T is bounded by

$$\text{total\_deficit}(T) \leq 801|D| + 2L. \tag{4.7}$$

Proof. All tiles with no obstruction markings present must have TM markings (but not TM signal markings) and visa-versa. Thus, assuming an n-square has correct obstruction markings, TM markings only appear on free tiles, and the TM signals only appear on tiles with obstruction markings going horizontally **or** vertically, but not both.

All tiles with only a horizontal (vertical) obstruction markings have a TM signal running vertically (horizontally). Since by definition 4.12, such tiles only appear in the appropriate free column (row), the TM signals only run along the free columns (rows). The TM signals propagate until they reach a free square, at which point they may change. If a TM signal changes between tiles, not mediated by a free tile, there must be a defect. Thus

$$\text{total\_deficit}(T) \leq \text{obstruction\_deficit}(T) + |D|. \tag{4.8}$$

□

4.6 Proofs of GSED Complexity Results

4.6.1 Classical hardness for $\mathsf{P}^{\mathsf{NEEXP}}$

In this section we set out to prove the following theorem:

Theorem 4.7. $\mathsf{P}^{\mathsf{NEEXP}} \subseteq \mathsf{EXP}^{\mathsf{GSED}}$*, for* GSED *as defined in definition 4.2, for a classical, nearest-neighbour, translationally invariant Hamiltonian.*

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

Specifying the Encoded TMs We want to enumerate over all input strings for a TM deciding some language, encode these using tiles, and arrange for the TMs running

on different inputs to be encoded within Robinson borders of different sizes. This is summed up as:

Lemma 4.5 (TMs in Robinson Squares)**.** *Let $x_n \in \{0,1\}^*$ be the $(n-n_0)^{th}$ string in lexicographic order where n_0 is a fixed integer, and let M be a non-deterministic TM. It is possible to construct a tile set such that all valid tilings of an $L \times L$ lattice consist of the pattern of nested squares formed by the Robinson tiling, such that within each complete n-border, $\forall n \geq n_0$, the tiles encode a valid computational evolution of $M(x_n)$ for time $2^{2^{c|x_n|}}$, $c \geq 1$.*

Proof. As per lemma 4.2, we are able to encode a TM in the $(2^n+1) \times (2^n+1)$ grid of free tiles of a Robinson n-squares. Section 3 of [GI09] proves that given a $L \times L$ grid with an appropriate border, it is possible to encode a computation of length kL and space L, for $k = O(1)$. Here the Robinson n-borders provide such a border.

We choose to encode a series of TMs as follows. This first TM is binary counter machine M_{BC} which after time step T, has T written in binary on the tape (see [Pat14] or [GI09, Section 3] for an explicit construction of this machine). This outputs the square size 2^n+1 in binary. Then run a TM computing $\log_4(y-1)-n_0$ on this output, which outputs $x \in \{0,1\}^*$, the $(n-n_0)^{th}$ string in lexicographic order. Finally encode a non-deterministic TM which takes input x and runs for $2^{2^{(k-2)|x|}} (\leq 2^{(k-2)n})$ steps. We can force M to run for $2^{2^{(k-2)|x|}}$ steps by employing a counter to limit the number of steps to $2^{2^{(k-2)|x|}}$; if the TM halts before reaching end of the allotted time, the final time step is copied to the next time step. If the timer runs out before the full grid space is used, the final time step of the encoded TM is copied forwards until the grid is filled. Choose n_0 to be the smallest integer such that these TMs have enough space to operate properly on a grid of size $(2^{n_0}+1) \times (2^{n_0}+1)$. □

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

4.6.1.1 Mapping Tiles to Hamiltonians

So far we have presented the problem in terms of a tiling problem and need to map this to a classical Local Hamiltonian problem. We use the tiling to local Hamiltonian mapping presented in section 3.7.1 to create such a Hamiltonian with translationally invariant, nearest neighbour local interactions on a 2D lattice. In particular. we map the tiling rules produced by lemma 4.5 to a Hamiltonian to get a nearest-neighbour, translationally invariant Hamiltonian. We add a term penalising rejecting instances of the verification computation; Π_{NO} is an additional term we add in which assigns an energy penalty to No problem instances.

We encapsulate the definition of the Hamiltonian in the following:

Definition 4.14 (Robinson + Computation Hamiltonian)**.**
Let $h^{col,Rob}, h^{row,Rob} \in \mathcal{B}(\mathbb{C}^R \otimes \mathbb{C}^R)$ be the local terms which encode the local matching rules for the Robinson tiling, obstruction rules and TM rules from lemma 4.5. Let $(\Pi_{NO})_{j,j+1}$ be a projector onto the reject state of the encoded TM, M, on a site in row j, and a Robinson border tile on the adjacent site in row $j+1$. Then the overall local terms are:

$$h^{row}_{i,i+1} = \Lambda h^{row,Rob}_{i,i+1} \tag{4.9}$$

$$h^{col}_{j,j+1} = \Lambda h^{col,Rob}_{j,j+1} + (\Pi_{NO})_{j,j+1} \tag{4.10}$$

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

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

Lemma 4.6. *Define $H(4^n)|_P$ to be the Hamiltonian on a $(4^n-1)\times(4^n-1)$ region described by the local terms given in eqs. (4.9) and (4.10), restricted to the subspace P corresponding to defect-free tilings of the region that contain a complete Robinson*

n-border. Let $x \in \{0,1\}^*$ *be the* $(n-n_0)^{th}$ *string in lexicographic order and let* M *be a non-deterministic Turing Machine running for time* $2^{2^{cm}}$ *on inputs of length* m, $c \geq 1$.

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

$$\lambda_0(H(4^n)|_P) = i_n := \begin{cases} 0 & M(x) \textit{outputs } \textsc{Yes} \\ 1 & M(x) \textit{outputs } \textsc{No}. \end{cases} \tag{4.11}$$

Proof. $H(4^n)|_P$ is restricted to the subspace of valid tiling configurations containing a complete Robinson n-border. Clearly, this border must run around the edge of the $(4^n-1)\times(4^n-1)$ region. By lemma 4.5 valid tilings encode the evolution of a non-deterministic TM $M(x)$, where x is the $(n-n_0)^{th}$ string in lexicographic order. By restricting to the subspace P we have ensured the encoded TM evolves correctly.

If x is a Yes instance, then $M(x)$ must have an accepting computational path, and so there must be a set of states that encode the correct evolution which finishes in an accepting state. Hence there is no energy penalty and the ground state is 0.

If x is a No instance, then there is no accepting path. Any correct evolution of $M(x)$ therefore enters the rejecting state, and the tile marking the rejecting state of the TM picks up an energy penalty of 1 from the term $(\Pi_{NO})_{k,k+1}$ (and no other state receives this energy penalty).

□

4.6.1.2 Robustness of the Ground State

We now want to find the ground state energy of the lattice with Hamiltonian from definition 4.14. The possible energy contributions come from tiling defects and energy penalties for No instances of the encoded computation. In the following, we use the square deficit bounds established in section 4.5 to show that it is energetically unfavourable to have too many tiling defects, regardless of how many No instances might be encoded in n-squares.

Lemma 4.7 (Robinson Square Bound)**.** *The number of n-borders in a Robinson tiling of* $\Lambda(L\times H) \subset \mathbb{Z}^2$ *using modified Robinson tiles is bounded by* $\geq (\lfloor H/2^{n+1} \rfloor -$

$1)\big(\lfloor L/2^{n+1}\rfloor - 1\big)$ *and* $\leq \big(\lfloor H/2^{n+1}\rfloor + 1\big)(\lfloor L/2^{n+1}\rfloor + 1)$ *for all* n.

Proof. A Robinson border is completely contained in an $L \times H$ lattice iff its top edge and its left edge are completely contained in the lattice. Lemma 48 of [CPGW15a] shows that the number of top edges of a Robinson n-square which are completely contained in the $L \times H$ lattice is $\geq \lfloor H/2^{n+1}\rfloor\big(\lfloor L/2^{n+1}\rfloor - 1\big)$ and $\leq \big(\lfloor H/2^{n+1}\rfloor + 1\big)\lfloor L/2^{n+1}\rfloor$. From this it is straightforward to see the number of left edges which are completely contained in the lattice is $\geq \big(\lfloor H/2^{n+1}\rfloor - 1\big)\lfloor L/2^{n+1}\rfloor$ and $\leq \lfloor H/2^{n+1}\rfloor\big(\lfloor L/2^{n+1}\rfloor + 1\big)$.

Combining these two bounds gives $\geq (\lfloor H/2^{n+1}\rfloor - 1)\big(\lfloor L/2^{n+1}\rfloor - 1\big)$ and $\leq \big(\lfloor H/2^{n+1}\rfloor + 1\big)(\lfloor L/2^{n+1}\rfloor + 1)$. □

We now want to check that the ground state of the Hamiltonian on the overall lattice is a tiling of the lattice with Robinson squares in which a verification TM is encoded as we expect, but potentially with a bounded number of defects.

Lemma 4.8. *Let* $h^{row}, h^{col} \in \mathcal{B}(\mathbb{C}^R \otimes \mathbb{C}^R)$ *be the local interactions that encode the tiling rules given by eqs.* (4.9) *and* (4.10)*. Let* $H^{\Lambda(L\times L)}$ *be the Hamiltonian with these local interactions on* $\Lambda(L \times L)$.

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

$$\left[\sum_{n=n_0}^{\lfloor \log_4(L/2)\rfloor} \left(\left\lfloor \frac{L}{2^{2n+1}}\right\rfloor - 1\right)^2 \lambda_0(H(4^n)|_P) + \Lambda|D| - k_1|D| - k_2 L,\right.$$
$$\left.\sum_{n=n_0}^{\lfloor \log_4(L/2)\rfloor} \left(\left\lfloor \frac{L}{2^{2n+1}}\right\rfloor + 1\right)^2 \lambda_0(H(4^n)|_P) + \Lambda|D| - k_1|D| - k_2 L\right] \tag{4.12}$$

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

Proof. From lemma 4.6, we see that in the ground state energy contribution from each sufficiently large, complete, Robinson n-square is $\lambda_0(H(4^n)|_P) \in \{0, 1\}$. By lemma 4.7, the number of n-borders of a given size in an $L \times L$ region with no defects is bounded by $\geq (\lfloor L/2^{2n+1}\rfloor - 1)^2$ and $\leq (\lfloor L/2^{2n+1}\rfloor + 1)^2$.

Let $N(D)$ denote the number of borders correctly encoding the TM evolution for some tile configuration T with defect set D. Let $N_{YES}(D)$, $N_{NO}(D)$ be the

number of borders which encode YES and NO instances, respectively. Hence $N(D) = N_{YES}(D) + N_{NO}(D)$. Let $E(|D| \text{ defects})$ be the energy of a configuration with $|D|$ defects. Then

$$E(|D| \text{ defects}) = \Lambda|D| + N_{NO}(D) \tag{4.13}$$

$$E(0 \text{ defects}) = N_{NO}(\emptyset) \tag{4.14}$$

Combining these:

$$E(|D| \text{ defects}) - E(0 \text{ defects}) = \Lambda|D| - (N_{NO}(\emptyset) - N_{NO}(D)) \tag{4.15}$$

$$E(|D| \text{ defects}) - E(0 \text{ defects}) \geq \Lambda|D| - (N(\emptyset) - N(D)) \tag{4.16}$$

where the fact $N(\emptyset) - N(D) \geq N_{NO}(\emptyset) - N_{NO}(D)$ has been used.

Lemma 4.4 gives $\text{total\_deficit}(T) = N(\emptyset) - N(D) \leq k_1|D| + k_2 L$ for constants k_1, k_2, hence

$$E(|D| \text{ defects}) - E(0 \text{ defects}) \geq \Lambda|D| - (k_1|D| + k_2 L) \tag{4.17}$$

Now choose the parameter Λ to be constant such that $\Lambda \gg k_1 + k_2$. If $|D| = \Omega(L)$, then for sufficiently large L,

$$E(|D|\text{defects}) - E(0\text{defects}) \geq (\Lambda - k_1 - k_2)\Omega(L)k = \Omega(L).$$

Thus, for sufficiently large L, the 0-defect case becomes the ground state.

If $|D| = O(L^{1-o(1)})$, then for sufficiently large L we have that

$$E(|D|\text{defects}) - E(0\text{defects}) \geq \Lambda|D| - k_1|D| - k_2 L = -O(L).$$

Thus we see the minimum lower bound occurs for $|D| = O(L^{1-o(1)})$

There is one energy contribution that has been omitted. Some Robinson squares will be too small to have the TM's encoded in them run correctly. However, there are only finitely many square sizes for which this is the case, and each square size

appears with constant density. So their contribution to to the ground state energy density is a constant which can be computed in constant time, and subtracted off with a 1-local term of the form $\sum_{i\in\Lambda(L\times L)} \alpha \mathbb{1}_i$. (Cf. [CPGW15a].) □

For simplicity of the exposition, we omit the above constant energy shift from the expressions and discussion, as it does not affect the analysis.

Lemma 4.9. *Consider an $L\times L$ lattice with a local Hamiltonian interactions given by eqs. (4.9) and (4.10), and let $H(4^n)|_P$ and i_n be defined as in lemma 4.6. In the limit of $L\to\infty$, the ground state energy density is*

$$\mathcal{E}_\rho = \frac{1}{4}\sum_{n=n_0}^{\infty} \frac{\lambda_0(H(4^n)|_P)}{16^n} = \frac{1}{4}\sum_{n=n_0}^{\infty} \frac{i_n}{16^n}. \tag{4.18}$$

Proof. By lemma 4.8, we have bounds on the ground state energy for the region:

$$\begin{aligned}
&\sum_{n=n_0}^{\lfloor \log_4(L/2)\rfloor} \frac{1}{L^2}\left(\left\lfloor \frac{L}{2^{2n+1}}\right\rfloor - 1\right)^2 \lambda_0(H(4^n)|_P) + (\Lambda - k_1)O(L^{-1}) + k_2 L^{-1} \\
&\leq \mathcal{E}_\rho(H^{\Lambda(L\times L)}) \\
&\leq \sum_{n=n_0}^{\lfloor \log_4(L/2)\rfloor} \frac{1}{L^2}\left(\left\lfloor \frac{L}{2^{2n+1}}\right\rfloor + 1\right)^2 \lambda_0(H(4^n)|_P)) + (\Lambda - k_1)O(L^{-1}) + k_2 L^{-1}
\end{aligned} \tag{4.19}$$

Taking the limit $L\to\infty$ gives

$$\lim_{L\to\infty} \mathcal{E}_\rho(H^{\Lambda(L\times L)}) = \mathcal{E}_\rho = \frac{1}{4}\sum_{n=n_0}^{\infty} \frac{\lambda_0(H(4^n)|_P)}{16^n} = \frac{1}{4}\sum_{n=n_0}^{\infty} \frac{i_n}{16^n}. \tag{4.20}$$

□

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

Theorem 4.8 ($\mathsf{P}^{\mathsf{NEEXP}}\subseteq\mathsf{EXP}^{\mathsf{GSED}}$)**.** $\mathsf{P}^{\mathsf{NEEXP}}\subseteq\mathsf{EXP}^{\mathsf{GSED}}$*, for* GSED *as defined in definition 4.2, for a classical, translationally invariant, nearest-neighbour Hamiltonian.*

Proof. Consider any polytime bounded TM M_1. We will show we can simulate M_1^{NEEXP} with M_2^{GSED} where M_2 is another exptime TM. If M_1^{NEEXP} takes an n-bit input, it can then make $O(\mathrm{poly}(n))$ queries. Denote these queries by $\{q_i\}_{i=1}^{O(\mathrm{poly}(n))}$.

Each individual query must have length $|q_i| = O(\text{poly}(n))$. The M_1 machine then runs for an $O(\text{poly}(n))$ time and produces some output.

To simulate this, M_2 takes the n-bit input and calculates each of the queries which M_1 makes: $\{q_i\}_{i=1}^{O(\text{poly}(n))}$. Each query q_i is made to a NEEXP oracle. So M_2 takes each query q_i, and reduces it to an instance of determining the output of a doubly-exponentially time non-deterministic TM, M, on input y_i. This reduction can be computed in polynomial time, as the problem of determining the output of double-exponential-time non-deterministic TMs is manifestly NEEXP-hard. (Note by using padding arguments we can reduce any language in NEEXP to $\text{NTIME}(2^{2^{cn}})$ for some $c > 1$ [Pap94]). This defines a new set of inputs to the non-deterministic machine M, $\{y_i\}_{i=1}^{O(\text{poly}(n))}$, such that $|y_i| = O(\text{poly}(n))$. Now order the $\{y_i\}_i$ lexicographically and take the largest one. Suppose the largest string, y_j, is the k^{th} string in lexicographic order. Then $k = O(2^{O(|y_j|)}) = O(2^{\text{poly}(n)})$.

We will use the GSED oracle for the Hamiltonian of definition 4.14 to perform a binary search in order to obtain a sufficiently precise approximation to the ground state energy density $\mathcal{E}_\rho$, such that we can extract the result of computing M on all inputs up to y_j. To do this, we need to query the GSED oracle on all the instances before it in lexicographic order, of which there are $k = O(2^{\text{poly}(n)})$ many.

By lemma 4.9, outputs i_n to the queries $\{y_i\}_i$ are encoded as

$$\mathcal{E}_\rho = \frac{1}{4}\sum_{n=n_0}^{\infty}\frac{i_n}{16^n}. \tag{4.21}$$

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

$$\mathcal{E}_\rho = \frac{1}{4}\sum_{n=2}^{\infty}\frac{1}{16^n} = \frac{1}{960} \tag{4.22}$$

and otherwise the minimum it can be is $1/64$. Hence we ask whether $\mathcal{E}_\rho \geq \beta_1 = 1/64$

or $\mathcal{E}_\rho \leq \alpha_1 = 1/960$. Thus

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

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

$$\begin{aligned} \beta_m &= \frac{1}{4}\left(\frac{1}{16^m} + \sum_{n=1}^{m-1} \frac{i_n}{16^n}\right) \\ \alpha_m &= \frac{1}{4}\left(\sum_{n=1}^{m-1} \frac{i_n}{16^n} + \sum_{n=m+1}^{\infty} \frac{1}{16^n}\right). \end{aligned} \tag{4.23}$$

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

4.6.2 Classical Containment in $\mathsf{EXP}^{\mathsf{NEXP}}$

We now need to show that for classical GSED, as defined in definition 4.2, $\mathsf{EXP}^{\mathsf{GSED}} \subseteq \mathsf{EXP}^{\mathsf{NEXP}}$. The first step is to show that the ground state energy density of a finite $L \times L$ part of the lattice is a good estimate for the energy density of the full lattice [CPGW15a]:

Lemma 4.10. *Consider a translationally invariant, nearest-neighbour Hamiltonian on $\Lambda(L \times L)$ lattice defined by local terms $h^{row}_{i,i+1}, h^{col}_{j,j+1}$. Let $\mathcal{E}_\rho(L)$ be the energy density of the Hamiltonian on this lattice, and let $\mathcal{E}_\rho$ be the energy density in the*

$L \to \infty$ *limit. Then*

$$|\mathcal{E}_\rho(L) - \mathcal{E}_\rho| = \frac{4\max\left\{\left\|h^{row}_{i,i+1}\right\|, \left\|h^{col}_{i,i+1}\right\|\right\}}{L}. \tag{4.24}$$

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

$$H_{grid}(L,t) = H(tL) - \sum_{i \mod t=0} h^{row}_{i,i+1} - \sum_{j \mod t=0} h^{row}_{j,j+1}. \tag{4.25}$$

The interaction graph of $H_{grid}(L,t)$ is a set of t^2 squares of size $L \times L$. Hence equation 4.25 gives

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

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

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

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

$$|\mathcal{E}_\rho(L) - \mathcal{E}_\rho| \leq \frac{4\max\left\{\left\|h^{row}_{i,i+1}\right\|, \left\|h^{col}_{i,i+1}\right\|\right\}}{L}. \tag{4.26}$$

□

Lemma 4.11. GSED $\in$ NEXP *for any classical, nearest-neighbour, translationally invariant Hamiltonian, for* GSED *as defined in definition 4.2.*

Proof. (α, β) is the input of the problem, $\beta - \alpha = \Omega(2^{-q(n)})$. We show an EXP machine will be able calculate $\mathcal{E}_\rho(L)$ (using the notation of lemma 4.10) using a classical witness for $L = 2^{p(n)}$, for a polynomial p.

First compute the ground state energy of an $L \times L$ square of the lattice. Take as the witness the ground state of the Hamiltonian restricted to an $L \times L$ region of

the lattice: $|\psi\rangle = |\phi_1\rangle \otimes |\phi_2\rangle \otimes \ldots |\phi_{L^2}\rangle$, where $|\phi_i\rangle \in \mathbb{C}^{|S|}$ is the state of the spin at lattice site i. Now,

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

where $\langle i, j \rangle$ denotes pairs of nearest-neighbours. $\langle \phi_i | \left\langle \phi_j \right| h_{i,j} \left| \phi_i \right\rangle \left| \phi_j \right\rangle$ can be computed in $O(1)$ time, and there are $O(L^2)$ such terms. Since $L = 2^{p(n)}$, the estimate $\mathcal{E}_\rho(L)$ can be computed in $O(L^2) = O(2^{2p(n)})$ time. By lemma 4.10, $|\mathcal{E}_\rho(L) - \mathcal{E}_\rho| = O(L^{-1})$, hence provided we choose $p(n)$ to be sufficiently large relative to $q(n)$, the approximation $\mathcal{E}_\rho(L)$ allows us to determine $\mathcal{E}_\rho > \beta$ or $\mathcal{E}_\rho < \alpha$ for $\beta - \alpha = \Omega(2^{-q(n)})$. □

Lemma 4.12 ($\mathsf{EXP}^{\mathsf{NEXP}}$ Containment)**.** $\mathsf{EXP}^{\mathsf{GSED}} \subseteq \mathsf{EXP}^{\mathsf{NEXP}}$*, for* GSED *as defined in definition 4.2, for a fixed, classical Hamiltonian.*

Proof. For $\mathsf{EXP}^{\mathsf{GSED}} \subseteq \mathsf{EXP}^{\mathsf{NEXP}}$ we show that, given an exponential time TM M_1 with access to an oracle GSED, its action can be simulated by an exponential time TM M_2 with oracle access to NEXP.

Consider the action of M_1^{GSED}. If it takes an n-bit input, it may make $O(\exp(n))$ queries, each of length $O(\exp(n))$, before outputting an answer based on these query outcomes. Each query must be in the form of an (α, β) such that $\beta - \alpha = \Omega(2^{-2^{p(n)}})$ for some polynomial p.

The (α, β) queries made by M_1 must have input length of $|q_i| = O(\exp(n))$ By lemma 4.11 determining whether $\mathcal{E}_\rho > \beta$ or $\mathcal{E}_\rho < \alpha$ for $\beta - \alpha = \Omega(2^{-|q_i|}) = \Omega(2^{-2^{-g(n)}})$ is contained in NEXP. Thus M_2^{NEXP} can simulate the queries to GSED by making querying the NEXP oracle, and hence the entire action of $\mathsf{EXP}^{\mathsf{GSED}}$. □

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

Lemma 4.13. *Let* $\mathsf{P}^{\mathsf{GSED}_h}$ *be the class of languages decided by a polynomial time*

oracle machine with access to a GSED *oracle for the Hamiltonian of definition 4.14 only. Let* $\mathsf{P}^{O}_{\log}$ *be the languages decided by a polytime oracle machine with oracle* O *which is only able to make* $\log(n)$ *length queries to the oracle for an n-bit input. Then* $\mathsf{P}^{\mathsf{GSED}_h} \subset \mathsf{P}^{\mathsf{NEEXP}}_{\log}$.

Proof. Let $M_1^{\mathsf{GSED}_h}$ be a polytime TM with oracle access to a GSED oracle for the Hamiltonian defined in definition 4.14 only. Let M_2^{NEEXP} an oracle machine which can only make $O(\log(n))$ length queries to the oracle. We will show the latter machine can simulate the former.

$M_1^{\mathsf{GSED}_h}$ can make at most $O(\mathrm{poly}(n))$ length queries to the oracle, corresponding to α, β queries such that $\beta - \alpha = \Omega(2^{-p(n)})$ for some polynomial p. After making at most $\mathrm{poly}(n)$ queries, it performs some post-processing and finally outputs an answer.

M_2^{NEEXP} can simulate this by simply calculating $\mathcal{E}_\rho$ for the Hamiltonian in definition 4.14 by querying the NEEXP oracle for the first $O(\log(n))$ instances, and then computing an estimate for $\mathcal{E}_\rho$, denoted $\tilde{\mathcal{E}}_\rho$, using equation eq. (4.18). By making sufficiently many queries to the NEEXP oracle, one can make it so $|\tilde{\mathcal{E}}_\rho - \mathcal{E}_\rho| = O(2^{-q(n)})$ for some polynomial q. Thus by making $q(n) \gg p(n)$, M_2 can then simulate all the queries that $M_1^{\mathsf{GSED}_h}$ makes, do the same post-processing, and output the same answer. □

Theorem 4.9. *Using the notation defined in lemma 4.13, if* $\mathsf{NP} \subseteq \mathsf{P}^{\mathsf{GSED}_h}$*, then the polynomial hierarchy collapses to* Σ_2^P.

Proof. From lemma 4.13, $\mathsf{P}^{\mathsf{GSED}_h} \subseteq \mathsf{P}^{\mathsf{NEEXP}}_{\log}$. Now note $\mathsf{P}^{\mathsf{NEEXP}}_{\log} \subseteq P/\mathrm{poly}$. This is true because for an input of length n, $\mathsf{P}^{\mathsf{NEEXP}}_{\log}$ can make at most $O(\mathrm{poly}(n))$ different queries. Hence we could simply give a TM a $O(\mathrm{poly}(n))$ length advice string giving the answers to each of these queries, such that the advice string only depends on the input length n.

Thus $\mathsf{P}^{\mathsf{GSED}_h} \subseteq \mathsf{P}^{\mathsf{NEEXP}}_{\log} \subseteq P/poly$. However, it is known that if $\mathsf{NP} \subseteq P/\mathrm{poly}$, then the polynomial hierarchy collapses to Σ_2^P [KL80].

□

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

4.6.2.1 Improving the Hardness Result

We can improve our containment and hardness results by using a PSPACE oracle machine. There is, however, some controversy as to how a PSPACE oracle machine should have access to its oracle; in particular whether the input tape to the oracle has a polynomial space bound or not [Bus88; Har+93; For94]. Here we consider both of these definitions and show how they can be used to tighten our complexity bounds on GSED.

Definition 4.15 (1^{st} PSPACE Oracle Machine Definition)**.**
A PSPACE^{O} *oracle machine is a* PSPACE *machine with access to an oracle input tape, for which it can make* ***polynomial length*** *queries to the oracle.*

For this definition we get:

Theorem 4.10. $\text{PSPACE}^{\text{NEEXP}} \subseteq \text{EXP}^{\text{GSED}}$.

Proof. Identical to the proof for theorem 4.8 except M_1 is now a PSPACE machine which needs to be simulated by the EXP^{GSED} oracle machine. □

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

Definition 4.16 (2^{nd} PSPACE Oracle Machine Definition)**.**
A PSPACE^{O} *oracle machine is a* PSPACE *machine with access to a write only oracle input tape, for which it can make* ***exponential length*** *queries to the oracle.*

This is the preferred definition of several authors [LL76; For94]. For this definition of oracle machine, we realise that one can do the binary search protocol used in the proof of theorem 4.8 to get:

Theorem 4.11. $\text{P}^{\text{NEEXP}} \subseteq \text{PSPACE}^{\text{GSED}}$.

Proof. The proof will be similar to the proof for theorem 4.8, except now the PSPACE machine will have to make exponentially long oracle calls to the GSED oracle for to extract the query results while using only polynomial space everywhere else.

Let M^{GSED} be a PSPACE machine with (for convenience) two work tapes[1] (bounded by polynomial space) and one unbounded oracle tape which is read only. Let the GSED oracle be the one for the Hamiltonian of definition 4.14. Let M^{GSED} have made $(k-1)$ queries to the oracle machine with outputs $i_1, i_2 \dots i_{k-1}$, for i_j as defined in lemma 4.6, such that it now needs to make a k^{th} query. To do so, it needs to calculate a pair (α_k, β_k) which will allow it to extract i_k. Assume M has the string $i_1 i_2 \dots i_{k-1}$ stored on one of the two work tapes. We need to write out the numbers α_k, β_k in binary as given in equation eq. (4.23).

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

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

$$a_k = \frac{1}{4}\left(\sum_{n=1}^{k-1} \frac{i_n}{16^n} + \frac{2}{16^{k+1}}\right) > \alpha_k. \tag{4.27}$$

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

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

Since M only needs to record the number of queries $k = O(\mathrm{poly}(n))$ and the

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

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

This result maybe should not be too surprising given that it is known how to do binary search procedures using exponentially less space.

The results from this section immediately give:

Corollary 4.2. $\mathsf{P}^{\mathsf{NEEXP}} \subseteq \text{PSPACE}^{\mathsf{GSED}} \subseteq \text{PSPACE}^{\mathsf{NEXP}}$

4.6.2.2 Complexity Results for FGSED

We show containment of the function problem version FGSED of the ground state energy density problem:

Theorem 4.12. $\mathsf{FGSED} \in \mathsf{FP}^{\mathsf{GSED}} \subseteq \mathsf{FP}^{\mathsf{NEXP}}$ *for classical* FGSED.

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

Lemma 4.14. $\mathsf{FP}^{\mathsf{NEEXP}} \subseteq \mathsf{FEXP}^{\mathsf{FGSED}} \subseteq \mathsf{FEXP}^{\mathsf{NEXP}}$ *for classical* FGSED.

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

To show $\mathsf{FP}^{\mathsf{NEEXP}} \subseteq \mathsf{FEXP}^{\mathsf{FGSED}}$, consider a polytime oracle machine M_3^{NEEXP} and an exptime oracle machine M_4^{FGSED}. M_3 can make at most $O(\text{poly}(n))$ queries to the NEEXP oracle of at most $O(\text{poly}(n))$ length, and then do post-processing

to output the relevant function. M_4 can simulate all of these queries by asking the FGSED oracle for an estimate for ϵ such that $|\epsilon| = O(\exp(n))$, from which it can extract all the NEEXP queries. It can then do the relevant post-processing and output the same function as M_3.

□

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

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

Lemma 4.15. $\mathsf{GSED} \in \mathsf{QMA}_{\mathsf{EXP}}$ *for any quantum, nearest-neighbour, translationally invariant Hamiltonian, for* GSED *as defined in definition 4.2.*

Proof. (α, β) is the input of the problem for $\beta - \alpha = \Omega(2^{-p(n)})$. Let $|\psi\rangle$ be the ground state an $L \times L$ section of the lattice, for $L = 2^{q(n)}$, which our $\mathsf{QMA}_{\mathsf{EXP}}$ machine will take as a witness. Perform quantum phase estimation of $e^{iH^{\Lambda(L)}}$ to $q(n)$ bits of precision, which gives an estimate $\tilde{\lambda_0}$ of $\lambda_0(H^{\Lambda(L)})$ such that $|\tilde{\lambda_0} - \lambda_0(H^{\Lambda(L)})| \leq 2^{-p(n)}$, and takes time $O(2^{q(n)})$ [NC10].

Since $\mathcal{E}_\rho(L) = \tilde{\lambda}_0$, and by lemma 4.10 that $|\mathcal{E}_\rho(L) - \mathcal{E}_\rho| = O(2^{-p(n)})$, choosing $q(n)$ to be sufficiently larger than $p(n)$ allows us to verify whether $\mathcal{E}_\rho > \beta$ or $\mathcal{E}_\rho < \alpha$. □

Corollary 4.3. $\mathsf{EXP}^{\mathsf{GSED}} \subseteq \mathsf{EXP}^{\mathsf{QMA}_{\mathsf{EXP}}}$ *for a fixed, nearest-neighbour, translationally invariant* ***quantum*** *Hamiltonian.*

Proof. The proof is identical to lemma 4.12, but making use of lemma 4.15. □

Since classical Hamiltonians are a subset of quantum Hamiltonians, the following result is an immediate corollary of theorem 4.8:

Corollary 4.4. $\mathsf{P}^{\mathsf{NEEXP}} \subseteq \mathsf{EXP}^{\mathsf{GSED}}$ *for a fixed, nearest-neighbour, translationally invariant* ***quantum*** *Hamiltonian.*

4.7 Discussion

Quantum GSED A natural question to ask is if tighter results can be found for GSED for quantum Hamiltonians. As we have seen, it follows straightforwardly

that $\mathsf{EXP}^{\mathsf{GSED}} \subseteq \mathsf{EXP}^{\mathsf{QMA}_{\mathsf{EXP}}}$, but a non-trivial quantum lower bound does not follow easily.

Our proof of a $\mathsf{P}^{\mathsf{NEEXP}}$ lower bound works as we can enumerate over NEEXP-complete problems. Attempting to prove a similar quantum lower bound (e.g. $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EEXP}}}$) runs into the problem that, since $\mathsf{QMA}_{\mathsf{EEXP}}$ is a promise class, for a given $\mathsf{QMA}_{\mathsf{EEXP}}$-complete problem there may be instances which do not satisfy the promise (so called "invalid queries"). This makes it impossible to enumerate over all instances of a given $\mathsf{QMA}_{\mathsf{EEXP}}$-complete problem without potentially including instances which do not satisfy the promise. It is not currently known how to avoid these instances from occurring, although some techniques exist, such as [GY19; GPY20; WBG20].

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

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

Chapter 5

Uncomputably Complex Renormalisation Group Flows

5.1 Introduction

We've seen in Chapter 2 that there can exist Hamiltonians that undergo phase transitions at points which are uncomputable to determine. However, there exists an enormously useful tool set in theoretical physics known as the "Renormalisation Group" (RG), which is used to extract macroscale properties of systems. A natural question to ask is why can't we use RG methods to solve the phase of the model seen in the previous chapter? The uncomputability result implies that any computable method *must* fail, but it's not clear why. Is it simply because this model does not have a well defined RG scheme available? Maybe any legitimate RG scheme fails to preserve key properties. Perhaps something more perverse is going on. In this chapter we consider the application of RG techniques to the undecidable Hamiltonian in [CPGW15a; CPGW15b] and demonstrate not only that a "good" RG scheme can exist, but that such an RG scheme presents novel and previously unexplored behaviour. Unsurprisingly, this new behaviour is also uncomputable.

Determining phases of many-body systems from an underlying model of the interactions between their constituent parts remains one of the major research areas in physics, from high-energy physics to condensed matter. Indeed, the question of how to go from a microscopic model of physics, in which for N particles there are

$O(2^N)$ degrees variables to a macroscopic description of a material which requires only a few variables to describe has long befuddled scientists. Even more confusing is the fact that many materials, although differing significantly in their microscopic description, display apparently very similar behaviour.

Many powerful techniques have been developed to tackle this problem. One of the most far-reaching and fundamental was the development by Wilson [Wil71; WK74] of renormalisation group techniques, building on early work by others [BP53; GML54]. At a conceptual level, RG analysis involves constructing a map that takes as input a description of the many-body system (e.g. a Hamiltonian, or an action, or a partition function, etc.), and outputs a description of a new many-body system (a new Hamiltonian, or action, or partition function, etc.), that can be understood as a "coarse-grained" version of the original system, in such a way that physical properties of interest are preserved but irrelevant details are discarded.

For example, the RG map may "integrate out" the microscopic details of the interactions between the constituent particles described by the full Hamiltonian of the system. This procedure produces a coarse-grained Hamiltonian that still retains the same physics at larger length-scales [Kad66]. By repeatedly applying the RG map, the original Hamiltonian is transformed into successively simpler Hamiltonians, where the physics may be far easier to extract. The RG map therefore describes a dynamical map on Hamiltonians, and consecutive applications of this map generates a "flow" in the space of Hamiltonians. Often, the form of the Hamiltonian is preserved, and the RG flow can be characterised as a trajectory for its parameters.

The development of RG methods not only allowed sophisticated theoretical and numerical analysis of a broad range of many-body systems. It also explained phenomena such as *universality*, whereby many physical systems, apparently very different, exhibit the same macroscopic behaviour, even at a quantitative level. This is explained by the fact that these systems "flow" to the same fixed point under the RG dynamics — many degrees of freedom present in the microscopic description are "integrated out" as the Hamiltonian flows and simply become irrelevant.

For many condensed matter systems — even complex strongly interacting ones

— the RG dynamics are relatively simple, exhibiting a finite number of fixed points to which the RG flow converges. Hamiltonians that converge to the same fixed point correspond to the same phase, so that the basins of attraction of the fixed points map out the phase diagram of the system. However, more complicated RG dynamics are also possible, including chaotic RG flows with highly complex structure [MBK82; SKS82; DEE99; DT91; MN03]. Nonetheless, as with chaotic dynamics more generally, the structure and attractors of such chaotic RG flows can still be analysed, even if specific trajectories of the dynamics may be highly sensitive to the precise starting point. This structure elucidates much of the physics of the system [GP83; ER85; SS14]. RG techniques have become one of the most important technique in modern physics for understanding the properties of complex many-body systems.

On the other hand, the work presented in chapter 3 of this thesis, as well as previous work [CPGW15a; CPGW15b; Bau+18b] has shown that determining the macroscopic properties of many-body systems, even given a complete underlying microscopic description, can be uncomputable or undecidable. The result implies that any RG technique which we may apply to this specific system to characterise the spectrum and other properties is bound to fail: there can be no RG scheme — or even more broadly, no algorithm — that can answer the spectral gap problem. Yet, it is unclear *how* such a negative result will emerge. In principle, this obstacle may be due to the fact that there may not exist an RG map which can compute a coarse-grained version of an intractable Hamiltonian, or which cannot retain its macroscopic properties at every iteration, or again whose fixed points are not well-defined (or do not exist to begin with). In this work we prove that there are no such critical obstacles: we show that a legitimate RG procedure actually exists by providing an explicit construction. At the same time we illustrate in which way it fails in determining the spectral gap. More specifically, we construct an RG map for the Hamiltonian of [CPGW15a] which has the following features: (i) The RG map is computable at each renormalisation step. (ii) The RG map preserves whether the Hamiltonian is gapped or gapless. (iii) The Hamiltonian is guaranteed to converge to one of two fixed points under the RG flow: one gapped, with low energy properties similar to

those of an Ising model with field; the other gapless, with low energy properties similar to the critical XY-model. (iv) The behaviour of the Hamiltonian under the RG mapping, the trajectory of the RG flow and which fixed point it converges to are all uncomputable. Our RG scheme is based on the the block renormalisation group (BRG) [Jul+78; JP79; PJP82; BS99].

Remarkably, our RG scheme exhibits a novel type of behaviour, displaying a qualitatively new and more extreme form of unpredictability than chaotic RG flows: unpredictability of chaotic systems arises from the fact that even a tiny difference in the initial system parameters — which in practice may not known exactly — can eventually lead to exponentially diverging trajectories (see fig. 5.1b). However, the more precisely the initial parameters are known, the longer it is possible to accurately predict it, and if the system parameters are perfectly known it is in principle possible to determine the long-time behaviour of the RG flow. The RG flow behaviour exhibited in this work is more intractable still. Even if we know the *exact* initial values of all system parameters, its RG trajectory and the fixed point it ultimately ends up at is provably impossible to predict. Moreover, no matter how close two initial parameters are, it is impossible to predict how long their trajectories will remain close together before abruptly diverging to different fixed points that correspond to separate phases (see fig. 5.1a). Thus, the structure of the RG flow — e.g. the basins of attraction of the fixed points — is so complex that it cannot be computed or approximated, even in principle. We note that a similar form of unpredictability has previously been seen in classical single-particle dynamics, in seminal work by Moore [Moo90; Moo91; Ben90], while our result shows for the first time that this extreme form of unpredictability can occur in RG flows of many-body systems.

It is worth emphasising here that work in this chapter is not attempting to use RG methods to determine properties of the Hamiltonian used in [CPGW15a], rather we wish to determine general properties about RG techniques by analysing their behaviour when applied to Hamiltonians with undecidable properties. Finally we note that this work can be found as a preprint here: [WOC21].

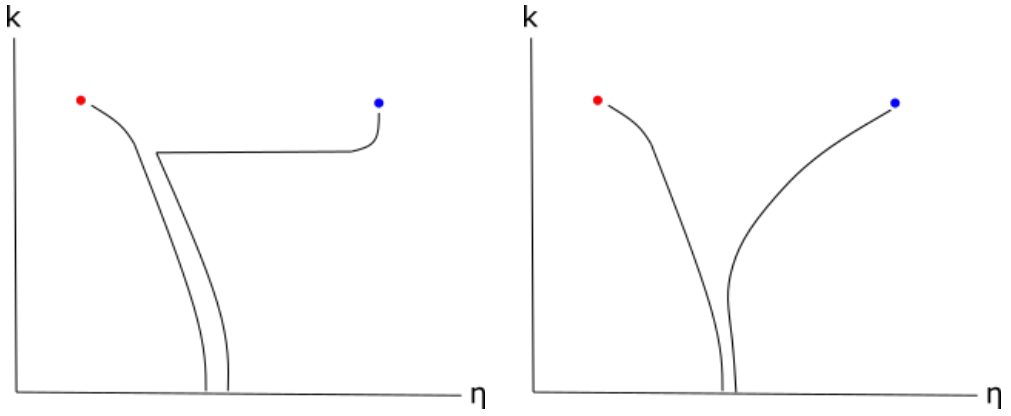

(a) Uncomputable RG flow. **(b)** Chaotic RG flow.

Figure 5.1: In both diagrams, k represents the number of RG iterations and η represents some parameter characterising the Hamiltonian; the blue and red dots are fixed points corresponding to different phases. We see that in the chaotic case, the Hamiltonians diverge exponentially in k, according to some Lyapunov exponent. In the undecidable case, the Hamiltonians remain arbitrarily close for some uncomputably large number of iterations, whereupon they suddenly diverge to different fixed points.

5.2 Preliminaries and Previous Work

5.2.1 Notation

We denote the renormalisation group map by $\mathcal{R}$, and the k-fold iteration of this map by $\mathcal{R}^{(k)}$. We will denote renormalised quantities and operators with R or $R^{(k)}$ prefix for the renormalised and k-times renormalised cases respectively. For example, denote the renormalised Hamiltonians terms as $R(h^{row})^{i,i+1}$ and $R(h^{col})^{j,j+1}$, and the local terms after k-fold iterations as $R^{(k)}(h^{row})^{i,i+1}$ and $R^{(k)}(h^{col})^{j,j+1}$. We then denote the Hamiltonian defined over the lattice by the renormalised interactions as $R(H)^{\Lambda(L)}$, and for the k-times iteration as $R^{(k)}(H)^{\Lambda(L)}$. We note that in general $\mathcal{R}(h^{row}_{i,i+1}) \neq R(h^{row})^{i,i+1}$, and similarly for the other terms.

If the initial local Hilbert space is $\mathcal{H}$, then the local Hilbert space after k iterations of the RG map is denoted $R^{(k)}(\mathcal{H})$. Throughout, we will denote a canonical set of local basis states by $\mathfrak{B}$, and after the renormalisation mapping has been applied k times it becomes $\mathfrak{B}^{(k)}$, so that $R^{(k)}(\mathcal{H}) = \text{span}\{|x\rangle \in \mathfrak{B}^{(k)}\}$.

It will occasionally be useful to distinguish h^{row} acting on given row j. When this is important, we write $h^{row}_{i,i+1}(j)$ to denote the interaction between columns i and

$i+1$ in the j^{th} row. Similarly $h^{col}_{j,j+1}(i)$ denotes the interaction between rows j and $j+1$ in the i^{th} column.

Finally, following [CPGW15a], we adopt the following precise definitions of gapped and gapless given in definition 3.1 and definition 3.2. In [CPGW15a] it was shown that the particular Hamiltonians they construct always fall into one of these clear-cut cases, allowing sharp spectral gap undecidability results to be proven.

5.2.2 Real Space Renormalisation Group Maps

The notion of what exactly constitutes a renormalisation group scheme is somewhat imprecise, and there is no universally agreed upon definition in the literature. We therefore start from a minimal set of conditions that we would like a mapping on Hamiltonians to satisfy, if it is to be considered a reasonable RG map. The RG scheme we define for the Hamiltonian from [CPGW15a] will satisfy all these conditions as well as additional desirable properties.

Definition 5.1 (Renormalisation Group (RG) Map)**.** *Let $\{h_i\}_i$ be an arbitrary set of r-local interactions $h_i \in \mathcal{B}((\mathbb{C}^d)^{\otimes r})$, for $r = O(1)$ and $d \in \mathbb{N}$. A renormalisation group (RG) map*

$$\mathcal{R}(\{h_i\}) = \{h'_i\} \tag{5.1}$$

is a mapping from one set of r-local interactions to a new set of r'-local interactions $h'_i \in \mathcal{B}((\mathbb{C}^{d'})'^{\otimes r'})$, with $r', d' \in \mathbb{N}$, satisfying the following properties:

1. *$\mathcal{R}(\{h_i\})$ is a computable map.*
2. *Let H and $R^{(k)}(H)$ be the Hamiltonian defined by the original local terms and the k-times renormalised local terms respectively. If H is gapless, then $R^{(k)}(H)$ is gapless, as per definition 3.2. If H is gapped, then $R^{(k)}(H)$ is gapped, as per definition 3.1.*
3. *If the order parameter for the system has a non-analyticity between two phases of H, then there is a renormalised order parameter which also has a non-analyticity between the two phases for $R^{(k)}(H)$.*

4. *If the initial local Hamiltonian terms can decomposed into as*

$$h_i = \sum_j \alpha_j O_j, \tag{5.2}$$

for some operator $\{O_j\}_j$*, then* k*-times renormalised local Hamiltonian terms are of the form*

$$R^{(k)}(h)_i = \sum_j \alpha_j^{(k)} R^{(k)}(O)_j, \tag{5.3}$$

where $\alpha_i^{(k)} = f(\{\alpha_i^{(k-1)}\}_i)$ *for some function* f*.*

The motivation for points 2 and 3 of definition 5.1 is that we want to preserve the quantum phase diagram of the system. Point 3 of definition 5.1 requires that if we start in phase A, the system should remain in phase A under the RG flow: a key property of any RG scheme. Furthermore, any indicators of a phase change still occur (e.g. non-analyticity of the order parameter). Point 4 asks that the "form" of the Hamiltonian is preserved.

Hamiltonians under RG flows have "fixed points" which occur where the Hamiltonian is left invariant by the action of the RG procedure. If H^* is the fixed point a particular Hamiltonian is converging to under the RG flow, and h^* is the corresponding local term, then the local terms away from the fixed point can be rewritten in terms of their deviation from the fixed point as:

$$h = h^* + \sum_i \beta_i O_i \tag{5.4}$$

and after renormalisation

$$R^{(k)}(h) = h^* + \sum_i \beta_i^{(k)} O_i', \tag{5.5}$$

where if $\beta_i^{(k)} \to 0$ as $k \to \infty$ then O_i is said to be an *irrelevant operator*; if $\beta_i^{(k)} \to \infty$, then O_i is a *relevant operator*; and if $\beta_i^{(k)} \to c$ for a constant c, then O_i a *marginal operator*.

We note that many well-known renormalisation group schemes fit the criteria given in definition 5.1 when applied to the appropriate Hamiltonians. In the following subsections, we review a number of these. However, in general, a given RG scheme may satisfy the conditions for the family of Hamiltonians it was designed for, but will not necessarily satisfy all the desired conditions when applied to an arbitrary Hamiltonian.

5.2.2.1 The Block Spin Renormalisation Group Map

We base our RG map on a blocking technique widely used in the literature to study spin systems, often called the Block Spin Renormalisation Group (BRG)[1] [Jul+78; JP79; PJP82; BS99]. Modifications and variations of this RG scheme have also been extensively studied [MDS96; WKL02].

The BRG is among the simplest RG schemes. The procedure works by grouping nearby spins together in a block, and then determining the associated energy levels and eigenstates of this block by diagonalisation. Having done this, high energy (or otherwise unwanted) states are removed resulting in a new Hamiltonian.

As an explicit example, suppose there exists a Hamiltonian on a 1D chain

$$H = \sum_{i=1}^{N-1} K^{(0)} h^{(0)}_{i,i+1} + C^{(0)} \sum_{i=1}^{N} \mathbb{1}_i. \tag{5.6}$$

The BRG first groups the lattice points into pairs

$$H = K^{(0)} \sum_{i \text{ odd}}^{N-1} h^{(0)}_{i,i+1} + K^{(0)} \sum_{i \text{ even}}^{N-1} h^{(0)}_{i,i+1} + C^{(0)} \sum_{i=1}^{N} \mathbb{1}_i. \tag{5.7}$$

We then diagonalise the operators for odd i. (In higher dimensional geometries we group the terms into blocks of neighbouring qudits.) Having done this, remove all "high energy states" within each block, either by introducing an energy cut-off or just keeping a chosen subset of the lowest energy states. The produces a renormalised

[1] This is also sometimes called the "quantum renormalisation group".

Hamiltonian

$$R^{(1)}(H) = K^{(1)} \sum_{i=1}^{N/2-1} h_{i,i+1}^{(1)} + b^{(1)} \sum_{i=1}^{N/2} h_i^{(1)} + C^{(1)} \sum_{i=1}^{N/2} \mathbb{1}_i. \tag{5.8}$$

For each further RG iteration the same process is repeated: the terms $h_{i,i+1}$ for odd i are diagonalised and the high energy states are removed.

After k iterations, the RG procedure returns a Hamiltonian of the same form, but now with different coupling constants:

$$R^{(n)}(H) = K^{(n)} \sum_{i=1}^{N/2-1} h_{i,i+1}^{(n)} + b^{(n)} \sum_{i=1}^{N/2} h_i^{(n)} + C^{(n)} \sum_{i=1}^{N/2} \mathbb{1}_i. \tag{5.9}$$

5.2.3 Properties of the Spectral Gap Undecidability Construction

Constructing a mathematically rigorous RG flow for the undecidable Hamiltonian exhibited in [CPGW15b; CPGW15a] presents particular challenges, since its properties are uncomputable. Nonetheless, we are able to construct such a scheme by carefully analysing the local structure and properties of this Hamiltonian. We review the structure here.

We start by stating the main result in [CPGW15a], where the authors construct a Hamiltonian depending on one external parameter, which is gapped iff a universal Turing Machine halts on an input related to the Hamiltonian parameter. The spectral gap problem for this Hamiltonian is therefore equivalent to the Halting Problem, hence undecidable.

Definition 5.2 (From theorem 3 of [CPGW15a])**.** *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', Π and a rational number β which can be as small as desired, with the following properties:*

1. *A is diagonal with entries in $\mathbb{Z}$.*
2. *A' is Hermitian with entries in $\mathbb{Z} + \frac{1}{\sqrt{2}}\mathbb{Z}$,*
3. *B, C have integer entries,*

4. *D is diagonal with entries in $\mathbb{Z}$,*

5. *D' is Hermitian with entries in $\mathbb{Z}$.*

6. *Π is a diagonal projector.*

For each natural number n, define:

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

where $\alpha(n) \leq \beta$ is an algebraic number computable from n and $|\varphi|$ denotes the length of the binary representation of φ. Then:

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

2. *If UTM halts on input n, then the associated family of Hamiltonians $\{H^{\Lambda(L)}(n)\}$ is gapped with gap $\gamma \geq 1$.*

3. *If UTM does not halt on input n, then the associated family of Hamiltonians $\{H^{\Lambda(L)}(n)\}$ is gapless.*

We first explain the overall form of the Hamiltonian and the Hilbert space structure, and later how the individual parts fit together.

5.2.3.1 Local Interaction Terms and Local Hilbert Space Structure

The Hamiltonian $H_u(\varphi)$ is constructed such that its ground state is composed of two components: a classical "tiling layer" and a highly entangled "quantum layer". The local Hilbert space decomposes as:

$$\mathcal{H}_u = \mathcal{H}_c \otimes (\mathcal{H}_q \oplus |e\rangle), \tag{5.10}$$

where $\mathcal{H}_c$ is the Hilbert space corresponding ot the classical tiling layer and $\mathcal{H}_q \oplus |e\rangle$ is the "quantum" layer. The local terms h_u are constructed as

$$h_u = h_T^{(i,i+1)} \otimes \mathbb{1}_{eq}^{(i)} \otimes \mathbb{1}_{eq}^{(i+1)} + \mathbb{1}_c^{(i)} \otimes \mathbb{1}_c^{(i+1)} \otimes h_q^{(i,i+1)} + \text{"coupling terms"}. \tag{5.11}$$

Let $h_u^{(i,j)} \in \mathcal{B}(\mathbb{C}^d \otimes \mathbb{C}^d)$ be the local terms of the Hamiltonian H_u, $h_d^{(i,j)} \in \mathcal{B}(\mathbb{C}^2 \otimes \mathbb{C}^2)$ be the local interactions of the 1D critical XY model, and let H_d be the Hamiltonian composed of XY interactions along the rows of the lattice. This has a dense spectrum in the thermodynamic limit [LSM61]. $h_u^{(i,j)} = h_u^{(i,j)}(\varphi)$ is designed so that $H_u(\varphi) = \sum h_u(\varphi)$ has a ground state energy which depends on whether a universal Turing Machine (UTM) halts when given on input φ supplied in binary. In particular, on a lattice of size $L \times L$, the ground state energy is

$$\lambda_0(H_u^{\Lambda(L)}) = \begin{cases} -\Omega(L) & \text{if UTM does not halt on input } \varphi, \\ +\Omega(L^2) & \text{if UTM does halt on input } \varphi. \end{cases} \tag{5.12}$$

Since the halting problem is undecidable, determining which of the two ground state energies of $H_u(\varphi)$ occurs is undecidable.

The local Hilbert space of the overall Hamiltonian can be decomposed as:

$$\mathcal{H} = |0\rangle \oplus \mathcal{H}_u \otimes \mathcal{H}_d. \tag{5.13}$$

Here $|0\rangle$ is a zero-energy filler state, $\mathcal{H}_d$ is the Hilbert space associated with the dense spectrum Hamiltonian h_d, and $\mathcal{H}_u$ is the Hilbert space associated with the Hamiltonian with undecidable ground state energy h_u.

The local interactions along the edges and on the sites of the lattice are act on this local Hilbert space as:

$$h(\varphi)^{(i,j)} = |0\rangle\langle 0|^{(i)} \otimes (\mathbb{1} - |0\rangle\langle 0|)^{(j)} + h_u^{(i,j)}(\varphi) \otimes \mathbb{1}_d^{(i,j)} + \mathbb{1}_u^{(i,j)} \otimes h_d^{(i,j)} \tag{5.14}$$

$$h(\varphi)^{(1)} = -(1+\alpha_2)\Pi_{ud}, \tag{5.15}$$

where Π_{ud} is a projector onto $\mathcal{H}_u \otimes \mathcal{H}_d$, and $\alpha_2 = \alpha_2(|\varphi|)$ is a constant depending

only on $|\varphi|$. Importantly, the spectrum of the overall lattice Hamiltonian composed of these local interactions is

$$\operatorname{spec} H(\varphi) = \{0\} \cup \{\operatorname{spec}(H_u(\varphi)) + \operatorname{spec}(H_d)\} \cup S, \tag{5.16}$$

for a set S with all elements > 1. This means that if $\lambda_0(H_u^{\Lambda(L)}) \to -\infty$ then the overall Hamiltonian has a dense spectrum, while if $\lambda_0(H_u^{\Lambda(L)}) \to +\infty$ the overall Hamiltonian has a spectral gap > 1.

In the $\lambda_0(H_u^{\Lambda(L)}(\varphi)) = +\Omega(L^2)$ case, the ground state of the entire Hamiltonian is $|0\rangle^{\Lambda}$. In the $\lambda_0(H_u^{\Lambda(L)}(\varphi)) = -\Omega(L)$ case, the overall ground state is $|\psi_u\rangle \otimes |\psi_d\rangle$ where $|\psi_u\rangle$ and $|\psi_d\rangle$ are the ground states of $H_u(\varphi)$ and $H_d = \sum_{i\in\Lambda} h_d^{i,i+1}$ respectively.

We now explain the terms h_T and h_q as well as the cumulative effects of the coupling terms.

The Tiling Hamiltonian In [CPGW15a] a Wang tile set is chosen to be a slightly modified version the Robinson tiles from [Rob71], shown in fig. 5.2a. We refer the reader to section 4.4 for a summary of Robinson tiles. When placed on a 2D grid such that the tiling rules are satisfied, the markings on the tiles form an aperiodic tiling consisting of a series of nested squares of sizes $4^n + 1$, for all $n \in \mathbb{N}$, as shown in fig. 5.2b.

This set of tiles can then be mapped to a 2D, translationally invariant, nearest neighbour, classical Hamiltonian by the procedure explained previously in section 3.7.1. Then, the ground state of the entire 2D lattice, $|T\rangle_c$, corresponds to the Robinson tiling pattern as shown in fig. 5.2b. Any other configuration must violate a tiling rule and thus receives an energy penalty. We denote the local interaction terms of this Hamiltonian as $h_T \in \mathcal{B}(\mathcal{H}_c \otimes \mathcal{H}_c)$.

The Quantum Hamiltonian $H_u(\varphi)$ is constructed so that its ground state energy encodes the halting or non-halting of a computation as per the circuit-to-Hamiltonian construction used in chapter 3. The fundamental ingredient required is the "QTM-to-Hamiltonian" mapping, as explained in section 2.3.1.1 [GI09; CPGW15a]. This takes a given quantum QTM and creates a corresponding Hamiltonian which has a

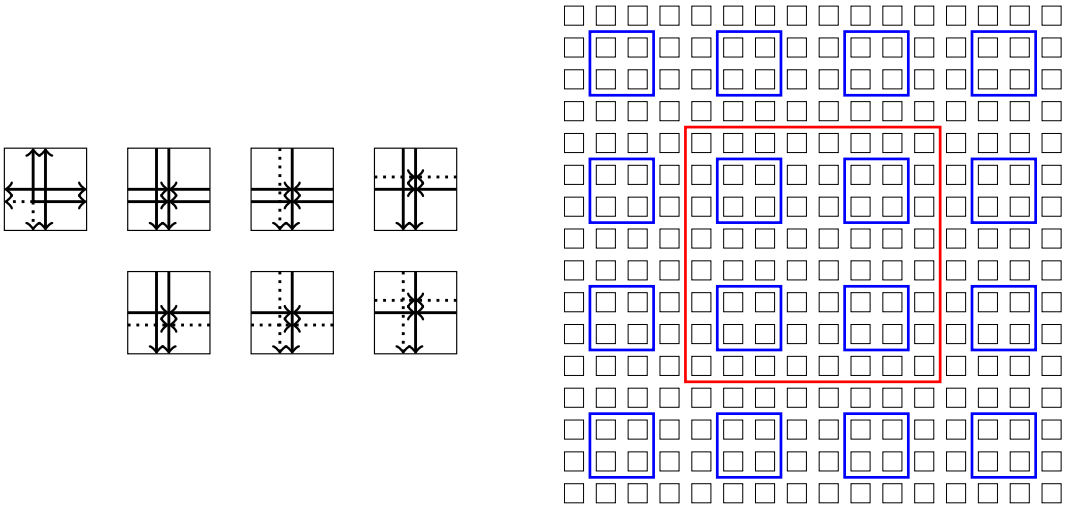

(a) The modified Robinson tiles used in [CPGW15a].

(b) Ground state $|T\rangle_c$ of the classical Hamiltonian.

Figure 5.2

ground state which encodes its evolution. It is then possible to add a local projector term to the Hamiltonian which gives an additional energy penalty to certain outcomes of the computation. In particular, [CPGW15a] penalise the halting state, so that if the QTM which is encoded halts at some point, the Hamiltonian defined by h_q picks up an additional energy contribution. As a result, the energy of the ground state differs depending on whether or not the QTM halts within time T.

In particular, [CPGW15a] adapt the QTM-to-Hamiltonian mapping originally developed by Gottesman and Irani [GI09], which takes a QTM and maps its evolution to 1D, translationally invariant, nearest neighbour, Hamiltonian. By H_q we denote this modified version of the *Gottesman-Irani Hamiltonian* (cf. section 5.2.4).

The length of the computation encoded on a chain of length L is $T(L) \sim \text{poly}(L)2^L$, and the associated ground state energy is

$$\lambda_0(H_q(L)) = \begin{cases} 0 & \text{if QTM is non-halting within time } T(L), \\ \theta(1/T^2) & \text{if QTM halts within time } T(L). \end{cases} \tag{5.17}$$

We give a more detailed analysis of the construction at the beginning of section 5.5.

Combining h_T, h_q and the Coupling Terms The terms h_u are designed so that all eigenstates of $H_u^{\Lambda(L)}$ are product states $|T\rangle_c \otimes |\psi\rangle_{eq}$ where $|T\rangle \in \mathcal{H}_c^{\otimes(L\times L)}$ and

$|\psi\rangle \in \mathcal{H}_{eq}^{\otimes(L\times L)}$ [CPGW15a, Lemma 51].

Furthermore, the coupling terms are chosen such that the ground state has the following properties:

1. the classical part of the ground state $|T\rangle_c$ corresponds to a perfect Robinson tiling. The pattern created has a series of nested red Robinson squares as per fig. 5.2b.

2. the quantum part of the ground state $|\psi\rangle_{eq}$ has the following structure: along the top of every red Robinson square there is a history state (as defined in definition 2.14); everywhere which is not along the top of a square is in the zero energy filler state $|e\rangle_e$.

The consequence of this is that ground states of $H_q(\ell)$ of all lengths $\ell \in 4^N+1, n \in \mathbb{N}$, appear with a constant density across the lattice. If, for any length, the encoded computation halts, then the ground state picks up a constant energy density, so that the energy scales as $\Omega(L^2)$. However, if the encoded computation never halts, then for all lengths the ground state of the Gottesman-Irani Hamiltonian has zero energy, and (due to boundary effects), the ground state has energy $-\Omega(L)$ [CPGW15a].

5.2.4 The Gottesman-Irani Hamiltonian

The particular circuit-to-Hamiltonian mapping used in the previous section will be important when it comes to renormalising the overall Hamiltonian. The overall structure used in [CPGW15a] is a modification of the one used in [GI09]. Following [GI09], the QTM can be encoded into a 1D, translationally-invariant, nearest-neighbour Hamiltonian, which we refer to as *a Gottesman-Irani Hamiltonian*, denoted by $H_q(L) \in \mathcal{B}((\mathbb{C}^d)^{\otimes L})$.

This is summarised by theorem 32 of [CPGW15a]; we write out a slightly simpler version here.

Theorem 5.1 (Informal Version of Theorem 32 of [CPGW15a])**.**
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 $\left|\circledless\right\rangle, \left|\circledgtr\right\rangle$. Denote the orthogonal complement of span$(\left|\circledless\right\rangle, \left|\circledgtr\right\rangle)$ *in $\mathbb{C}^d$ by $\mathbb{C}^{d-2}$. Let d, Q and C all be fixed.*

For any well-formed unidirectional Quantum Turing Machine $M = (\Sigma, Q, \delta)$ *and any constant* $K > 0$*, we can construct a two-body interaction* $h \in \mathcal{B}(\mathbb{C}^d \otimes \mathbb{C}^d)$ *such that the 1-dimensional, translationally-invariant, nearest-neighbour Hamiltonian* $H(L) = \sum_{i=1}^{L-1} h^{(i,i+1)} \in \mathcal{B}(\mathcal{H}(L))$ *on the chain of length* L *has the following properties:*

1. *d depends only on the alphabet size and number of internal states of M.*

2. *$h \geq 0$, and the overall Hamiltonian $H(L)$ is frustration-free for all L.*

3. *Denote $\mathcal{H}(L-2) := (\mathbb{C}^{d-2})^{\otimes L-2}$ and define $\mathcal{S}_{br} = \text{span}(|\circledcirc\rangle) \otimes \mathcal{H}(L-2) \otimes \text{span}(|\circledcirc\rangle) \subset \mathcal{H}$. Then the unique ground state of $H(L)|_{\mathcal{S}_{br}}$ is a computational history state (cf. definition 2.14 for a definition) encoding the evolution of M.*

Moreover, the action of M satisfies:

1. *The computational history state always encodes $\Omega(2^L)$ time-steps. If M halts in fewer than the number of encoded time steps, exactly one $|\psi_t\rangle$ has support on a state $|\top\rangle$ that encodes a halting state of the QTM. The remaining time steps of the evolution encoded in the history state leave M's tape unaltered, and have zero overlap with $|\top\rangle$.*

2. *If M runs out of tape within a time T less than the number of encoded time steps, 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.*

We refer the reader to [CPGW15a; GI09] for a detailed overview. We provide in the following a more detailed sketch of how the modified Gottesman-Irani construction works, and refer the reader to [CPGW15a; GI09] for a detailed overview. We begin by considering the general setup. Our basis states for $(\mathbb{C}^d)^{\otimes L}$ (i.e. the chain of length L) will have the following structure:

⋖	⋯	Track 1: Clock oscillator	⋯	⋗
⋖	⋯	Track 2: Counter TM head and state	⋯	⋗
⋖	⋯	Track 3: Counter TM tape	⋯	⋗
⋖	⋯	Track 4: QTM head and state	⋯	⋗
⋖	⋯	Track 5: QTM tape	⋯	⋗
⋖	⋯	Track 6: Time-wasting tape	⋯	⋗

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

The outline of the construction is the following: tracks 1 encodes the evolution of an oscillator which goes back and forth along its track as per fig. 5.3 Tracks 2 and 4 contain the heads of a classical and quantum TM respectively. These heads are only able to move when the oscillator on track 1 passes by their heads – in this way their evolution can be encoded with only local Hamiltonian terms. Tracks 3 and 5 are the read/write tapes for the respective TMs.

The classical TM encoded by the track 2 head will be a simple counter: it will write out binary number on its tape (on track 3) and then increment it by one to the next binary number. The head on track 2 is only able to make a transition when the oscillator head passes next to it in the track above. This continues until the tape is filled, at which point it halts along with the clock oscillator.

The QTM on tracks 4 and 5 will be a generic QTM. The QTM evolves as per its transition rules until either: (a) the counter TM runs out of space and hence the

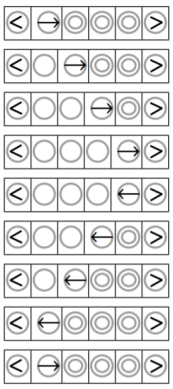

Figure 5.3: Evolution of the Track 1 clock oscillator, from [GI09].

oscillator stops, or (b) the QTM finishes its computation and halts. If the QTM halts before the counter TM runs out of steps, it places a halting marker on track 5. The head then moves to track 6 where it performs some arbitrary time wasting computation which is guaranteed not to halt before the counter TM.

We also note that tracks 1-3 evolved entirely classically whereas tracks 4-6 will contain quantum states. As such, we decompose the local Hilbert space into a classical and quantum part $\mathbb{C}^C \otimes \mathbb{C}^Q$.

5.2.5 Order Parameters

As per claim 3 of definition 5.1, we now discuss order parameters in more detail. It can be checked that the Hamiltonian in [CPGW15a] has an order parameter for its two phases[2] (which we label A and B for convenience) which can be distinguished by an order parameter $O_{A/B}$, defined as:

$$O_{A/B} = \frac{1}{|\Lambda|} \sum_{i \in \Lambda} |0\rangle\langle 0|^{(i)} . \tag{5.18}$$

Note this order parameter is identical to the one for the Hamiltonian in chapter 3. In particular, upon moving from one phase to another, the expectation value of the order parameter is expected to undergo a non-analytic change. In the case $\lambda_0(H_u^{\Lambda(L)}(\varphi)) = +\Omega(L^2)$ the ground state of the entire Hamiltonian is then $|0\rangle^\Lambda$ and hence $\langle O_{A/B}\rangle = 1$, and otherwise $\langle O_{A/B}\rangle = 0$. This is true even if we restrict $O_{A/B}$ to subsections of the lattice, hence $O_{A/B}$ is a local order parameter (as opposed to the global order parameters required to distinguish topological phases). Thus $O_{A/B}$ undergoes a non-analytic change between phases, which itself demonstrates a phase transition. More generally for a ball $B(r)$ of radius r, and for a state $|\nu\rangle \in \mathcal{H}^{\otimes\Lambda}$ we can define a local observable

$$O_{A/B}(r) = \frac{1}{|B(r)|} \sum_{i \in B(r)} |0\rangle\langle 0|^{(i)} , \tag{5.19}$$

which acts as a local order parameter.

[2] Phase in this context refers to the state of matter, not a quantum mechanical phase factor (of the form $e^{i\theta}$).

5.3 Results and Overview of RG Procedure

The main results of this chapter are the following set of theorems and corollaries:

Theorem 5.2 (Exact RG flow for Undecidable Hamiltonian)**.** *Let $H(\varphi)$ be the Hamiltonian defined in [CPGW15a]. We construct a renormalisation group procedure for the Hamiltonian which has the following properties:*

1. *$\mathcal{R}$ is computable.*

2. *If $H(\varphi)$ is gapless, then $R^{(k)}(H(\varphi))$ is gapless, and if $H(\varphi)$ is gapped, then $R^{(k)}(H(\varphi))$ is gapped (where gapped and gapless are defined in definition 3.1 and definition 3.2).*

3. *For the order parameter $O_{A/B}(r)$ (as defined in eq. (5.19)) which distinguishes the phases of $H^{\Lambda(L)}$ and is non-analytic at phase transitions, there exists a renormalised observable $R^{(k)}(O_{A/B}(r))$ which distinguishes the phases of $R^{(k)}(H)^{\Lambda(L)}$ and is non-analytic at phase transitions.*

4. *Under an arbitrary number of iterations, the renormalised local interactions belong to a family $\mathcal{F}(\varphi, \tau_1, \tau_2, \{\alpha_i\}_i, \{\beta_i\}_i)$, and for any finite k all of the parameters are computable.*

5. *If $H(\varphi)$ initially has algebraically decaying correlations, then $R^{(k)}(H(\varphi))$ also has algebraically decaying correlations. If $H(\varphi)$ initially has zero correlations, then $R^{(k)}(H(\varphi))$ also has zero correlations.*

Theorem 5.3 (Uncomputability of RG flows)**.** *Let $h(\varphi)$, $\varphi \in \mathbb{Q}$, be the full local interaction of the Hamiltonian from [CPGW15a]. $H(\varphi) := \sum h(\varphi)^{(i,j)}$ is gapped if the UTM corresponding to $h(\varphi)$ halts on input φ, and gapless if the UTM never halts, where gapped and gapless are defined in definition 3.1 and definition 3.2. Using the RG scheme defined in theorem 5.2 then under k iterations of the RG scheme (defined later in definition 5.14) acting on $H(\varphi)$, the renormalised local terms are given by $R^{(k)}(h(\varphi))$, which can be parameterised as part of a family $\mathcal{F}(\varphi, \tau_1, \tau_2, \{\alpha_i\}_i, \{\beta_i\}_i)$ (defined later in corollary 5.4). If the UTM is non-halting on input φ, for all $k > k_0(\varphi)$,*

$\tau_2(k) = -2^k$, for some computable $k_0(\varphi)$. If the UTM is halting on input φ, then there exists an uncomputable $k_h(\varphi)$ such that for $k_0(\varphi) < k < k_h(\varphi)$, $\tau_2(k) = -2^k$, and for all $k > k_h(\varphi)$ then $\tau_2(k) = -2^k + \Omega(4^{k-k_h(\varphi)})$.

A direct consequence of this is:

Corollary 5.1. *Determining which fixed point the Hamiltonian flows to under this RG scheme is undecidable.*

The overall RG scheme is explicitly given in definition 5.14, and the family $\mathcal{F}(\varphi, \tau_1, \tau_2, \{\beta_i\})$ which the renormalised Hamiltonians belong to is given in corollary 5.4. One of the consequences of theorem 5.3 is that the Hamiltonian is guaranteed to flow towards one of two fixed points. However, determining which fixed point it flows to for a given value of φ is undeciable.

The undecidability of the fixed point follows implicitly from undecidability of the spectral gap [CPGW15a; CPGW15b], since the fixed point depends on the gappedness of the unrenormalised Hamiltonian. However, theorem 5.3 shows precisely how the trajectory of the Hamiltonian in parameter space diverges in an uncomputable manner under RG flow.

We note a subtlety in the statement of theorem 5.2. It is important that we are able to explicitly construct the RG scheme, rather than just prove the existence of such an RG scheme. If only existence were proven, it would leave open the possibility that finding the RG scheme is itself an uncomputable task, thus meaning it cannot actually be determined.

5.3.1 Overview of the Proof of the Main Results

The renormalisation group scheme we will employ will be a variant of the BRG described in section 5.2.2.1, where we block 2×2 groups of spins to a single "super-spin" which preserves some of the properties of the original set. Due to the complexity of the Hamiltonian in consideration, we will first renormalise the different parts $h_u, h_d, |0\rangle$ of the Hamiltonian separately, then combine these RG maps into the complete map. For a finite size lattice, h_u has a ground state which is product between $\mathcal{H}_C$ and $\mathcal{H}_q \oplus |e\rangle$. This key property allows us to essentially renormalise

the tiling Hamiltonian and the Gottesman-Irani Hamiltonian separately.

Renormalising the Tiling Hamiltonian

Section 4.4.1 (and fig. 4.3) shows that the ground state of the tiling Hamiltonian corresponds to a particular pattern; notably the Robinson tiling creates a self-similar pattern for across all sizes of squares, where smaller squares are nested within larger ones. We design a blocking RG procedure which takes a set of 2×2 Robinson tiles, then maps them onto a single new tile which has the same markings and tiling rules as one in the original set of Robinson tiles. Doing this we recover a set of tiles which recreate the Robinson tiling pattern, but now with the smallest squares "integrated out". Repeated iterations of this process still preserve the Robinson tiling pattern. The details are give in section 5.4.

Renormalising the Gottesman-Irani Hamiltonian

The Gottesman-Irani Hamiltonian h_q is a 1D Hamiltonian which serves as a QTM-to-Hamiltonian map. As noted in section 5.2.3, in the ground state of $\sum h_u$, ground states of Gottesman-Irani Hamiltonians appear along the top edge of the Robinson tiles. We aim to design an RG scheme such that the energy of the Gottesman-Irani ground state attached to a square remains the same even when the square size is halved. To do this, we map pairs of spins in the Gottesman-Irani Hamiltonian to a new "combined spin" which now has local Hilbert space dimension d^2 if the original dimension is d. As with the BRG, we consider the new 1-local terms and diagonalise them. Since we know the form of the ground state explicitly, it is possible to identify states which pick up too much energy to have overlap with the ground state. We can truncate the local Hilbert space by removing these states and hence reduce the dimension of the combined spin to something $< d^2$ (but still $> d$). This blocking procedure will preserve whether the Hamiltonian has a zero energy ground state or a ground state with energy > 0.

In mathematical terms, the procedure is implemented by a series of isometries which are used to map the original states to the new blocked states, and then subspace restrictions which remove the high energy states. This is summarised in lemma 5.4. We refer the reader to section 5.5 for full details.

Renormalising h_u Since $h_u = h_T^{(i,i+1)} \otimes \mathbb{1}_{eq}^{(i)} \otimes \mathbb{1}_{eq}^{(i+1)} + \mathbb{1}_c^{(i)} \otimes \mathbb{1}_c^{(i+1)} \otimes h_q^{(i,i+1)} +$ coupling terms, to renormalise it, we do the following:

- Choose a 2×2 block of spins.
- Renormalise the classical tiling part of the Hamiltonian as above.
- To renormalise the quantum part of the Hamiltonian, break the 2×2 block into two 2×1 blocks. Renormalise these two sections as the above renormalisation for the Gottesman-Irani Hamiltonian. The 2×2 block is now a 2×1 block.
- Trace out part of the Hilbert space such the 2×1 block is now a single site in the renormalised Hilbert space such that we are left with 1-local and 2-local projector terms which introduce an energy shift. This energy shift exactly compensates for any energy lost in the integrating out operation.

The above can be shown to preserve the ground state energy in the desired way. See definition 5.13 in section 5.6.1 for the complete description.

Renormalising the Entire Hamiltonian We renormalise h_d and $|0\rangle\langle 0|$ in a straightforward way such that their properties are preserved. Thus the overall renormalisation scheme acts on h_u as above, and essentially leaves h_d and $|0\rangle\langle 0|$ unchanged.

Since h_u, h_d and $|0\rangle\langle 0|$ have their respective ground state energies preserved (approximately), whether the ground state is $|0\rangle^{\Lambda}$ or the more complex ground state of the tiling+quantum Hamiltonian, is preserved. Importantly it can be shown the spectral gap of both cases is preserved. The RG process can then be iterated arbitrarily many times: we show the relevant properties are preserved throughout. Determining the properties of the ground state and spectral gap are undecidable for the unrenormalised Hamiltonian, and since these properties are preserved by the RG mapping, it is also undecidable for the renormalised Hamiltonians.

The renormalisation of the entire Hamiltonian is given in detail in section 5.6.

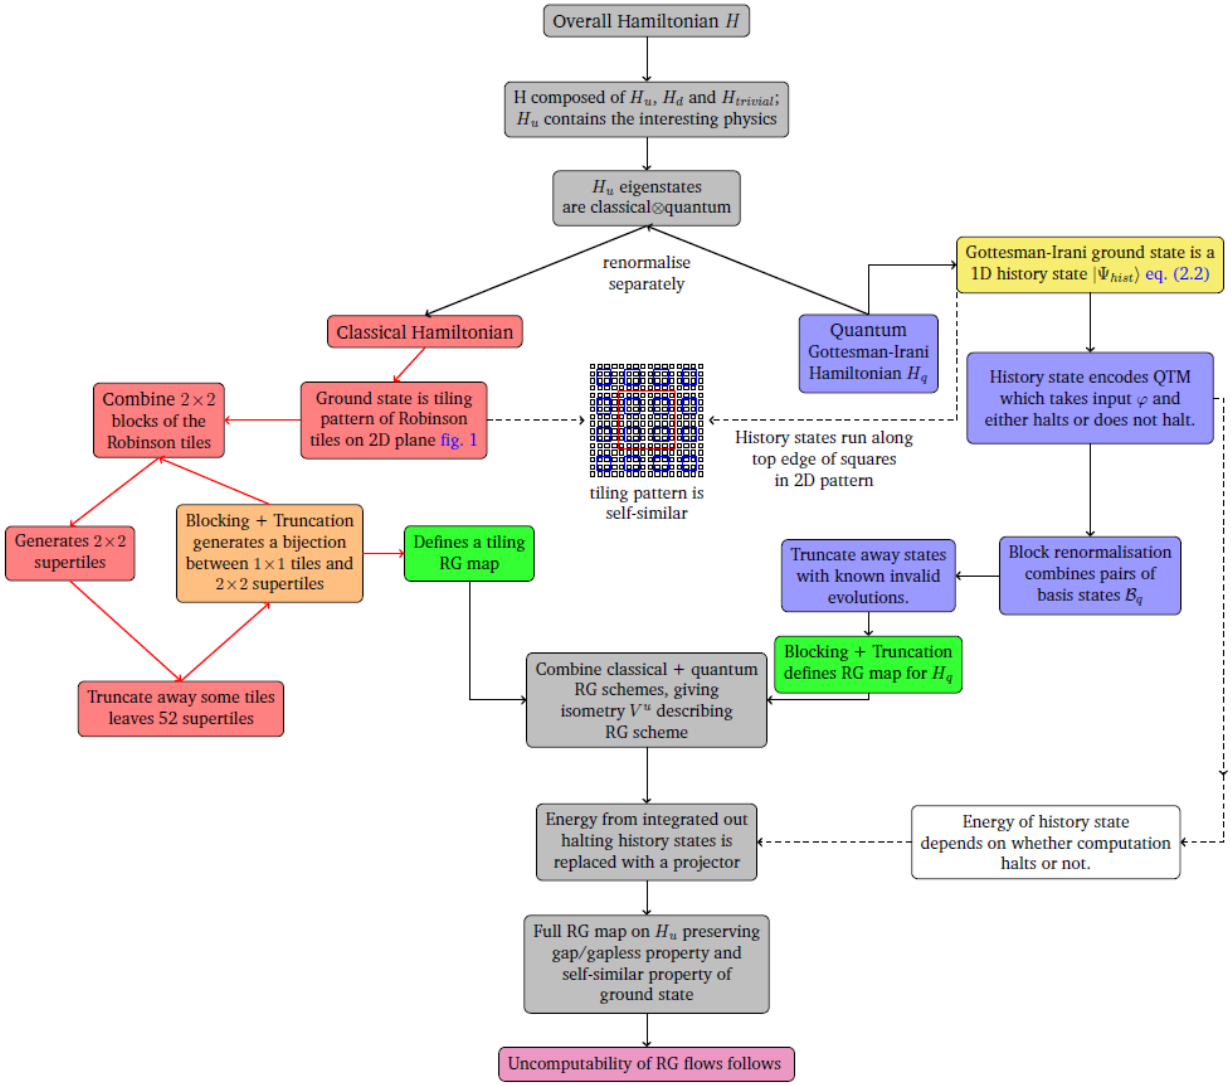

Figure 5.4: Flow diagram of the proof.

5.4 Renormalisation of the Robinson Tiling Hamiltonian

The renormalisation of the Robinson tiling and the corresponding Hamiltonian were primarily the work of Emilio Onorati, and as such we will only include a brief sketch of the result in this thesis to reduce the overall length. We refer the reader to [WOC21] for full details of the classical renormalisation procedure.

To renormalise the Robinson tiling an RG map will be constructed under which the two graphs representing respectively the *adjacency relations* (roughly speaking, the rules telling us what tiles can stay above / below / left / right of a

given tile) for the Robinson tiles and for a specific subset of 2×2 supertiles are isomorphic. This implies that the pattern produced by the tiling of the 2D plane using Robinson tiles is scale-invariant. This property is crucial in order to ensure that the density of the Gottesman-Irani ground states (corresponding to the top edges of the squares appearing in the pattern) which encode the QTM is preserved under the renormalisation procedure.

More formally, we have that:

Theorem 5.4. *(Adjacency Rules Isomorphism) Let T_1 be the set of Robinson tiles and A_1 be the corresponding adjacency rules. Let T_2 be the set of 2×2 supertiles, obtained from all combinations allowed by A_1 of four Robinson tiles placed in a 2×2 square, and A_2 be the adjacency rules of T_2, derived from the principle that two supertiles can be placed next to each other only if the Robinson tiles on the edges that are put adjacent respect A_1. Then there exists a subset $T_2' \subset T_2$, $|T_2'| = |T_1| = 56$, with tiling rules $A_2' = A_2|_{T_2'}$, and a bijection $T_2' \to T_1$ under which A_1 and A_2' are equivalent.*

From this result it follows that:

Corollary 5.2. *(Scale Invariance of the Robinson Tiling) Under the bijection in theorem 5.4, the Robinson tiling pattern is preserved under the $2\times2 \to 1\times1$ renormalisation of the grid.*

We can then translate this scale invariance into a statement about the properties of the Hamiltonian which describes the Robinson tiling, i.e.,

Theorem 5.5 ((Informal) Robinson Tiling Hamiltonian Renormalisation)**.** *Let $h_T \in \mathbb{C}^T \otimes \mathbb{C}^T$ be the local interactions which describe the Robinson tiling Hamiltonian. Then there exists a renormalisation group mapping $\mathcal{R}_T$ such that $R_T(h_T) \in \mathbb{C}^T \otimes \mathbb{C}^T$, such that $R_T(h_T)$ preserves both the ground state energy and the tiling pattern.*

More formally we can write:

Definition 5.3 (Tiling Hamiltonian Renormalisation)**.** *Let $h_T^{col}, h_T^{row} \in \mathcal{B}(\mathbb{C}^T \otimes \mathbb{C}^T)$ be the local interactions describing the tiling Hamiltonian. Let $h_{i,i+1}^{row}(j)$ denote the row*

interaction between sites $(i,j),(i+1,j)$ *and similarly let* $h^{col}_{j,j+1}(i)$ *be the interaction between* $(i,j),(i,j+1)$. *Let the* 2×2 *supertiles be assigned at* $(i,j),(i+1,j),(i,j+1),(i+1,j+1)$ *and sites consistent with it. Then the renormalised Hamiltonian has local terms* $R(h^{col}_T),R(h^{col}_T)\in\mathcal{B}(\mathbb{C}^T\otimes\mathbb{C}^T)$,

$$R(h^{col}_T)_{\lceil j/2-1\rceil,\lceil j/2-1\rceil+1} = V_{(i,i+1),(j+2,j+3)}V_{(i,i+1),(j,j+1)}\left(h^{col}_{T,j+1,j+2}(i)+h^{col}_{T,j+1,j+2}(i+1)\right)\Big|_{T'_2} \tag{5.20}$$

$$\times V^{\dagger}_{(i,i+1),(j,j+1)}V^{\dagger}_{(i,i+1),(j+2,j+3)} \tag{5.21}$$

$$R(h^{row}_T)_{\lceil i/2-1\rceil,\lceil i/2-1\rceil+1} = V_{(i+2,i+3),(j,j+1)}V_{(i,i+1),(j,j+1)}\left(h^{row}_{T,i+1,i+2}(j)+h^{row}_{T,i+1,i+2}(j+1)\right)\Big|_{T'_2} \tag{5.22}$$

$$\times V^{\dagger}_{(i,i+1),(j,j+1)}V^{\dagger}_{(i+2,i+3),(j,j+1)} \tag{5.23}$$

where V *is an isometry which can be explicitly constructed. In the above we have used the standard abbreviation that each local term is implicitly tensored with the appropriate identity terms, e.g.* $h^{col}_{T,j+1,j+2}(i)$ *is actually* $\mathbb{1}_{i,j}\otimes\mathbb{1}_{i+1,j}\otimes h^{col}_{T,j+1,j+2}(i)\otimes\mathbb{1}_{i,j+3}\otimes\mathbb{1}_{i+1,j+3}$.

The result is:

Lemma 5.1. *The matrix form of the initial and renormalised Hamiltonian are the same, i.e.,*

$$R(h^{row}_T)_{i,i+1} = h^{row}_{T,i,i+1} \qquad \text{and} \qquad R(h^{col}_T)_{j,j+1} = h^{col}_{T,j,j+1}. \tag{5.24}$$

5.5 Renormalisation of the Quantum Hilbert Space

In this section we will deal with the renormalisation of the quantum Hamiltonian. For this, we will need a number of definitions from [CPGW15a].

Definition 5.4 (Standard Basis States)**.** *Let the single site Hilbert space be* $\mathcal{H}=\otimes_i\mathcal{H}_i$ *and fix some orthonormal basis for the single site Hilbert space. Label the set of single site basis states for site* i *as* $\mathfrak{B}^{(i)}_q$. *Then a* standard basis state *for* $\mathcal{H}^{\otimes L}$ *are product states over the single site basis.*

Definition 5.5 (Penalty Terms and Transition Rules)**.** *The two-local quantum Hamiltonian will contain two types of terms:* penalty terms *and* transition rule *terms. Penalty terms have the form* $|ab\rangle\langle ab|$ *where* $|a\rangle$ *and* $|b\rangle$ *are standard basis states. This adds a positive energy contribution to any configuration containing the state* $|ab\rangle$*, which we call an* illegal pair*. Transition rule terms take the form* $\frac{1}{2}(|ab\rangle - |cd\rangle)(\langle ab| - \langle cd|)$ *with* $|ab\rangle \neq |cd\rangle$*, where* $|ab\rangle$ *and* $|cd\rangle$ *act on the same pair of adjacent sites.*

Definition 5.6 (Legal and Illegal States)**.** *We call a standard basis state* legal *if it does not contain any illegal pairs, and* illegal *otherwise*

We then define a standard form Hamiltonian on the joint system

$$\mathcal{H}_C \otimes \mathcal{H}_Q := (\mathbb{C}^C \otimes \mathbb{C}^Q)^{\otimes L} = (\mathbb{C}^C)^{\otimes L} \otimes (\mathbb{C}^Q)^{\otimes L}. \tag{5.25}$$

Importantly, the Gottesman-Irani Hamiltonian we will be considering will be of *standard form*. Standard form Hamiltonians are a particular type of Hamiltonian. Due to the length of the definition, we leave it to definition A.2 and simply claim that certain properties apply to standard form Hamiltonians.

The 1D Gottesman-Irani Hamiltonian $H_q(L) \in \mathcal{B}(\mathbb{C}^d)^{\otimes L}$, introduced in section 5.2.4, can be written as:

$$H_q = H_{trans} + H_{in} + H_{pen} + H_{halt}, \tag{5.26}$$

where H_{trans} contains transition rule terms, H_{pen} is a set of penalty terms which penalise states that should not appear in correct history states, H_{in} penalises states which are incorrectly initialised, and H_{halt} penalises states which encode a halting computation. Moreover, it has a six-fold tensor product form

$$\mathcal{H}_q = \bigotimes_{j=1}^{6} (\mathcal{H}_q)_j. \tag{5.27}$$

where each $(\mathcal{H}_q)_j$ is identified with a different track.

Lemma 43 of [CPGW15a] identifies three subspaces of states, which are closed under the action of H_q.

1. **Illegal Subspace**, $\mathcal{S}_1$: All $|x\rangle \in \mathcal{S}_1 \subset \mathfrak{B}^{\otimes L}$ are in the support of H_{pen} and hence $\langle x| H |x\rangle \geq 1$. By [CPGW15a] Lemma 43, the minimum eigenvalue of these subspaces is $\lambda_0(H|_{\mathcal{S}_1}) \geq 1$.

2. **Evolve-to-Illegal Subspace**, $\mathcal{S}_2$: All standard basis states $|x\rangle \in \mathcal{S}_2 \subset \mathfrak{B}^{\otimes L}$ will evolve either forwards or backwards in time to an illegal state in $O(L^2)$ steps under the transition rules. As per lemma 5.8 of [Wat19], the minimum eigenvalue of these subspaces is $\lambda_0(H|_{\mathcal{S}_2}) = \Omega(L^{-2})$.

3. **Legal Subspace**, $\mathcal{S}_3$: all standard basis states in $\mathcal{S}_3$ are legal and *do not* evolve to illegal states. By [CPGW15a] lemma 43, they have zero support on H_{pen} or H_{in}.

In our renormalisation procedure we seek to preserve only the low energy subspace, hence at any point where we can locally identify states as being in subspace $\mathcal{S}_1$ or $\mathcal{S}_2$, we will remove them from the state space in the renormalisation step.

However, we note that in the general case we cannot locally identify all such states in $\mathcal{S}_2$. That is, determining the whether a state evolves to an illegal under the action of the transitions may be impossible if we only look at what the state looks like on a $O(1)$-subset of the sites.

5.5.0.1 The Ground States

From [CPGW15a] we know that there are two cases we need to consider: the QTM encoded in $H_q(L)$ halts or does not halt.

Lemma 5.2. *Let a given UTM be encoded in the Gottesman-Irani Hamiltonian $H_q(L)$. Then $H_q(L)$ has a ground state energy that is either 0 if the UTM does not halt within time $T(L)$ or $1 - \cos\left(\frac{\pi}{2T}\right)$ if the UTM does halt within $T(L)$. $T(L)$ is a fixed, predetermined function. In the non-halting case, the ground state is*

$$|\Psi_{hist}(L)\rangle = \frac{1}{\sqrt{T}} \sum_{t=1}^{T(L)} |t\rangle |\psi_t\rangle, \tag{5.28}$$

and in the halting case it is

$$|\Psi_{halt}(L)\rangle = \sum_{t=1}^{T(L)} 2\cos\left(\frac{(2t+1)\pi t}{4T}\right)\sin\left(\frac{\pi}{4T}\right)|t\rangle\,|\psi_t\rangle\,, \tag{5.29}$$

where $|t\rangle$ *is the state of the clock register and* $|\psi_t\rangle = \prod_{j=1}^{t} U_j\,|\psi_0\rangle$ *and* $|\psi_0\rangle$ *is the initial state of the computational register and the* $\{U_t\}$ *represent the action of the QTM at time step* t.

Proof. Combine the standard form property of H_q from [CPGW15a] with Lemma 5.10 of [Wat19]. □

5.5.1 Block Renormalisation of the Gottesman-Irani Hamiltonian

In this section we will construct a renormalisation scheme for the Gottesman-Irani Hamiltonian. For a given spin at site i, we write each possible conventional basis state (i.e. basis state before the RG procedure has started) as $\left|\begin{array}{|c|}\hline a \\ \hline \alpha \\ \hline\end{array}\right\rangle_{(i)} \in \mathbb{C}^C \otimes \mathbb{C}^Q$, where the top cell indicates the classical tracks of the construction encoded in [CPGW15a], while the bottom cell indicates the quantum tracks (see section 5.2.4).

We then define a pair of operations: the blocking operation $\mathcal{B}_q$ and the truncation operation $\mathcal{T}_q$. Given a line of qudits $\mathcal{B}_q$ will essentially combine pairs of lattice sites into single sites with a larger local Hilbert space dimension, while $\mathcal{T}_q$ will remove any of the new single site states which can be locally detected to have zero overlap with the ground state. Thus $\mathcal{T}_q$ reduces the local Hilbert space dimension.

We do not truncate all high energy states since in the halting case this would remove the ground state of the Gottesman-Irani Hamiltonian (note this is a slight difference from the standard definition of the BRG). Instead, we removed states based on a combination of high energy and a priori knowledge of the ground state.

Blocking $\mathcal{B}_q$

The blocking part of the renormalisation procedure is defined as follows.

Definition 5.7 (Gottesman-Irani Blocking, $\mathcal{B}_q$)**.** *Let* $|\psi\rangle \in \mathcal{H}_q^{(i)} \otimes \mathcal{H}_q^{(i+1)}$, $i \in \mathbb{N}$. *The blocking operation,* $\mathcal{B}_q : \mathcal{H}_q^{(i)} \times \mathcal{H}_q^{(i+1)} \to \mathcal{H}_q'^{(i/2)}$, *is given by the action of the unitary*

$U_{i,i+1} : \mathcal{H}_q^{(i)} \times \mathcal{H}_q^{(i+1)} \to R(\mathcal{H}_q)'$ *as*

$$\mathcal{B}_q^{(i,i+1)} : |\psi\rangle \mapsto U_{i,i+1} |\psi\rangle \tag{5.30}$$

where

$$U_{i,i+1} = \sum_{|x\rangle,|y\rangle \in \mathfrak{B}} |xy\rangle_{i/2} \langle x|_i \langle y|_{i+1} . \tag{5.31}$$

We extend this to $|\chi\rangle \in \mathcal{H}_q^{\otimes L}$ *as*

$$\mathcal{B}_q : |\chi\rangle \mapsto U |\chi\rangle , \tag{5.32}$$

where $U = \bigotimes_{i \in 2\mathbb{N}}^{i \leq L/2} U_{i,i+1}$.

This can be expressed more intuitively in terms of basis states

$$\mathcal{B}_q^{(i,i+1)} : \left| \begin{array}{|c|}\hline a \\ \hline \alpha \\ \hline \end{array} \right\rangle_{(i)} \otimes \left| \begin{array}{|c|}\hline b \\ \hline \beta \\ \hline \end{array} \right\rangle_{(i+1)} \longrightarrow \left| \begin{array}{|c|c|}\hline a & b \\ \hline \alpha & \beta \\ \hline \end{array} \right\rangle_{(i/2)} . \tag{5.33}$$

Note that $\mathcal{B}_q$ is just a relabelling of the space, so the local Hilbert space dimension is now $\mathbb{C}^{d^2}$ and part of the tensor product structure is lost. We denote by $\mathcal{H}'_q$ this new local Hilbert space spanned by the basis $\mathfrak{B}'^{(1)}$.

From here forward if $|a\rangle, |b\rangle$ are the initial states, then the image under the blocking will be denoted $|ab\rangle$, and if the map is repeated $|ab\rangle . |cd\rangle$ is mapped to $|abcd\rangle$, etc.

Truncation $\mathcal{T}_q$

The truncation part of the RG map truncates the local Hilbert space to discard those states which locally have support on the penalty terms.

Definition 5.8 (Gottesman-Irani Truncation Mapping, $\mathcal{T}_q$)**.** *Let* $\mathfrak{B}^{(1)}$ *be the set of basis states defined by* $\mathcal{B}_q$ *such states with a preimage* $|a\rangle |b\rangle$*, such that* $|a\rangle, |b\rangle \in \mathfrak{B}$

cannot be locally identified as being in subspace $\mathcal{S}_1$ or $\mathcal{S}_2$. That is

$$\langle a|\langle b|h_{pen}^{i,i+1}|a\rangle|b\rangle = \langle a|\langle b|h_{in}^{i,i+1}|a\rangle|b\rangle = 0, \tag{5.34}$$

$$\langle a|\langle b|h_{trans}^{(i,i+1)}h_{pen}^{(i,i+1)}h_{trans}^{(i,i+1)}|a\rangle|b\rangle = 0. \tag{5.35}$$

The truncation mapping is then $\mathcal{T}_q^{(i,i+1)} : R(\mathcal{H}_q)' \to R(\mathcal{H}_q)$ for $R(\mathcal{H}_q) = \text{span}\{\mathfrak{B}_q^{(1)}\} \subset R(\mathcal{H})_q'$. Then the full restriction is $\mathcal{T}_q : \mathcal{H}_q'^{\otimes L/2} \to R(\mathcal{H}_q)^{\otimes L/2}$.

We now combine the unitary and subspace restriction to give an isometry which implements $\mathcal{T}_q \circ \mathcal{B}_q$.

Lemma 5.3 (Renormalisation Unitary Structure)**.** *Let the renormalisation isometry $V_{i,i+1}^{GI}$ be the unitary map follow by subspace restriction previously described. Define $V^{GI} : \mathcal{H}_q^{\otimes L} \to R(\mathcal{H}_q)^{\otimes L/2}$ to implement the mapping $\mathcal{T}_q \circ \mathcal{B}_q$ on a state in $\mathcal{H}_q^{\otimes L}$, as*

$$\mathcal{T}_q \circ \mathcal{B}_q : |\chi\rangle \mapsto U|\chi\rangle|_{R(\mathcal{H}_q)^{\otimes L/2}} =: V^{GI}|\chi\rangle. \tag{5.36}$$

where U is defined in definition 5.7 and $R(\mathcal{H}_q)$ is defined in definition 5.8. Then V^{GI} can be defined as and decomposed as

$$V^{GI} := \bigotimes_{i\in 2\mathbb{N}}^{i\leq\lfloor L/2\rfloor} V_{i,i+1}^{GI} = \bigotimes_{i\in 2\mathbb{N}}^{i\leq\lfloor L/2\rfloor}\left(\bigotimes_{j=1}^{6} V_{i,i+1}^{GI\,(j)}\right), \tag{5.37}$$

with

$$V_{i,i+1}^{GI} : \mathcal{H}_q^{\otimes 2} \to R(\mathcal{H}_q) \tag{5.38}$$

and where each part of the decomposition acts on one of the six different tracks,

$$V_{i,i+1}^{GI\,(j)} : \mathcal{H}_{q,j}^{\otimes 2} \to R(\mathcal{H}_q)_j. \tag{5.39}$$

Proof. The decomposition $V^{GI} = \bigotimes_{i\in 2\mathbb{N}}^{i\leq\lfloor L/2\rfloor} V_{i,i+1}^{GI}$ is evident from the block procedure. The decomposition $V_{i,i+1}^{GI} = \bigotimes_{j=1}^{6} V_{i,i+1}^{GI\,(j)}$ arises from the fact that the procedure keeps each basis state as a product across the different tracks and hence the different $\mathcal{H}_{q,j}$. □

We now need to define how the Hamiltonian acts with respect to the RG procedure. We want to break down the Hamiltonian into different subspaces and renormalise them separately while preserving the ground state (in both the halting and non-halting cases) and its energy.

Lemma 5.4 (Renormalised Gottesman-Irani Hamiltonian)**.** *Let h_q be the local terms of a nearest neighbour, translationally invariant Hamiltonian*

$$H_q(L) = \sum_{i=1}^{L} h_q^{(i,i+1)} = H_{trans} + H_{pen} + H_{in} + H_{out}, \tag{5.40}$$

such that $H(L)$ is standard form. Let $V : \mathbb{C}^d \otimes \mathbb{C}^d \to \mathbb{C}^{f(d)}$, be the isometry from lemma 5.3. Then the renormalised Hamiltonian, defined as

$$\mathcal{R}(H_q(L)) = V^{GI} H_q(L) V^{GI\dagger} = \sum_{i=1}^{L-1} V^{GI} h_q^{(i,i+1)} V^{GI\dagger} = R(H_q)(L), \tag{5.41}$$

is a translationally invariant, nearest-neighbour Hamiltonian with local interactions $R(h_q)^{(i/2,i/2+1)} = V^{GI}(h_q^{(i-1,i)} + h_q^{(i+1,i+2)})V^{GI\dagger}$ and $R(h_q)^{i/2} = V^{GI} h_q^{(i,i+1)} V^{GI\dagger}$. Furthermore, $R(H_q)(L)$ has the following properties:

1. *$R(H_q)(L)$ is a standard form Hamiltonian.*
2. *$R(H_{trans})$ encodes a transition $V^{GI}(|ab\rangle |\psi_{abcd}\rangle) \to V^{GI}(|cd\rangle U_{abcd} |\psi_{abcd}\rangle)$ iff H_{trans} encodes the transition $|ab\rangle |\psi_{abcd}\rangle \to |cd\rangle U_{abcd} |\psi_{abcd}\rangle$.*
3. *$R(H_{pen}), R(H_{in}), R(H_{out})$ have support on a renormalised basis state $V^{GI}(|ab\rangle |\psi\rangle)$ iff H_{pen}, H_{in}, H_{out} respectively have non-zero support on $|ab\rangle |\psi\rangle$.*
4. *$\lambda_0(H_q(L)) = \lambda_0(R(H_q)(L/2))$ (the ground state energy is preserved).*
5. *$R(H_q)$ maintains the six-fold tensor product structure of the original Hamiltonian H_q in eq. (5.27), that is, $R(\mathcal{H}_q) = \bigotimes_{j=1}^{6} R(\mathcal{H}_q)_j$.*

Proof. First note that for all $i \in 2\mathbb{N}$, $V_{i,i+1}^{GI} h^{(i,i+1)} V_{i,i+1}^{GI\dagger} \in \mathcal{B}(\mathbb{C}^{f(d)})$ is now a 1-local term in the new renormalised Hamiltonian, for some $f : \mathbb{N} \to \mathbb{N}$. However,

$V^{GI}_{i+2,i+3}V^{GI}_{i,i+1}h^{(i+1,i+2)}V^{GI\dagger}_{i,i+1}V^{GI\dagger}_{i+2,i+3} \in \mathcal{B}(\mathbb{C}^{f(d)} \otimes \mathbb{C}^{f(d)})$

Claims 1 and 2: From the linearity of V^{GI}, we see that $\mathcal{R}(H_q(L)) = R(H_{trans}) + R(H_{pen}) + R(H_{in}) + R(H_{out})$. It is trivial to see that $R(H_{trans}) = V^{GI}H_{trans}V^{GI\dagger} = \sum_{ab\to cd}(V^{GI}|cd\rangle \otimes U_{abcd} - V^{GI}|ab\rangle)(\langle cd| \otimes U^{\dagger}_{abcd}V^{GI\dagger} - \langle ab|V^{GI\dagger})$, and hence encodes transitions between the renormalised states. This also shows $R(H_{trans})$ satisfies Claim 2. Due to the decompositional properties of V^{GI}, as shown in lemma 5.3, we preserve that H_{trans} acts diagonally on the states in $\mathbb{C}^C$. Likewise, it preserves the form of H_{pen}, H_{in}, H_{out} as projectors onto a subset of states.

Claim 3: Consider the penalty terms: given a renormalised state $V^{GI}|\psi\rangle$, it is clear that

$$(\langle\psi|V^{GI\dagger})V^{GI}H_{pen}V^{GI\dagger}(V^{GI}|\psi\rangle) = \langle\psi|H_{pen}|\psi\rangle = 1,$$

hence $V^{GI}|\psi\rangle$ is penalised by the renormalised Hamiltonian iff $|\psi\rangle$ is penalised by the unrenormalised Hamiltonian. The same applied to H_{in} and H_{out}.

Claim 4: First note that any state

$$|\Psi\{a_t\}\rangle = \sum_{t=1}^{\tau} a_t(|t\rangle|\psi_t\rangle). \tag{5.42}$$

which encodes a valid evolution is in the kernel of H_{in}, H_{pen}, and is contained in subspace $\mathcal{S}_3$. Thus, $V^{GI}|\Psi\{a_t\}\rangle \in R(\mathcal{H})^{\otimes L/2}$, and after the RG procedure $\mathcal{T}_q \circ \mathcal{B}_q$ the corresponding renormalised state is

$$|\Psi'\{a_t\}\rangle = \sum_{t=1}^{\tau} a_t V^{GI}(|t\rangle|\psi_t\rangle). \tag{5.43}$$

To see the energy of such states is preserved note

$$\langle\Psi'\{a_t\}|V^{GI}H_q(L)V^{GI\dagger}|\Psi'\{a_t\}\rangle = \langle\Psi\{a_t\}|H_q(L)|\Psi\{a_t\}\rangle. \tag{5.44}$$

From lemma 5.2 the ground states are of the form $|\Psi\{a_t\}\rangle$. We know that the state

$V^{GI}|\Psi\{a_t\}\rangle$ has the same energy. Since the minimum eigenvalue is given by

$$\begin{aligned}
\lambda_0(H_q(L)) &= \min_{x\in\mathcal{H}_q^{\otimes L}} \frac{\langle x|H_q(L)|x\rangle}{\langle x|x\rangle} \\
&= \min_{x\in\mathcal{H}_q^{\otimes L}} \frac{\langle x|UU^\dagger H_q(L)U^\dagger U|x\rangle}{\langle x|U^\dagger U|x\rangle} \qquad (5.45)\\
&\leq \min_{\substack{x\in\mathcal{H}_q^{\otimes L}\\ V^{GI}|x\rangle\neq 0}} \frac{\langle x|V^{GI}V^{GI\dagger}H_q(L)V^{GI\dagger}V^{GI}|x\rangle}{\langle x|V^{GI\dagger}V^{GI}|x\rangle} \qquad (5.46)\\
&= \lambda_0(R(H_q)(L/2)),
\end{aligned}$$

where going from eq. (5.45) to eq. (5.46) we have used the fact that we have restricted the subspace to remove the states that are integrated out by V^{GI}. Since $\lambda_0(R(H_q)(L/2)) = \lambda_0(H_q(L/2))$, then we can confirm $V^{GI}|\psi_{halt}\rangle$ and $V^{GI}|\psi_{hist}\rangle$ are the appropriate ground states after the renormalisation procedure.

Claim 5: The preservation of the structure in eq. (5.27) follows directly from the tensor product form of the isometry given in eq. (5.37) applied according to the renormalisation method described by eq. (5.41).

□

Consecutive steps of the RG procedure can be derived straightforwardly. The Hilbert space obtained after k-th RG steps of can be constructed by induction

$$(\mathcal{T}_q \circ \mathcal{B}_q)^{\circ(k)} = \mathcal{T}_q \circ \mathcal{B}_q \circ (\mathcal{T}_q \circ \mathcal{B}_q)^{\circ(k-1)}$$

We can thus concatenate multiple renormalisations of the Gottesman-Irani Hamiltonian in one isometry, $V^{GI}[k] : R^{(k-1)}(\mathcal{H}_q)^{\otimes 2L} \to R^{(k)}(\mathcal{H}_q)^{\otimes L}$, given by

$$V^{GI}[k] = \Pi_{j=1}^{k} V^{GI}_{L/2^j}$$

where $V^{GI}_{L/2^j}$ is the isometry outlined in lemma 5.3, but now acting on the appropriate local Hilbert space, and the subscript $L/2^j$ indicates that the operator is acting on a 1D chain of $L/2^j$ sites.

Accordingly, the renormalised Hamiltonian is then

$$R^{(k)}(H_q(L)) = V^{GI}[k]H_q(L)V^{GI\dagger}[k].$$

It follows immediately from lemma 5.4 is that this RG mapping takes standard form Hamiltonians to standard form Hamiltonians while preserving the energy of the ground state.

5.6 Putting it all Together

In this section we combine the renormalisation group schemes for the separate parts of the Hamiltonian. First recall Lemma 51 of [CPGW15a] which characterises the ground state of the Hamiltonian defined by the local terms h_u:

Lemma 5.5 (Tiling + quantum layers, Lemma 51 of [CPGW15a])**.** *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 section 5.5. Then there is a Hamiltonian on a 2D square lattice with nearest-neighbour interactions $h_u^{\text{row}}, h_u^{\text{col}} \in \mathcal{B}(\mathbb{C}^{C+Q+1} \otimes \mathbb{C}^{C+Q+1})$ with the following properties: For any region of the lattice, the restriction of the Hamiltonian to that region has an eigenbasis of the form $|T\rangle_c \otimes |\psi\rangle_q$, where $|T\rangle_c$ is a* product *state representing a classical configuration of tiles. Furthermore, for any given $|T\rangle_c$, the lowest energy choice for $|\psi\rangle_q$ consists of ground states of $H_q(r)$ on segments between sites in which $|T\rangle_q$ 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 $|e\rangle$'s everywhere else.*

The $|L\rangle$ and $|R\rangle$ tiles are identified in [CPGW15a] with the right-down and left-down red cross in the Robinson tiles respectively (see section 5.4). The ground state can then be shown to be the ground state of the Robinson tiling Hamiltonian plus a "quantum layer" in which the Gottesman-Irani ground states appear only over the tops of the Robinson squares. Everywhere else in the quantum layer is a filler state $|e\rangle$.

A key point is that the eigenstates are all product states across $\mathcal{H}_c$ and $\mathcal{H}_{eq}$. We wish for the RG mapping to preserve this property. This restricts the type of renormalisation isometries we use, as detailed in the following lemma.

Lemma 5.6 (Separable Eigenstates)**.** *Let $H_u^{\Lambda(2L)}$ denote the Hamiltonian in lemma 5.5. Then for an isometry $Z = Z_c \otimes Z_{eq}$ where $Z_c : \mathcal{H}_c^{\otimes 2\times 2} \to R(\mathcal{H}_c)$ and $Z_{eq} : \mathcal{H}_{eq}^{\otimes 2\times 2} \to R(\mathcal{H}_{eq})$, the operator $ZH_u^{\Lambda(2L)}Z^\dagger$ also has eigenstates of the form $|T'\rangle_c \otimes |\psi\rangle_{eq}$ for $|T'\rangle_c \in R(\mathcal{H}_c)^{\otimes\Lambda(L)}$ and $|\psi\rangle_{eq} \in R(\mathcal{H}_{eq})^{\otimes\Lambda(L)}$.*

Proof. As per lemma 5.5, the eigenstates of $H_u^{\Lambda(2L)}$ decompose as product states $|T_c\rangle \otimes |\psi_i\rangle_{eq}$, hence we can write

$$H_u^{\Lambda(2L)} = \sum_i \lambda_i |T_i\rangle\langle T_i| \otimes |\psi_i\rangle\langle\psi_i|_{eq} . \tag{5.47}$$

Applying the renormalisation isometry Z gives

$$ZH_u^{\Lambda(2L)}Z^\dagger = \sum_i \lambda_i Z_C |T_i\rangle\langle T_i|_c Z_C^\dagger \otimes Z_{eq} |\psi_i\rangle\langle\psi_i|_{eq} Z_{eq}^\dagger \tag{5.48}$$

$$=: \sum_i \lambda_i \left|T_i'\right\rangle\left\langle T_i'\right|_{c'} \otimes \left|\psi_i'\right\rangle\left\langle\psi_i'\right|_{eq'} . \tag{5.49}$$

Thus the product structure across the two subspaces is preserved. □

Here we show that renormalising the full Hamiltonian preserves this Robinson tiling plus Gottesman-Irani ground state structure.

5.6.1 Renormalising $\mathcal{H}_T \otimes (\mathcal{H}_e \oplus \mathcal{H}_q)$

From lemma 5.6, we know the eigenstates of the Hamiltonian defined by h_u are product states across the classical-quantum Hilbert space partition and this structure is preserved under a tensor product of isometries on the two subspace separately. Thus we can consider the basis states of $\mathcal{H}_T$ and $\mathcal{H}_{eq}$ separately and then later show this preserves the desired properties.

Blocking Operation $\mathcal{B}_u$ We know that V^C from lemma 5.1 will renormalise the classical state space by mapping sets of 2×2 tiles to new tiles which recreate the

tiling pattern at all but the lowest level. We use this isometry unchanged, acting on the classical part of the Hilbert space.

Consider the quantum Hilbert space $\mathcal{H}_{eq}$. First note that the Gottesman-Irani Hamiltonian to be renormalised is a standard form Hamiltonian, and so can be renormalised as per section 5.5.1. However, the blocking procedure from section 5.5.1 is not sufficient for our purposes as it (a) takes a set of 2×1 lattice sites to a single lattice site and so is not appropriate for a 2D lattice, and (b) does not include the filler state $|e\rangle_e$. To remedy this we need an isometry which acts as:

$$V^{eq}_{(i,i+1)(j,j+1)} : \mathcal{H}^{(i,j)}_{eq}\otimes\mathcal{H}^{(i+1,j)}_{eq}\otimes\mathcal{H}^{(i,j+1)}_{eq}\otimes\mathcal{H}^{(i+1,j+1)}_{eq} \to (\mathcal{H}'_{eq}\otimes\mathcal{H}'_{eq})^{(i/2,j/2)}. \quad (5.50)$$

We will find it useful to define the following notation:

Definition 5.9 (k-times Blocked Basis States)**.** *Let* $|x_1\rangle, |x_2\rangle, \ldots, |x_{2^k}\rangle \in \mathfrak{B}\cup|e\rangle_e$, *then we denote the corresponding renormalised basis state after k applications of the RG mapping as* $|x_1x_2\ldots x_{2^k}\rangle$.

Now define $V^q_{(i,i+1)}(j)$ as follows, where $V^{GI}_{i,i+1}$ is the isometry used in lemma 5.4:

$$V^q_{(i,i+1)}(j) = V^{GI}_{i,i+1} + |ee\rangle_{i/2,j/2}\langle e|_{i,j}\langle e|_{i+1,j} \quad (5.51)$$

$$+ |xe\rangle_{i/2,j/2}\langle x|_{i,j}\langle e|_{i+1,j} + |ex\rangle_{i/2,j/2}\langle e|_{i,j}\langle x|_{i+1,j}\,. \quad (5.52)$$

This defines a new set of quantum basis states which now reflect the fact $|e\rangle_e$ is part of the Hilbert space. Denote this

$$\mathfrak{C}^{(1)} := \mathfrak{B}^{(1)}\cup|ee\rangle\bigcup_{x\in\mathfrak{B}}|ex\rangle\bigcup_{x\in\mathfrak{B}}|xe\rangle\,. \quad (5.53)$$

V^q only maps 2×1 spins to a single spin. We need an operator which maps a 2×2 block to a single spin. Define $W : \mathcal{H}'^{(i/2,j)}_{eq}\otimes\mathcal{H}'^{(i/2,j+1)}_{eq} \to (\mathcal{H}'_{eq}\otimes\mathcal{H}'_{eq})^{(i/2,j/2)}$, as simply

$$W_{(i,i+1)(j,j+1)} = \sum(|x\rangle_{q_1}\otimes|y\rangle_{q_2})_{i/2,j/2}\langle x|_{i/2,j}\otimes\langle y|_{i/2,j+1}\,. \quad (5.54)$$

This unitary acts to map the 1×2 set of sites to a single lattice site in the renormalised lattice.

The isometry:

$$V^{eq}_{(i,i+1)(j,j+1)} := W_{(i,i+1)(j,j+1)}\left(V^{q}_{(i,i+1)}(j)\otimes V^{q}_{(i,i+1)}(j+1)\right), \tag{5.55}$$

then maps 2×2 spins to a single spin.

The overall blocking map $\mathcal{B}_u$ is then given by:

Definition 5.10 (Blocking Isometry, V^b, $\mathcal{B}_u$)**.** *Let V^C and V^{eq} be the isometries from definition 5.3 and eq.* (5.55) *respectively. Then the blocking isometry for H_u is given by*

$$V^{b}_{(i,i+1)(j,j+1)} = V^{C}_{(i,i+1)(j,j+1)}\otimes V^{eq}_{(i,i+1)(j,j+1)}. \tag{5.56}$$

We now need to consider the full renormalisation process: the isometry defined above will map a certain subset of states to states on the renormalised lattice. However, some parts of the Hilbert space will be "integrated out". For convenience we will sometimes use indices I, J to indicate row and column indices on the new lattice after the RG transformation.

Let $h_q^{(i,i+1)}(j), h_q^{(i,i+1)}(j+1)$ be the local terms of the quantum Hamiltonian before renormalisation, then we see that

$$\begin{aligned} &V^{eq}_{(i,i+1)(j,j+1)}\left(h_q^{(i,i+1)}(j+1)+h_q^{(i,i+1)}(j)\right)V^{eq\dagger}_{(i,i+1)(j,j+1)} \\ &= h_q^{(1)\prime(I,J)}\otimes\mathbb{1}_{q_2}+\mathbb{1}_{q_1}\otimes h_q^{(1)\prime(I,J)} \end{aligned} \tag{5.57}$$

and

$$\begin{aligned} &V^{eq}_{(i+2,i+3)(j,j+1)}V^{eq}_{(i,i+1)(j,j+1)}\left(h_q^{(i+1,i+2)}(j)+h_q^{(i+1,i+2)}(j)\right)V^{eq\dagger}_{(i+2,i+3)(j,j+1)} \\ &\times V^{eq\dagger}_{(i,i+1)(j,j+1)} = h'^{(I,I+1)}_{q_1}\otimes\mathbb{1}^{(I,J)}_{q_2}\otimes\mathbb{1}^{(I+1,J)}_{q_2}+\mathbb{1}^{(I,J)}_{q_1}\otimes\mathbb{1}^{(I+1,J)}_{q_1}\otimes h'^{(I,I+1)}_{q_2}. \end{aligned} \tag{5.58}$$

Truncation Operation $\mathcal{T}_u$ The operator W has essentially merged two sites into a single site. We now wish to integrate out one of these sites and restrict to the set of "allowed states" in the other. We will implement this using the 1-local projector $\Pi_{gs}(k)$

Definition 5.11 (Truncation Operation $\mathcal{T}_u$)**.** *Let $|\psi\rangle \in \mathcal{H}_c \otimes \mathcal{H}_{eq}$, then*

$$\mathcal{T}_u : |\psi\rangle \mapsto (\mathbb{1}_c \otimes \mathbb{1}_{q_1} \otimes \Pi_{gs}(k)) |\psi\rangle, \tag{5.59}$$

where

$$\Pi_{gs}(k) = \begin{cases} \left|e^{\times 2^k}\right\rangle\left\langle e^{\times 2^k}\right| & k \text{ even} \\ \left|\psi_{hist}(4^n+1)e^{\times 2^k - 4^n - 1}\right\rangle\left\langle \psi_{hist}(4^n+1)e^{\times 2^k - 4^n - 1}\right| & \text{if } k \text{ odd, } 2^{k-1} < 4^n+1 < 2^k, \\ & \text{and non-halting} \\ \left|\psi_{halt}(4^n+1)e^{\times 2^k - 4^n - 1}\right\rangle\left\langle \psi_{halt}(4^n+1)e^{\times 2^k - 4^n - 1}\right| & \text{if } k \text{ odd, } 2^{k-1} < 4^n+1 < 2^k, \\ & \text{and halting}, \end{cases} \tag{5.60}$$

and where $|\psi_{hist}(L)\rangle$ and $|\psi_{halt}(L)\rangle$ are defined in lemma 5.2. This extends to states $|\chi\rangle \in (\mathcal{H}_c \otimes \mathcal{H}_{eq})^{\otimes \Lambda(L)}$, as

$$\mathcal{T}_u : |\chi\rangle \mapsto \bigotimes_{(I,J)\in\Lambda(L)} (\mathbb{1}_c^{(I,J)} \otimes \mathbb{1}_{q_1}^{(I,J)} \otimes \Pi_{gs}^{(I,J)}(k)) |\chi\rangle. \tag{5.61}$$

Definition 5.12 (Renormalisation Isometry, V^u)**.** *Let $V^b_{(i,i+1)(j,j+1)}$ and Π_{gs} be as defined in definition 5.10 and eq. (5.60) respectively. We define the isometry implementing the entire renormalisation scheme as*

$$V^u_{(i,i+1)(j,j+1)} := (\mathbb{1}_c \otimes \Pi_{gs}) V^b_{(i,i+1)(j,j+1)}. \tag{5.62}$$

To see why this is appropriate note that the Hamiltonian after the application of the blocking isometries has two sets of local terms: a 1-local term and a 2-local term (see definition 5.13 and the discussion following). First consider the 1-local term

$h_q^{(1)\prime(I,J)} \otimes \mathbb{1}_{q_2} + \mathbb{1}_{q_1} \otimes h_q^{(1)\prime(I,J)}$ and examine how it transforms under $\mathcal{T}_u$ and Π_{gs}. The idea is that Π_{gs} will "integrate out" the q_2 subspace by removing all states which are not the ground state while maintaining the energy contribution from this subspace. If the site is large enough to contain a full history state of length $4^n + 1$, for some $n \in \mathbb{N}$, then we keep only that state and the relevant renormalised $|e\rangle$ states. Otherwise we keep only the renormalised $|e\rangle$ states. Hence

$$\Pi_{gs}^{(I,J)}(k)(h_{q_1}^{(1)\prime(I,J)} \otimes \mathbb{1}_{q_2}^{(I,J)} + \mathbb{1}_{q_1}^{(I,J)} \otimes h_{q_2}^{(1)\prime(I,J)})\Pi_{gs}^{(I,J)}(k) \tag{5.63}$$

$$= h_{q_1}^{(1)\prime(I,J)} \otimes \Pi_{gs}^{(I,J)}(k) + \mathrm{tr}\left(\Pi_{gs}^{(I,J)}(k) h_{q_2}^{\prime(I,J)}\right) \mathbb{1}_{q_1}^{(I,J)} \otimes \Pi_{gs}^{(I,J)}(k). \tag{5.64}$$

Since Π_{gs} is a projector onto a 1-dimensional subspace, we will often omit it when writing the Hamiltonian. Thus obtain the term

$$h_q^{(1)\prime(I,J)} + \mathrm{Tr}\left(\Pi_{gs}(k) h_{q_2}^{\prime(I,J)}\right) \mathbb{1}_q. \tag{5.65}$$

Now examine how the 2-local terms transform:

$$\Pi_{gs}(k)^{(I,J)} \otimes \Pi_{gs}(k)^{(I+1,J)} \left(h_q^{\prime(I,I+1)} \otimes \mathbb{1}_{q_2}^{(I,J)} \otimes \mathbb{1}_{q_2}^{(I+1,J)}\right. \tag{5.66}$$

$$\left. + \mathbb{1}_{q_1}^{(I,J)} \otimes \mathbb{1}_{q_1}^{(I+1,J)} \otimes h_q^{\prime(I,I+1)}\right) \Pi_{gs}(k)^{(I,J)} \otimes \Pi_{gs}(k)^{(I+1,J)} \tag{5.67}$$

$$= h_q^{\prime(I,I+1)} \otimes \Pi_{gs}(k)^{(I)} \otimes \Pi_{gs}(k)^{(I+1)} \tag{5.68}$$

$$+ \mathrm{tr}\left(h_q^{\prime(I,I+1)} \Pi_{gs}(k)^{(I)} \otimes \Pi_{gs}(k)^{(I+1)}\right) \mathbb{1}_{q_1}^{(I,J)} \otimes \Pi_{gs}(k)^{(I)} \otimes \Pi_{gs}(k)^{(I+1)}. \tag{5.69}$$

Importantly $\mathrm{tr}\left(h_q^{\prime(I,I+1)} \Pi_{gs}(k)^{(I)} \otimes \Pi_{gs}(k)^{(I+1)}\right)$ only picks up a non-zero contribution from the terms proportional to $\mathbb{1}^{(I)} \otimes \mathbb{1}^{(I+1)}$ (we also note that this latter term is zero for interactions going along columns). Again the subspace spanned by $\Pi_{gs}(k)^{(I)} \otimes \Pi_{gs}(k)^{(I+1)}$ is a 1-dimensional subspace and hence we will often omit writing it explicitly. Thus the 2-local terms effectively become $h_q^{\prime(I,I+1)} + \mathrm{tr}\left(h_q^{\prime(I,I+1)} \Pi_{gs}(k)^{(I)} \otimes \Pi_{gs}(k)^{(I+1)}\right) \mathbb{1}_q^{(I,J)} \otimes \mathbb{1}_q^{(I+1,J)}$. This can be generalised straightforwardly to further iterations.

We formalise the overall RG mapping in the following definition:

Definition 5.13 (h_u Renormalisation Mapping)**.** *Let* $h_u^{col(i,i+1)}, h_u^{row(j,j+1)} \in \mathcal{B}(\mathbb{C}^d \otimes$

$\mathbb{C}^d$) *and* $V^u_{(i,i+1)(j,j+1)}$ *be as in definition 5.12. Then the renormalised local terms are given by*

$$\begin{aligned}
&\mathcal{R} : h_u^{row(i+1,i+2)}(j) + h_u^{row(i+1,i+2)}(j+1) \to V^u_{(i+2,i+3)(j,j+1)} V^u_{(i,i+1)(j,j+1)} \times \\
&\quad \left(h_u^{row(i+1,i+2)}(j) + h_u^{row(i+1,i+2)}(j+1) \right) V^{u\dagger}_{(i,i+1)(j,j+1)} V^{u\dagger}_{(i+2,i+3)(j,j+1)} \\
&\quad =: R(h_u^{row})^{(i,i+1)} \\
&\mathcal{R} : h_u^{col(j+1,j+2)}(i) + h_u^{col(j+1,j+2)}(i+1) \to V^u_{(i+2,i+3)(j,j+1)} V^u_{(i,i+1)(j,j+1)} \times \\
&\quad \left(h_u^{col(j+1,j+2)}(i) \ + h_u^{col(j+1,j+2)}(i+1) \right) V^{u\dagger}_{(i,i+1)(j,j+1)} V^{u\dagger}_{(i+2,i+3)(j,j+1)}, \\
&\quad =: R(h_u^{col})^{(i,i+1)}
\end{aligned}$$

$$\begin{aligned}
&\mathcal{R} : h_u^{row(i,i+1)}(j) + h_u^{row(i+1,i+2)}(j+1) + \sum_{\substack{k=0,1 \\ \ell=1,2}} \left(h_u^{(1)(i+k,j+\ell)} \right) \to \\
&V^u_{(i,i+1)(j,j+1)} \left(h_u^{row(i,i+1)}(j) + h_u^{row(i+1,i+2)}(j+1) + \sum_{\substack{k=0,1 \\ \ell=1,2}} \left(h_u^{(1)(i+k,j+\ell)} \right) \right) V^{u\dagger}_{(i,i+1)(j,j+1)} \\
&\qquad =: R(h_u^{(1)})^{(i)}.
\end{aligned}$$

$R^{(k)}(h_u^{row}), R^{(k)}(h_u^{col})^{(i,i+1)}, R^{(k)}(h_u^{(1)})^{(i)}$ *are defined in the same way but with the appropriate isometries for the* k^{th} *iteration of the RG mapping.*

Remark 5.1. $R^{(k)}(h_u^{(1)})^{(i)}$ *and* $R^{(k)}(h_u^{row})^{(i,i+1)}$ *have local projector terms of the form* $\sum_{m=1}^{k} 4^m \kappa^{(m)} \mathbb{1}^{(i)}$ *and* $\sum_{m=1}^{k} 2^m \gamma^{(m)} \mathbb{1}^{(i)} \otimes \mathbb{1}^{(i+1)}$*, where* $\gamma^{(k)}$ *and* $\kappa^{(k)}$ *are given by*

$$\kappa^{(k)} := \mathrm{Tr}\left(\Pi_{gs}(k) h_{q_2}^{\prime(I,J)} \right) \tag{5.70}$$

$$\gamma^{(k)} := \mathrm{tr}\left(h_q^{\prime(I,I+1)} \Pi_{gs}(k)^{(I)} \otimes \Pi_{gs}(k)^{(I+1)} \right). \tag{5.71}$$

We now examine the properties of the full Hamiltonian under this mapping, and show that its ground state energy and other properties are preserved.

Lemma 5.7 (H_u Renormalisation)**.** *Let* $H_u(L) = \sum h_u^{row(j,j+1)} + \sum h_u^{col(i,i+1)}$*, where*

$$
\begin{aligned}
h_{j,j+1}^{\text{col}} =& h_c^{\text{col}} \otimes \mathbb{1}_{eq}^{(j)} \otimes \mathbb{1}_{eq}^{(j+1)} && \text{(5.72a)}\\
h_{i,i+1}^{\text{row}} =& h_c^{\text{row}} \otimes \mathbb{1}_{eq}^{(i)} \otimes \mathbb{1}_{eq}^{(i+1)} && \text{(5.72b)}\\
&+ \mathbb{1}_c^{(i)} \otimes \mathbb{1}_c^{(i+1)} \otimes h_q && \text{(5.72c)}\\
&+ |L\rangle\langle L|_c^{(i)} \otimes (\mathbb{1}_{eq} - \left|\circledless\right\rangle\left\langle\circledless\right|)^{(i)} \otimes \mathbb{1}_{ceq}^{(i+1)} && \text{(5.72d)}\\
&+ (\mathbb{1}_c - |L\rangle\langle L|_c)^{(i)} \otimes \left|\circledless\right\rangle\left\langle\circledless\right|^{(i)} \otimes \mathbb{1}_{ceq}^{(i+1)} && \text{(5.72e)}\\
&+ \mathbb{1}_{ceq}^{(i)} \otimes |R\rangle\langle R|_c^{(i+1)} \otimes (\mathbb{1}_{eq} - \left|\circledgtr\right\rangle\left\langle\circledgtr\right|)^{(i+1)} && \text{(5.72f)}\\
&+ \mathbb{1}_{ceq}^{(i)} \otimes (\mathbb{1}_c - |R\rangle\langle R|)_c^{(i+1)} \otimes \left|\circledgtr\right\rangle\left\langle\circledgtr\right|^{(i+1)} && \text{(5.72g)}\\
&+ \mathbb{1}_c^{(i)} \otimes |0\rangle\langle 0|_e^{(i)} \otimes |R\rangle\langle R|_c^{(i+1)} \otimes \mathbb{1}_{eq}^{(i+1)} && \text{(5.72h)}\\
&+ |L\rangle\langle L|_c^{(i)} \otimes \mathbb{1}_{eq}^{(i)} \otimes \mathbb{1}_c^{(i+1)} \otimes |0\rangle\langle 0|_e^{(i+1)} && \text{(5.72i)}\\
&+ \mathbb{1}_c^{(i)} \otimes |0\rangle\langle 0|_e^{(i)} \otimes (\mathbb{1}_c - |L\rangle\langle L|)_c^{(i+1)} \otimes (\mathbb{1}_{eq} - |0\rangle\langle 0|)_e^{(i+1)} && \text{(5.72j)}\\
&+ (\mathbb{1}_c - |R\rangle\langle R|)_c^{(i)} \otimes (\mathbb{1}_{eq} - |0\rangle\langle 0|)_e^{(i)} \otimes \mathbb{1}_c^{(i+1)} \otimes |0\rangle\langle 0|_e^{(i+1)}, && \text{(5.72k)}\\
&+ \mathbb{1}_{ceq}^{(i)} \otimes \mathbb{1}_{ceq}^{(i+1)} && \text{(5.72l)}\\
h_i^{(1)} =& -(1 + \alpha_2(\varphi))\mathbb{1}_{ceq}^{(i)}, && \text{(5.72m)}
\end{aligned}
$$

where

$$
\alpha_2(\varphi) := \sum_{4^n+7 > |\varphi|} 4^{-2n-1} \lambda_0(H_q(4^n)), \tag{5.73}
$$

as defined in Proposition 53 of [CPGW15a]. Then the k times renormalised Hamiltonian $R^{(k)}(H_u)^{\Lambda(L\times W)}$ *has the following properties:*

1. *For any finite region of the lattice, the restriction of the Hamiltonian to that region has an eigenbasis of the form* $|T\rangle_c \otimes |\psi_i\rangle$ *where* $|T\rangle_c$ *is a classical tiling state (cf. Lemma 51 of [CPGW15a]).*

2. *Furthermore, for any given* $|T\rangle_c$*, the lowest energy choice for* $|\psi\rangle_q$ *consists of ground states of* $R^{(k)}(H_q)(r)$ *on segments between sites in which* $|T\rangle_c$ *contains an* $\left|R^{(k)}(L)\right\rangle$ *and an* $\left|R^{(k)}(R)\right\rangle$*, a 0-energy eigenstate on segments between an* $\left|R^{(k)}(L)\right\rangle$ *or* $\left|R^{(k)}(R)\right\rangle$ *and the boundary of the region, and* $|e\rangle$*'s everywhere*

else. Any eigenstate which is not an eigenstate of $R^{(k)}(H_q)(r)$ *on segments between sites in which* $|T\rangle_c$ *contains an* $\left|R^{(k)}(L)\right\rangle$ *and an* $\left|R^{(k)}(R)\right\rangle$ *has an energy* > 1 *(cf. Lemma 51 of [CPGW15a]).*

3. *The ground state energy is contained in the interval*

$$\Bigg[(g(k)-4^k\alpha_2(\varphi))LW-2^{-k}W+\sum_{n=1}^{\lfloor\log_4(L/2)\rfloor}\left(\left\lfloor\frac{W}{2^{2n+1(k\ mod 2)}}\right\rfloor\right.$$
$$\left.\times\left(\left\lfloor\frac{L}{2^{2n+1-(k\ mod 2)}}\right\rfloor-1\right)\right)\lambda_0(R^{(k)}(H_q)(4^{n-\lfloor(k\ mod 2)/2\rfloor})),$$
$$(g(k)-4^k\alpha_2(\varphi))LW-2^{-k}W+\sum_{n=1}^{\lfloor\log_4(L/2)\rfloor}\left(\left(\left\lfloor\frac{W}{2^{2n+1-(k\ mod 2)}}\right\rfloor+1\right)\right.$$
$$\left.\times\left\lfloor\frac{L}{2^{2n+1-(k\ mod 2)}}\right\rfloor\right)\lambda_0(R^{(k)}(H_q)(4^{n-\lfloor(k\ mod 2)/2\rfloor}))\Bigg]$$

where

$$g(k)=4^k\sum_{4^n+1<2^k}4^{-2n-1}\lambda_0(H_q(4^n)), \tag{5.74}$$

(cf. Lemma 52 of [CPGW15a]).

Proof. The proof of this lemma is long and does not contain much insight. For the sake of brevity we refer the reader to [WOC21, Appendix C]. The essentials of the proof are that (a) the classical RG map recreates the Robinson tiling Hamiltonian (b) the interactions between the "classical layer" and "quantum layer" are recreated and (c) the ground state of the Gottesman-Irani Hamiltonian is preserved. These factors mean the ground state is similar in form to the ground state of the original Hamiltonian, and thus almost all the relevant results about the original Hamiltonian apply to the renormalised Hamiltonian. □

Lemma 5.8. *Let* $S_{br}(k)$ *be the subspace spanned by states for which the left-most site is of the form* $\left|e^{\times p}\circledless\{x\}^{\times 2^k-p-1}\right\rangle$ *for a fixed integer* $1\leq p\leq 2^k$ *and the right-most*

site is of the form $\left|\{y\}^{\times 2^k-q-1} \circledS e^{\times q}\right\rangle$ *for fixed integer* $1 \leq q \leq 2^k$. *Then*

$$\lambda_0(R^{(k)}(H_q)(L)|_{S_{br}(k)}) = \min_{2^{k-1}L+1 \leq x \leq 2^k L} \lambda_0(H_q(x)) \tag{5.75}$$

Proof. $R^{(k)}(h_q)$ is block-diagonal with respect to the subspaces of $R^{(k)}(\mathcal{H}_{eq})^{\otimes 2}$ spanned by products of $\left|e^{\times p} \circledR \{x\}^{\times 2^k-p-1}\right\rangle$ and $\left|\{y\}^{\times 2^k-q-1} \circledS e^{\times q}\right\rangle$ for fixed p, q, together with the orthogonal complement thereof, while acting as identity on $R^{(k)}(\mathcal{H}_c)^{\otimes 2}$.

Thus the ground state energy is equal to $\min_{2^{k-1}L+1 \leq x \leq 2^k L} \lambda_0(H_q(x))$. □

Corollary 5.3. *If* $\lim_{L\to\infty} \lambda_0(H_u^{\Lambda(L)}) = +\infty$, *then* $\lim_{L\to\infty} \lambda_0(R^{(k)}(H_u)^{\Lambda(L)}) = +\infty$ *for all* $k \geq k_0(|\varphi|)$, *and* $k_0(|\varphi|)$ *is the smallest integer such that* $2^{k_0} > |\varphi|+7$. *If* $\lim_{L\to\infty} \lambda_0(H_u^{\Lambda(L)}) = -\infty$, *then* $\lim_{L\to\infty} \lambda_0(R^{(k)}(H_u)^{\Lambda(L)}) = -\infty$ *for all* $k \geq k_0(\varphi)$.

Proof. Consider applying the RG mapping $k > k_0(\varphi)$ times, then we see that

$$\begin{aligned} g(k) &= 4^k \sum_{4^n+1<2^k} 4^{-2n-1} \lambda_0(H_q(4^n)) && (5.76)\\ &= 4^k \sum_{4^n+1<2^{k_0}} 4^{-2n-1} \lambda_0(H_q(4^n)) + 4^k \sum_{2^{k_0}<4^n+1<2^k} 4^{-2n-1} \lambda_0(H_q(4^n)) && (5.77)\\ &= 4^k \alpha_2(\varphi) + 4^k \sum_{2^{k_0}<4^n+1<2^k} 4^{-2n-1} \lambda_0(H_q(4^n)). && (5.78) \end{aligned}$$

From lemma 5.7, the interval the ground state energy is contained in is

$$\begin{aligned} &\Bigg[LH \sum_{2^{k_0}<4^n+1<2^k} 4^{-2n-1} 4^k \lambda_0(H_q(4^n)) - 2^{-k} H \\ &+ \sum_{n=1}^{\lfloor \log_4(L/2) \rfloor} \left(\left\lfloor \frac{H}{2^{2n+1-(k \bmod 2)}} \right\rfloor \left(\left\lfloor \frac{L}{2^{2n+1-(k \bmod 2)}} \right\rfloor - 1 \right) \right) \lambda_0(R^{(k)}(H_q)(4^{n-\lfloor (k \bmod 2)/2 \rfloor})), \\ &LH \sum_{2^{k_0}<4^n+1<2^k} 4^{-2n-1} 4^k \lambda_0(H_q(4^n)) - 2^{-k} H \\ &+ \sum_{n=1}^{\lfloor \log_4(L/2) \rfloor} \left(\left(\left\lfloor \frac{H}{2^{2n+1-(k \bmod 2)}} \right\rfloor + 1 \right) \left\lfloor \frac{L}{2^{2n+1-(k \bmod 2)}} \right\rfloor \right) \lambda_0(R^{(k)}(H_q)(4^{n-\lfloor (k \bmod 2)/2 \rfloor})) \Bigg]. \end{aligned} \tag{5.79}$$

From lemma 5.8, if $\lambda_0(H_q(4^n+1))=0$ for all n, then $\lambda_0(R^{(k)}(H_q)(4^n+1))=0$ for all n. In this case the ground state energy becomes $\lambda_0(R^{(k)}(H)^{\Lambda(L)}) = -2^{-k}L \xrightarrow{L\to\infty} -\infty$.

We see that if for any n_0, $\lambda_0(H_q(4^{n_0}+1))>0$, then $\lambda_0(R^{(k)}(H_q)(4^n+1))>0$ $\forall n \geq n_0'$ (n_0' not necessarily equal to n_0). Define $g(k) =: \eta(k)+4^k\alpha_2(\varphi)$, where $g(k)$ is defined in eq. (5.74), then $\eta(k)\geq 0$, and we see that the lower bound of the ground state is

$$L^2\eta(k) - 2^{-k}L + \sum_{n=1}^{\lfloor \log_4(L/2)\rfloor} \left(\left\lfloor \frac{L}{2^{2n+1-(k \bmod 2)}} \right\rfloor \left(\left\lfloor \frac{L}{2^{2n+1-(k \bmod 2)}} \right\rfloor - 1\right)\right) \times \lambda_0(R^{(k)}(H_q)(4^{n-\lfloor (k \bmod 2)/2\rfloor})) \xrightarrow{L\to\infty} +\infty.$$

□

For $2^{k_0} \leq |\varphi|+7$ the above relationship is not necessarily preserved. To see why, note that for lengths $\ell \leq |\varphi|+7$ the Gottesman-Irani Hamiltonian will not encode the correct computation and hence will pick up some energy. Since $\lambda_0(R^{(k)}(H_q)(L)|_{S_{br}}) = \min_{2^{k-1}L+1\leq x\leq 2^k L}\lambda_0(H_q(x))$ rather than $\lambda_0(R^{(k)}(H_q)(L)|_{S_{br}}) = \lambda_0(H_q(x))$, the energies in the summation term and the α_2 term will not exactly cancel out until we reach higher order steps of the RG flow. This is only rectified once we reach $2^{k_0} > |\varphi|+7$ as the energy integrated out by the projector Π_{gs}, as given in definition 5.11, is exactly $\lambda_0(H_q(x))$, not $\lambda_0(R^{(k)}(H_q)(L)|_{S_{br}})$.

5.6.2 Renormalising H_d

The only part of the Hamiltonian acting on $\mathcal{H}_d$ is H_d; there is no coupling to other parts of the Hilbert space and so we can renormalise this part independently. Indeed, we can choose H_d to be any Hamiltonian with a dense spectrum that is a fixed point of an appropriate RG scheme. That is, H_d should be preserved by the RG scheme.

For concreteness, following [CPGW15a] (but with small changes), we will let H_d be the critical XY-model with local terms $X_i\otimes X_{i+1} + Y_i\otimes Y_{i+1} + \mu/2(Z_i\otimes \mathbb{1}^{(i+1)} +$

$\mathbb{1}^{(i)} \otimes Z_{i+1}$), which can be written as:

$$h_d^{row(i,i+1)} = X_i \otimes X_{i+1} + Y_i \otimes Y_{i+1}, \tag{5.80}$$

$$h_d^{col(i,i+1)} = 0, \tag{5.81}$$

$$h_d^{(1)(i)} = \mu Z_i. \tag{5.82}$$

This has zero gap for any $0 \geq \mu < 1$. Since the critical XY model is critical, it forms a fixed point in any reasonable RG scheme.

In particular, Penson, Jullien, and Pfeuty apply the BRG to renormalise this model [PJP82]. Notably, they show that there are multiple fixed points depending on the block size used: here since we are interested in blocking 2×2 blocks, we choose a block size of 2, and set μ equal to one of the relevant fixed point values, ensuring a gapless spectrum. The authors demonstrate that the coefficient of $h_d^{row(i,i+1)}$ and Z_i terms maintain a constant ratio to each other. This RG scheme can then be expressed in terms of an RG isometry V^d. Since the two coefficients maintain a constant ratio, the renormalisation unitary simply rescales these two terms.

5.6.3 Renormalising $|0\rangle$

If we wish to preserve the form of the possible ground states depending, it is straightforward to see that this can be done if the states $|0\rangle$ simply get mapped to themselves $|0\rangle^{\otimes(2\times2)} \to |0\rangle$ under the RG operation. This can be implemented using the isometry

$$V^0_{(i,i+1),(j,j+1)} := |0\rangle_{(i/2,j/2)} \langle 0|_{(i,j)} \langle 0|_{(i+1,j)} \langle 0|_{(i,j+1)} \langle 0|_{(i+1,j+1)}. \tag{5.83}$$

5.6.4 The Overall Renormalised Hamiltonian

Accounting for the renormalisation of all the different parts of the Hamiltonian, we can now define renormalisation group mapping for the entire Hamiltonian. Recall

that the original local terms are

$$h(\varphi)^{(i,j)} = |0\rangle\langle 0|^{(i)} \otimes (\mathbb{1} - |0\rangle\langle 0|)^{(j)} + (\mathbb{1} - |0\rangle\langle 0|)^{(i)} \otimes |0\rangle\langle 0|^{(j)} \tag{5.84}$$

$$+ h_u^{(i,j)}(\varphi) \otimes \mathbb{1}_d^{(i,j)} + \mathbb{1}_u^{(i,j)} \otimes h_d^{(i,j)} \tag{5.85}$$

$$h(\varphi)^{(1)} = -(1 + \alpha_2(\varphi))\Pi_{ud}, \tag{5.86}$$

where $\alpha_2(\varphi)$ is defined in lemma 5.7.

Definition 5.14 (Full Renormalisation Group Mapping)**.** *Let V^u, V^0, V^d be the isometries defined in definition 5.12, eq. (5.83), and section 5.6.2 respectively. Define:*

$$V^r_{(i,i+1),(j,j+1)} := V^0_{(i,i+1),(j,j+1)} \oplus \left(V^u_{(i,i+1),(j,j+1)} \otimes V^d_{(i,i+1),(j,j+1)} \right). \tag{5.87}$$

Then the overall RG mapping of local Hamiltonian terms is given by

$$\mathcal{R} : h(\varphi)^{(i,i+1)} \mapsto V^{r\dagger}_{(i,i+1),(j,j+1)} h(\varphi)^{(i,i+1)} V^r_{(i,i+1),(j,j+1)} \tag{5.88}$$

$$\mathcal{R} : h(\varphi)^{(i+1,i+2)} \mapsto V^{r\dagger}_{(i+2,i+3),(j,j+1)} h(\varphi)^{(i+1,i+2)} V^r_{(i,i+1),(j,j+1)} V^r_{(i+2,i+3),(j,j+1)} \tag{5.89}$$

Lemma 5.9. *Applying the RG mapping from definition 5.14 to the terms in eq. (5.84) we see that the renormalised 1- and 2-local terms become*

$$R^{(k)}(h(\varphi))^{(i,j)} = 2^k (|0\rangle\langle 0|^{(i)} \otimes \Pi_{ud}^{(j)} + \Pi_{ud}^{(i)} \otimes |0\rangle\langle 0|^{(j)}) \tag{5.90}$$

$$+ R^{(k)}(h_u(\varphi))^{(i,j)} \otimes \mathbb{1}_d^{(i,j)} + \mathbb{1}_u^{(i,j)} \otimes h_d^{(i,j)} \tag{5.91}$$

$$R^{(k)}(h(\varphi))^{(1)} = (g(k) - 4^k \alpha_2(\varphi) - 2^k)\Pi_{ud}^{(i)} + R^{(k)}(h_u^{(1)})^{(i)} \tag{5.92}$$

where $g(k)$ is defined in lemma 5.7. All the terms are computable.

Proof. Note that the RG isometry acts block-diagonally with respect to the subspaces spanned by $|0\rangle^{\otimes(2\times 2)}$ and those spanned by states in $(R^{(k)}(\mathcal{H}_u) \otimes \mathcal{H}_d)^{\otimes(2\times 2)}$. Furthermore, any state which is not in one of the two subspaces is projected out. The $h_u(\varphi)$, h_d and 1-local terms transform as they would in the absence of the $|0\rangle$ state, thus giving the terms seen above. The term $g(k)$ is computable for any k by calculating

the $\lambda_0(H_q)(4^n+1)$ for all $n \leq 2k+1$. Since this is a finite dimensional matrix for any finite n, this is a computable quantity.

The form of the overall renormalisation isometry means the $|0\rangle\langle 0|^{(i)} \otimes \Pi_{ud}^{(j)}$ term must be preserved in form, however, we note that because all states of 2×2 blocks in different subspaces in the previous RG step must be in $|0\rangle^{\otimes(2\times 2)}$ or $(R^{(k)}(\mathcal{H}_u) \otimes R^{(k)}(\mathcal{H})_d)^{\otimes(2\times 2)}$, then two neighbouring blocks must pick up an energy penalty of $\times 2$ of the previous local terms. □

Corollary 5.4. *The local terms of the initial Hamiltonian $h(\varphi)$ and all further renormalised local terms belong to a family of Hamiltonians $\mathcal{F}(\varphi, \tau_1, \tau_2, \{\alpha_i\}_i, \{\beta_i\}_i)$, which all take the form*

$$R^{(k)}(h(\varphi))^{(i,j)} = \tau_1(|0\rangle\langle 0|^{(i)} \otimes \Pi_{ud}^{(j)} + \Pi_{ud}^{(i)} \otimes |0\rangle\langle 0|^{(j)}) \tag{5.93}$$

$$+ R^{(k)}(h_u(\varphi, \{\beta_t\}_t))^{(i,j)} \otimes \mathbb{1}_d^{(i,j)} + \mathbb{1}_u^{(i,j)} \otimes R^{(k)}(h_d)^{(i,j)} \tag{5.94}$$

$$R^{(k)}(h(\varphi))^{(1)} = \tau_2 \Pi_{ud} + R^{(k)}(h_u(\varphi, \{\alpha_t\}_t))^{(1)}, \tag{5.95}$$

where the sets $\{\alpha_t\}_t$, $\{\beta_i\}$ characterises the parameters of the renormalised Gottesman-Irani Hamiltonian. Furthermore, for any $k \in \mathbb{N}$, the coefficients $\tau_1(k)$, $\tau_2(k)$, $\{\alpha_t(k)\}_t$ and $\{\beta_t(k)\}_t$ are computable.

Proof. Follows immediately from lemma 5.9. □

Lemma 5.10. *Let $R^{(k)}(h(\varphi))^{(i,j)}, R^{(k)}(h(\varphi))^{(1)}$ be the local terms defined by the RG mapping in definition 5.14 for any $k > k_0(|\varphi|)$. The Hamiltonian $R^{(k)}(H)$ defined by these terms then has the following properties:*

1. *If the unrenormalised Hamiltonian $H(\varphi)$ has a zero energy ground state with a spectral gap of 1/2, then $R^{(k)}(H)$ also has a zero energy ground state with zero correlations functions, and has a spectral gap of $\geq 2^{k-1}$.*

2. *If the unrenormalised Hamiltonian $H(\varphi)$ has a ground state energy $-\infty$ with a dense spectrum above this, then $R^{(k)}(H)$ also a ground state energy of $-\infty$ with a dense spectrum, and has algebraically decaying correlation functions.*

Proof. First examine the spectrum of the renormalised Hamiltonian from lemma 5.9: for convenience let

$$R^{(k)}(h_0)^{(i,j)} := 2^k(|0\rangle\langle 0|^{(i)} \otimes \Pi_{ud}^{(j)} + |0\rangle\langle 0|^{(i)} \otimes \Pi_{ud}^{(j)}). \tag{5.96}$$

Further let

$$R^{(k)}(H_0^{\Lambda(L)}) := \sum_{\langle i,j\rangle} R^{(k)}(h_0)^{(i,j)}, \tag{5.97}$$

$$R^{(k)}(\tilde{H}_u)^{\Lambda(L)} := \sum_{\langle i,j\rangle} \mathbb{1}_d^{(i,j)} \otimes R^{(k)}(h_u)^{(i,j)} \tag{5.98}$$

$$R^{(k)}(\tilde{H}_d)^{\Lambda(L)} := \sum_{\langle i,j\rangle} \mathbb{1}_u^{(i,j)} \otimes R^{(k)}(h_d)^{(i,j)} \tag{5.99}$$

We note $R^{(k)}(H_0)^{\Lambda}$,$R^{(k)}(\tilde{H}_d)^{\Lambda}$, $R^{(k)}(\tilde{H}_u)^{\Lambda}$ all commute. Further note that

$$\operatorname{spec} R^{(k)}(H_0)^{\Lambda} \subset 2^k\mathbb{Z}_{\geq 0}. \tag{5.100}$$

If $\lambda_0(H(\varphi)) = 0$, then it implies $\lambda_0(H_u(\varphi)) \to +\Omega(L^2)$ (see section 5.2.3). By corollary 5.3, this implies $\lambda_0(R^{(k)}(H_u(\varphi))) \to +\Omega(L^2)$ too. Hence the ground state is the zero-energy $|0\rangle^{\Lambda(L)}$ state. Since $\operatorname{spec} R^{(k)}(H_0)^{\Lambda} \subset 2^k\mathbb{Z}_{\geq 0}$, then the first excited state (provided L is sufficiently larger) has energy at least 2^k. Finally, the state $|0\rangle^{\Lambda(L)}$ has zero correlations.

If $\lambda_0(H(\varphi)) = -\Omega(L)$, then $\lambda_0(H_u(\varphi)) \to -\Omega(L)$ (see section 5.2.3). By corollary 5.3, this implies $\lambda_0(R^{(k)}(H)) \to -\Omega(L)$. Since $\operatorname{spec}(R^{(k)}(H_0)) \subset 2^k\mathbb{Z}_{\geq 0}$, then the ground state is the ground state of $R^{(k)}(\tilde{H}_d)^{\Lambda(L)} + R^{(k)}(\tilde{H}_u)^{\Lambda(L)}$. Since $\operatorname{spec}(R^{(k)}(\tilde{H}_d)^{\Lambda(L)})$ becomes dense in the thermodynamic limit, we see that the Hamiltonian has a dense spectrum in the thermodynamic limit. Let $|\psi\rangle_u$ and $|\phi\rangle_d$ be the ground states of $R^{(k)}(H_u)^{\Lambda(L)}$ and $R^{(k)}(H_d)^{\Lambda(L)}$ respectively, then the ground state of $R^{(k)}(\tilde{H}_d)^{\Lambda(L)} + R^{(k)}(\tilde{H}_u)^{\Lambda(L)}$ is $|\psi\rangle_u |\phi\rangle_d$. Since $R^{(k)}(H_d)^{\Lambda(L)}$ is just the critical XY-model and its ground state has algebraically decaying correlations [LSM61], hence the overall ground state has algebraically decaying correlations. □

5.6.5 Order Parameter Renormalisation

In section 5.2.5 we saw that the observable $O_{A/B}(r)$ functioned as an order parameter which distinguished the two phases. Defining $V_r := V^0_{(i,i+1),(j,j+1)} \oplus \left(V^u_{(i,i+1),(j,j+1)} \otimes V^d_{(i,i+1),(j,j+1)}\right)$, and $V_r[k]$ as the corresponding isometry for the k^{th} step of the RG process, then define

$$R^{(k)}(O_{A/B})(r) := V^r[k] O_{A/B}(2^k r) V^{r\dagger}[k]. \tag{5.101}$$

The following lemma then holds:

Lemma 5.11. *Let $\left|\psi_{gs}\right\rangle$ be the ground state of $R^{(k)}(H_u)$. The expectation value of the order parameter satisfies:*

$$\left\langle\psi_{gs}\right| R^{(k)}(O_{A/B})(r) \left|\psi_{gs}\right\rangle = \begin{cases} 1 & \text{if } \lambda_0(R^{(k)}(H)) = 0 \\ 0 & \text{if } \lambda_0(R^{(k)}(H)) = \Omega(L). \end{cases} \tag{5.102}$$

Proof. If $\lambda_0(R^{(k)}(H)) \longrightarrow -\Omega(L)$, then the ground state is that of $H_u^{(\Lambda(L))}$, and hence the state $|0\rangle$ does not appear anywhere in the ground state.
If $\lambda_0(R^{(k)}(H)) = 0$, the ground state is $|0\rangle^{\Lambda(L)}$. Since, under $V_r[k]$, $|0\rangle^{\otimes 2^k \times 2^k} \mapsto |0\rangle$, the lemma follows. □

Thus the renormalised order parameter still acts as an order parameter for the renormalised Hamiltonian. In particular, it still undergoes a non-analytic change when moving between phases.

5.6.6 Uncomputability of RG flows

We finally have all the ingredients for the proof of our two main results.

Theorem 5.6 (Exact RG flow for undecidable Hamiltonian)**.** *Let H be the Hamiltonian defined in [CPGW15a]. The renormalisation group procedure, defined in definition 5.14, has the following properties:*

1. *$\mathcal{R}(h)$ is computable.*

2. *If $H(\varphi)$ is gapless, then $R^{(k)}(H(\varphi))$ is gapless, and if $H(\varphi)$ is gapped, then $R^{(k)}(H(\varphi))$ is gapped.*

3. *For the order parameter of the form $O_{A/B}(r)$ which distinguished the phases of $H^{\Lambda(L)}$, there exists a renormalised observable $R^{(k)}(O_{A/B})(r)$ which distinguishes the phases of $R^{(k)}(H)^{\Lambda(L)}$ and is non-analytic at phase transitions.*

4. *For k iterations, the renormalised local interactions of $R^{(k)}(H)$ are computable and belong to the family $\mathcal{F}(\varphi, \tau_1, \tau_2, \{\beta_i\})$, as defined in corollary 5.4.*

5. *If $H(\varphi)$ initially has algebraically decaying correlations, then $R^{(k)}(H(\varphi))$ also has algebraically decaying correlations. If $H(\varphi)$ initially has zero correlations, then $R^{(k)}(H(\varphi))$ also has zero correlations.*

Proof. Claim 1 follows from definition 5.14, where the renormalisation isometries and subspace restrictions are explicitly written down and are manifestly computable, and hence for any k the coefficients in lemma 5.9 are computable. Claim 2 follows from lemma 5.10: we see that, for all $k > k_0$ the spectrum below energy 2^{k-1} is either dense with a ground state with energy at $-\infty$, or is empty except for a single zero energy state, corresponding to the gapped and gapless cases of $H(\varphi)$. Claim 3 follows from lemma 5.11. Claim 4 follows from corollary 5.4. Claim 5 follows from the properties of the ground states in the cases $\lambda_0(H_u^{\Lambda(L)}) \to \pm\infty$ and by lemma 5.10. □

Theorem 5.7 (Uncomputability of RG flow)**.** *Let $h(\varphi)$, $\varphi \in \mathbb{Q}$, be the full local interaction of the Hamiltonian from [CPGW15a]. Consider k iterations of the RG map from definition 5.14 acting on $H(\varphi)$, such that the renormalised local terms are given by $R^{(k)}(h(\varphi))$, which can be parameterised as per corollary 5.4.*

If the UTM is non-halting on input φ, then for all $k > k_0(\varphi)$ we have that $\tau_2(k) = -2^k$, for some computable $k_0(\varphi)$. If the UTM halts on input φ, then there exists an uncomputable $k_h(\varphi)$ such that for $k_0(\varphi) < k < k_h(\varphi)$ we have $\tau_2(k) = -2^k$, and for all $k > k_h(\varphi)$ then $\tau_2(k) = -2^k + \Omega(4^{k-k_h(\varphi)})$.

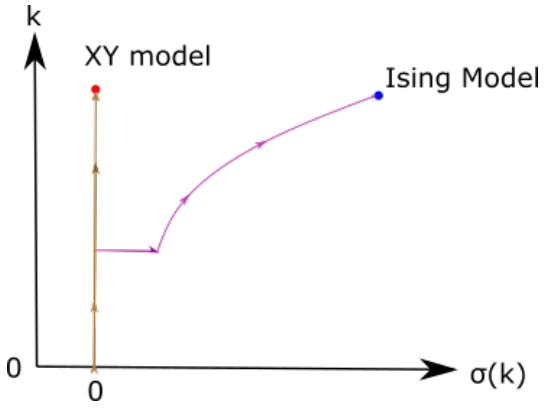

Figure 5.5: A schematic picture of the flow of Hamiltonians in parameter space. $\sigma(k)$ is defined in Orange represents some value of $\varphi = \varphi_0$ for which the QTM does not halt on input, while purple represents $\varphi = \varphi_0 + \epsilon$ for any algebraic number ϵ for which the QTM halts. For small k, the orange and purple lines coincide. Then at a particular value of k, $\sigma(k)$ becomes non-zero and then increases exponentially.

Proof. Consider the expression for τ_2 from lemma 5.9:

$$\tau_2(k) = 4^k \sum_{4^n+1<2^k} 4^{-2n-1}\lambda_0(H_q(4^n)) + 4^k\alpha_2(\varphi) - 2^k. \tag{5.103}$$

From the definition of $\alpha_2(\varphi)$, we see that there is a $k_0(\varphi) \in \mathbb{N}$ such that $g(k_0(\varphi)) = \alpha_2(\varphi)$, and hence we get

$$\tau_2(k) = -2^k + 4^k \sum_{2^{k_0(\varphi)}<4^n+1<2^k} 4^{-2n-1}\lambda_0(H_q(4^n)). \tag{5.104}$$

If the encoded QTM never halts, then by lemma 5.2 $\lambda_0(H_q(4^n)) = 0$ for all n such that $4^n + 1 > 2^{k_0(\varphi)}$. If the encoded UTM halts then by lemma 5.2 there exists an n_0 such that $\lambda_0(H_q(4^n)) > 0$ for all $n > n_0$. Then $k_h(\varphi)$ is defined as the minimum k such that $4^{n_0} + 1 < 2^{k_h(\varphi)}$. Thus determining $k_h(\varphi)$ is at least as hard as computing the halting time and thus is an uncomputable number.

□

5.7 Fixed points of the RG flow

Theorem 5.6 shows that our RG scheme satisfies the expected properties. We now qualitatively examine the Hamiltonian for large values of k.

5.7.1 Fixed Point for Gapped Instances

Here we show that for gapped instances the Hamiltonian becomes "Ising-like", for appropriately small energy scales. From corollary 5.4 the renormalised Hamiltonian is

$$\begin{aligned}
R^{(k)}(h^{row}(\varphi))^{(i,j)} =& 2^k(|0\rangle\langle 0|^{(i)} \otimes \Pi_{ud}^{(j)} + |0\rangle\langle 0|^{(i)} \otimes \Pi_{ud}^{(j)}) && (5.105)\\
&+ R^{(k)}(h_u^{row}(\varphi)')^{(i,j)} \otimes \mathbb{1}_d^{(i,j)} + \mathbb{1}_u^{(i,j)} \otimes R^{(k)}(h_d)^{(i,j)} && (5.106)\\
&+ 2^k \Pi_{ud}^{(i)} \otimes \Pi_{ud}^{(j)} && (5.107)\\
R^{(k)}(h^{col}(\varphi))^{(i,j)} =& 2^k(|0\rangle\langle 0|^{(i)} \otimes \Pi_{ud}^{(j)} + |0\rangle\langle 0|^{(i)} \otimes \Pi_{ud}^{(j)}) && (5.108)\\
&+ R^{(k)}(h_u^{col}(\varphi)')^{(i,j)} \otimes \mathbb{1}_d^{(i,j)} && (5.109)\\
R^{(k)}(h(\varphi))^{(1)} =& (g(k) - 4^k\alpha_2(\varphi) - 2^k)\Pi_{ud} + R^{(k)}(h_u^{(1)}(\varphi)), && (5.110)
\end{aligned}$$

where here we have explicitly separated out $\Pi_{ud}^{(i)} \otimes \Pi_{ud}^{(j)}$ from the term $R^{(k)}(h_u^{row}(\varphi))^{(i,j)} = R^{(k)}(h_u^{row}(\varphi)')^{(i,j)} + \Pi_{ud}^{(i)} \otimes \Pi_{ud}^{(j)}$.

Define the Ising-like Hamiltonian with local terms:

$$\begin{aligned}
h'^{row}_{Ising}(k)^{(i,j)} &:= 2^k \left(|0\rangle\langle 0|^{(i)} \otimes \Pi_{ud}^{(j)} + \Pi_{ud}^{(j)} \otimes |0\rangle\langle 0|^{(i)} + \Pi_{ud}^{(i)} \otimes \Pi_{ud}^{(j)} \right)\\
h'^{col}_{Ising}(k)^{(i,j)} &:= 2^k \left(|0\rangle\langle 0|^{(i)} \otimes \Pi_{ud}^{(j)} + \Pi_{ud}^{(j)} \otimes |0\rangle\langle 0|^{(i)} \right)\\
h'_{Ising}(k)^{(1)} &:= B(k)\Pi_{ud}.
\end{aligned}$$

This is reminiscent of the Ising interaction with both an ferromagnetic $|0\rangle\langle 0|^{(i)} \otimes |1\rangle\langle 1|^{(j)} + |1\rangle\langle 1|^{(i)} |0\rangle\langle 0|^{(j)}$ along the rows and columns and an anti-ferromagnetic $|1\rangle\langle 1|^{(i)} \otimes |1\rangle\langle 1|^{(j)}$ term along just the rows, with local field $B(k) = (g(k) - 4^k\alpha_2(\varphi) - 2^k)\,|1\rangle\langle 1|$, but with the orthogonal projector Π_{ud} playing the role of the projector onto the $|1\rangle\langle 1|$ state. However, note that Π_{ud} projects onto a larger dimensional subspace than $|1\rangle\langle 1|$, so e.g. the partition function of this Ising-like Hamiltonian is not identical to that of an Ising model.

We now show the following:

Proposition 5.1. *Let E be a fixed energy cut-off and $H'_{Ising}(k) = \sum_{\langle i,j\rangle} h'_{Ising}(k)^{(i,j)}$.*

Then

$$\left\| R^{(k)}(H(\varphi))|_{\leq E} - H'_{Ising}(k)|_{\leq E} \right\|_{op} \leq \left(\frac{E}{2^k}\right)^2. \tag{5.111}$$

Proof. Consider the local interaction term $h_0 = |0\rangle\langle 0| \otimes \Pi_{ud} + \Pi_{ud} \otimes |0\rangle\langle 0|$. This commutes with all other terms in both the $R^{(k)}(H(\varphi))$ Hamiltonian and the Ising-like Hamiltonian, and hence the eigenstates of both of the overall Hamiltonians are also eigenstates of $|0\rangle\langle 0| \otimes \Pi_{ud} + \Pi_{ud} \otimes |0\rangle\langle 0|$. As a result, for each eigenstate, a given site $p \in \Lambda$ either has support only on $|0\rangle_p$ or only on $R^{(k)}(\mathcal{H}_{ud})$. Therefore, an eigenstate defines regions (domains) of the lattice where all points in the domain are in $\mathcal{H}_{ud}$.

For a given eigenstate $|\psi\rangle$, let $D \coloneqq \left\{ i \in \mathbb{Z}^2 \middle| \operatorname{tr}\left(|0\rangle\langle 0|^{(i)} |\psi\rangle\langle\psi| \right) = 0 \right\}$ denote the region of the lattice where the state is supported on $R^{(k)}(\mathcal{H}_{ud})$, and ∂D be the set of sites on the boundary of D. Then we see that the terms in eq. (5.106) act non-trivially only within D, and that the boundaries of D receive an energy penalty of $2^k|\partial D|$ from terms in eq. (5.105) and eq. (5.108).

Note that $\left\| R^{(k)}(h_d)^{(i,j)} \right\|_{op}, \left\| R^{(k)}(h_u(\varphi)')^{(i,j)} \right\|_{op}, \left\| R^{(k)}(h_u^{(1)}(\varphi)) \right\|_{op} \leq 2$. For $\left\| R^{(k)}(h_d)^{(i,j)} \right\|_{op}$ this is straightforward to see. For $\left\| R^{(k)}(h_u(\varphi)')^{(i,j)} \right\|_{op}$, any states which pick up non-zero energy, other than those which receive a penalty due to halting, are removed from the local Hilbert space (as per section 5.5).

Let $m \in \mathbb{N}$ be a cut-off such that $|\partial D| \leq m$, hence $|D| \leq m^2/16$. Since for each boundary term we get an energy penalty of at least 2^k from h_0, we can relate m to the energy cut-off E to m as $E \coloneqq 2^k m$. If we consider the Hamiltonians restricted to

a subspace with energy $\leq E := 2^k m$, then

$$\left\| R^{(k)}(H(\varphi))|_{\leq E} - H'_{Ising}(k)|_{\leq E} \right\|_{op} \tag{5.112}$$

$$= \left\| \sum_{\langle i,j \rangle} \left(R^{(k)}(h_u(\varphi)')^{(i,j)} \otimes \mathbb{1}_d^{(i,j)} + \mathbb{1}_u^{(i,j)} \otimes R^{(k)}(h_d)^{(i,j)} \right) \Big|_{\leq E} \right\|_{op} \tag{5.113}$$

$$\leq \frac{m^2}{16} \left(\left\| R^{(k)}(h_u(\varphi)')^{(i,j)} \right\|_{op} + \left\| R^{(k)}(h_d)^{(i,j)} \right\|_{op} + \left\| R^{(k)}(h_u^{(1)}(\varphi)) \right\|_{op} \right) \tag{5.114}$$

$$\leq \frac{m^2}{2} \tag{5.115}$$

$$< \left(\frac{E}{2^k} \right)^2 . \tag{5.116}$$

Going from eq. (5.113) to eq. (5.114) we have used the fact that the terms in the sum are only non-zero within domains, and $|D| \leq m^2/16$. Going from eq. (5.114) to eq. (5.116) we have used the bound on the individual norms of the local terms. □

Thus, for appropriately small energies, we expect only small deviations from the "Ising-like" Hamiltonian. And these deviations vanish as the RG process is iterated. In particular, the spectrum will look like fig. 5.6.

5.7.2 Fixed Point for Gapless Instances

For a φ for which $H(\varphi)$ is gapless, $R^{(k)}(H(\varphi))$ is also gapless and we see that the ground state is that of $R^{(k)}(H_u(\varphi))$. If we restrict to a low energy subspace, one can see that excited states are either the excited states of the Gottesman-Irani Hamiltonians or the excited states of the critical XY-model. Indeed, let $E(k)$ be the subspace of states with energy less than 2^k, then for sufficiently large k we see that

$$R^{(k)}(H)^{\Lambda}|_{E(k)} = R^{(k)}(H_u(\varphi))^{\Lambda}|_{E(k)} \otimes \mathbb{1}_d^{\Lambda} + \mathbb{1}^{\Lambda} \otimes R^{(k)}(H_d)^{\Lambda}|_{E(k)}. \tag{5.117}$$

Since $R^{(k)}(H_d)^{\Lambda}|_{E(k)}$ has the same spectrum as H_d, the spectrum of $R^{(k)}(H)^{\Lambda}|_{E(k)}$ is also dense in the thermodynamic limit. Furthermore, $R^{(k)}(H)^{\Lambda}|_{E(k)}$ has algebraically decaying correlations since $R^{(k)}(H_d)^{\Lambda}|_{E(k)}$ also has algebraically decaying

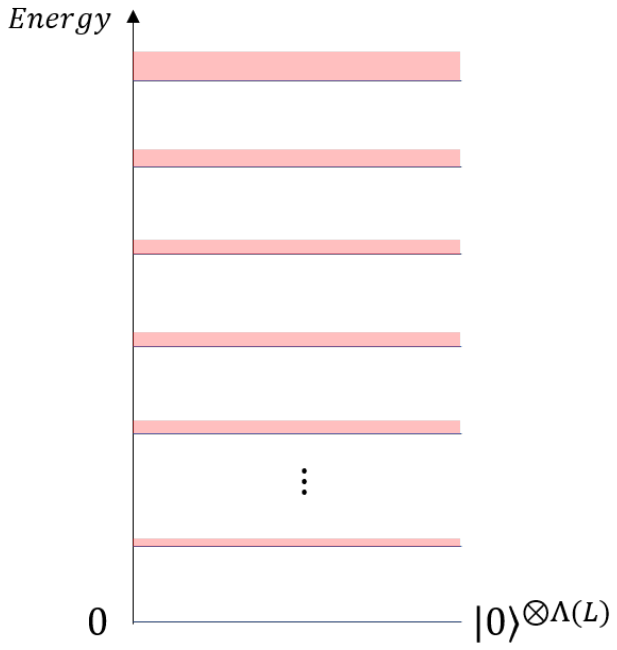

Figure 5.6: The energy level diagram of $R^{(k)}(H)$. The blue levels represent excitations of the $2^k(|0\rangle\langle 0|^{(i)} \otimes \Pi_{ud}^{(j)} + \Pi_{ud}^{(i)} \otimes |0\rangle\langle 0|^{(j)})$ term, while the red area represents the excited states of $R^{(k)}(h_u(\varphi)')^{(i,j)}$, $R^{(k)}(h_d)^{(i,j)}$, and $R^{(k)}(h_u^{(1)}(\varphi))$. The size of the red region increases as the domains get larger, and hence there are more high energy states. The ground state has no associated red region due to the presence of the spectral gap. The blue lines have an energy spacing of integer multiples of 2^k (although they are not necessarily as regular).

correlations [LSM61].

5.8 Discussion

We have seen under the renormalisation group procedure constructed here, the Hamiltonian flows towards either an Ising-like Hamiltonian or an XY-like Hamiltonian. Which case occurs depends on the parameter τ_2 in eq. (5.95). Let k be the number of iterations of the RG procedure, then from theorem 5.7 we see that there are two cases: $\tau_2 = -2^k$ always, or $\tau_2 = -2^k$ initially, and once a sufficiently large value of k is reached it begins to diverge as $\tau_2 > -2^k + \Omega(4^k)$. Determining which case occurs is undecidable. Moreover, the value of k at which we go from the first case to the second is uncomputable. Thus, determining the trajectory of the system for an arbitrary value of φ is uncomputable. Even if φ were known exactly, we see that the Hamiltonian's path in parameter space would be unpredictable.

Contrast this with chaotic behaviour: for chaotic systems, a tiny difference in the initial system parameters can lead to large diverges in trajectories later. Here the difficulty in predicting behaviour arises as it is usually difficult to determine the initial system parameters exactly. However, if the system parameters are known exactly, it should theoretically be possible to ascertain the long-time system. RG flows which undergo chaotic behaviour have been demonstrated before [MBK82; SKS82; DEE99; DT91; MN03].

The behaviour of the RG trajectory shown here is stronger than this in that even if the initial parameters characterising the microscopic interactions are known *exactly*, determining which fixed point the system may flow to is not possible to determine. We compare this to a similar uncomputability result in [Moo90] which showed that computing the trajectory of a particle in a potential is uncomputable.

The Hamiltonian discussed in this work is highly artificial and the RG scheme reflects this. Indeed, this Hamiltonian has an enormous local Hilbert space dimension and its matrix elements are functions of both φ and the binary length of φ, $|\varphi|$. Both of these factors are unlikely to be present in naturally occurring Hamiltonians. Thus an obvious route for further work is to consider RG schemes for more natural Hamiltonians which display undecidable behaviour.

Furthermore, although the RG scheme is essentially a simple BRG scheme, the details of its construction and analysis rely on knowledge of the structure of the ground states. Due to the behaviour of this undecidable model, any BRG scheme will have to exhibit similar behaviour to the one we have analysed rigorously here. But it would be interesting to find a simpler RG scheme for this Hamiltonian (or other Hamiltonians with undecidable properties) which is able to truncate the local Hilbert space to a greater degree, without using explicit a priori knowledge of the ground state, for which it is still possible to prove this rigorously.

The Hamiltonian and RG scheme constructed here could also be used to prove rigorous results for chaotic (but still computable) RG flows. Indeed, if we modify the Hamiltonian $H(\varphi)$ so that tead of running a universal Turing Machine on input φ, it carries out a computation of a (classical) chaotic process (e.g. repeated application

of the logistical map), then two inputs which are initially very close may diverge to completely different outputs after some time. By penalising this output qubit appropriately, the Hamiltonian will still flow to either the gapped or gapless fixed point depending on the outcome of the chaotic process under our RG map, but the RG flow will exhibit chaotic rather than uncomputable dynamics.

Given the RG scheme here, it is also relevant to ask is whether we can apply a similar scheme to the Hamiltonians designed in [Bau+18b] or chapter 3. Although we do not prove it here, we expect to be able to apply the modified BRG developed in this work to these Hamiltonians in an analogous way.

Chapter 6

Complexity of Measuring Local Observables for Systems with Circuit-to-Hamiltonian Mappings

6.1 Introduction

Given that much of condensed matter physics is devoted to determining the low-energy properties of materials, a natural question to ask is whether one can easily compute expectation values of observables at low temperatures. Since local measurements are the primary tools available to experimentalists for examining these systems, this is an extremely important problem. With this in mind, we consider the problem Approximate Simulation (APX-SIM), which asks how difficult it is to estimate the expectation of a local measurement against the low energy subspace of a local Hamiltonian.

Introduced in [Amb14], APX-SIM was shown to be $\mathsf{P}^{\mathsf{QMA}[\log]}$-complete for $\mathrm{O}(1)$-local Hamiltonians with 1-local measurements (with inverse polynomial precision) in [Amb14; GY19]. Here, $\mathsf{P}^{\mathsf{QMA}[\log]}$ is the class of decision problems decidable by a polynomial time (deterministic) Turing machine with access to logarithmically many adaptive queries to a QMA oracle (this is believed to be strictly larger class than QMA). More recently [GPY20] showed $\mathsf{P}^{\mathsf{QMA}[\log]}$-completeness for APX-SIM on physically motivated 2D models, such as for the Heisenberg interaction, and on

(non-translationally invariant) 1D chains.

The more realistic a system the Hamiltonian describes while retaining hard-to-compute properties, the more insight about the intrinsic complexity of the system one gains. For instance, a local spin dimension of 2 is often seen in nature (e.g. an electron spin up/down); a spin dimension of 2^{12} less so. Furthermore, condensed matter systems in real life often feature symmetries, such as a regular lattice structure with nearest-neighbour couplings that are both isotropic and translationally invariant. A goal of Hamiltonian complexity theory is thus to find increasingly "simple" systems that retain hard-to-compute properties. This renders claims more generic, and the resulting implications stronger.

Beyond their relevance for real-world systems, translationally-invariant Hamiltonians in particular are widely believed to be simpler than general Hamiltonians. Intuitively, due to the spatial invariance of the system, the degrees of freedom available to encode complex behaviour are limited; and less information can be encoded into the couplings throughout the system (assuming they are specified to the same precision).

As for spatial structure, the most basic lattice model is the one-dimensional spin chain, for which any hardness results often immediately imply respective hardness results for two- or higher-dimensional systems (by simply repeating the system along the extra dimensions). And just like translational symmetry, having only one dimension often renders systems more tractable, which is supported by empirical and theoretical evidence [LSM61; Aff+87; Fra17]. For instance, algorithms like DMRG [Whi92] to approximate ground state energies have long been known to work well in practice in 1D. Indeed, this eventually led to rigorous polynomial-time algorithms to solve ground state energy problems for gapped one-dimensional spin chains [LVV15]. Beyond that, there often exist closed form solutions in 1D, such as for fermionic 1D systems described by the Fermi-Hubbard model or 1D Heisenberg model [Bet31]; similar systems are notoriously difficult to simulate in higher dimensions.

So are one-dimensional, nearest neighbour, translationally-invariant (TI) systems tractable, or at least "more tractable" than their higher-dimensional counterparts?

Alas, results such as [GI09] show that finding the ground state of a 1D spin chain is $\mathsf{QMA}_{\mathsf{EXP}}$-complete ($\mathsf{QMA}_{\mathsf{EXP}}$ is a quantum analogue of NEXP), and the question of existence of a spectral gap in one spatial dimension—even for couplings with translational invariance—remains undecidable [Bau+18b]. Yet for other natural questions, such as APX-SIM, the verdict is still open.

Here, motivated by the goal of showing hardness of APX-SIM in the even simpler setting of 1D TI systems, we give a much more general framework for "lifting" hardness results about ground state energies (i.e. for LH) to hardness results for APX-SIM. Formally, these are given via the Lifting Lemma (lemma 6.4) and applications in section 6.3.5; here, we informally state the general premise as follows.

Theorem 6.1 (LH to APX-SIM (informal))**.** *If the family of Hamiltonians $\mathcal{F}$ admits a circuit-to-Hamiltonian mapping such that approximating the ground state energy is* C*-hard, then the APX-SIM problem for $\mathcal{F}$ is either* $\mathrm{P}^{\mathrm{C}[\log]}$ *or* P^{C}*-complete, depending on how the input is encoded.*

In contrast to previous approaches for showing hardness for APX-SIM, which were custom-designed based on the circuit-to-Hamiltonian constructions in mind, here we obtain a black-box mapping (lemma 6.4) which requires minimal assumptions, and which automatically *preserves structural properties of $\mathcal{F}$*, such as locality, geometry, translational invariance, etc. This demonstrates that the Local Hamiltonian problem fundamentally characterises the complexity of computing properties (such as simulating measurements) of the low energy states of a family of Hamiltonians.

6.2 Preliminaries

6.2.1 Approximate Simulation (APX-SIM)

The approximate simulation problem (APX-SIM) is concerned with the properties of ground states. However, it is usually more natural to consider a low-energy subspace. With this in mind we follow [GPY20] and define a more symmetric variant of the above problem APX-SIM which concerns itself with the low-energy subspace rather than the ground state:

Definition 6.1 ($\forall$-APX-SIM$(H, A, k, l, a, b, \delta)$ [GPY20])**.** *Given a k-local Hamiltonian $H = \sum_i H_i$ acting on N qubits, an l-local observable A, and real numbers a, b, and δ such that $b - a \geq N^{-c}$ and $\delta \geq N^{-c'}$, for $c, c' > 0$ constant, decide:*

YES. If for all $|\psi\rangle$ satisfying $\langle\psi| H |\psi\rangle \leq \lambda_{\min}(H) + \delta$, it holds that $\langle\psi| A |\psi\rangle \leq a$.

NO. If for all $|\psi\rangle$ satisfying $\langle\psi| H |\psi\rangle \leq \lambda_{\min}(H) + \delta$, it holds that $\langle\psi| A |\psi\rangle \geq b$.

The difference between APX-SIM and $\forall$-APX-SIM is that one requires *all* states below a threshold energy δ above the ground state energy to have expectation value upper-bounded by a. Throughout this work we consider the translationally invariant version of $\forall$-APX-SIM which simply for translationally invariant Hamiltonians:

Definition 6.2 ($\forall$-TI-APX-SIM)**.** *Defined analogously to $\forall$-APX-SIM, except the input Hamiltonian is specified via local term h of a translationally invariant Hamiltonian $H = \sum h$ acting on N qubits, where N is specified in binary, and each local term is describable in* $\mathrm{O}(\log(N))$ *bits.*

6.2.2 Useful Lemmas

We now state the Extended Projection Lemma, which consists of three claims, the first of which was given in [KKR06]. The lemma was later extended to include the second and third claims [GY19].

Lemma 6.1 (Extended Projection Lemma ([KKR06; GY19]))**.** *Let $H = H_1 + H_2$ be the sum of two Hamiltonians operating on some Hilbert space $\mathcal{H} = \mathcal{S} + \mathcal{S}^\perp$. The Hamiltonian H_1 is such that $\mathcal{S}$ is a zero eigenspace and the eigenvectors in $\mathcal{S}^\perp$ have eigenvalue at least $J > 2\|H_2\|_\infty$. Let $K \coloneqq \|H_2\|_\infty$. Then, for any $\delta \geq 0$ and $|\psi\rangle$ satisfying $\langle\psi| H |\psi\rangle \leq \lambda_{\min}(H) + \delta$, there exists a $|\psi'\rangle \in \mathcal{S}$ such that:*

- *(Ground state energy bound)*

$$\lambda_{\min}(H_2|_{\mathcal{S}}) - \frac{K^2}{J - 2K} \leq \lambda_{\min}(H) \leq \lambda_{\min}(H_2|_{\mathcal{S}}),$$

where $\lambda_{\min}(H_2|_{\mathcal{S}})$ denotes the smallest eigenvalue of H_2 restricted to space $\mathcal{S}$.

- *(Ground state deviation bound)*

$$|\langle\psi|\psi'\rangle|^2 \geq 1 - \left(\frac{K + \sqrt{K^2 + \delta(J-2K)}}{J-2K}\right)^2.$$

- *(Energy obtained by perturbed state against H)*

$$\langle\psi'|H|\psi'\rangle \leq \lambda_{\min}(H) + \delta + 2K\frac{K + \sqrt{K^2 + \delta(J-2K)}}{J-2K}.$$

Next, we state a quantum analogue of the union bound for commuting operators (see, e.g. [SKO19]).

Lemma 6.2 (Commutative Quantum Union Bound)**.** *Let $\{P_i\}_{i=1}^m$ be a set of pairwise commuting projectors, each satisfying $0 \leq P_i \leq I$. Then for any quantum state ρ,*

$$1 - \mathrm{tr}(P_m \cdots P_1 \rho P_1 \cdots P_m) \leq \sum_{i=1}^{m} \mathrm{tr}((I - P_i)\rho).$$

The following is a standard fact (see, e.g., Equation 1.33 [Gha13] for a proof):

$$\||v\rangle\langle v| - |w\rangle\langle w|\|_{\mathrm{tr}} = 2\sqrt{1 - |\langle v|w\rangle|^2} \leq 2\||v\rangle - |w\rangle\|_2. \tag{6.1}$$

6.2.3 Relevant (Oracle) Complexity Classes

As discussed extensively in [Koh+20, Sec. 4], the natural complexity class for the local Hamiltonian problem for a translationally-invariant system is $\mathsf{QMA}_{\mathsf{EXP}}$. Intuitively, this is because for translationally invariant system, the only parameter is the system size N. If N is the input, then it can be encoded in a string which has length $O(\log(N))$. Together with a promise gap which closes $\propto 1/T^2$, where T is the run-time of the embedded computation, and a $1/\mathrm{poly}(N)$ promise gap allowed in the definition of the local Hamiltonian problem, this means that we can only saturate this bound if we allow $T = \mathrm{poly}(N)$—i.e., the encoded computation runs in time exponential in the input size ($\mathrm{poly}(\log N)$), which naturally gives $\mathsf{QMA}_{\mathsf{EXP}}$.[1]

[1] We take care to distinguish this from the class $\mathsf{QMA}_{\mathsf{exp}}$ of [FL16] which is for an exponentially small promise gap in the input size, but polynomial length run time, also called PreciseQMA.

In [GPY20], the authors prove that APX-SIM is $\mathrm{P}^{\mathrm{QMA}[\log]}$-complete by showing that APX-SIM is $\mathrm{P}^{||\mathrm{QMA}}$-complete, and in a second step that $\mathrm{P}^{||\mathrm{QMA}} = \mathrm{P}^{\mathrm{QMA}[\log]}$. For the translationally invariant version of APX-SIM, we are essentially dealing with the same encoded computation, yet with exponentially input less information (i.e. the problem is succinctly encoded), just like in the case for the local Hamiltonian problem. Here, the base class now has a runtime $\mathrm{poly}(N)$, but an input of length $O(\log(N))$ and hence is an *exponential* time computation. The right class for which TI-APX-SIM is complete is thus $\mathrm{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$. Technically, this class comes out of the $\mathsf{EXP}^{||\mathsf{QMA}}$-completeness result, but it is also motivated in another way. Since $\mathsf{QMA}_{\mathsf{EXP}}$ has verification time $\mathrm{poly}(N)$, its circuit-to-Hamiltonian construction will have norm $\mathrm{poly}(N)$. Thus, the number of adaptive queries to the $\mathsf{QMA}_{\mathsf{EXP}}$ oracle to estimate the ground state energy within $1/\mathrm{poly}(N)$ error is $\log(N) = O(\mathrm{poly}(n))$, which is polynomial in the input size. So the P machine, which runs in time $\mathrm{poly}(\log(N)) = O(\mathrm{poly}(n))$, makes $\mathrm{poly}(\log(N)) = O(\mathrm{poly}(n))$ queries to the oracle, which yields $\mathsf{QMA}_{\mathsf{EXP}}$.

Oracle complexity classes. The classes $\mathrm{P}^{\mathrm{QMA}[\log]}$, $\mathrm{P}^{||\mathrm{QMA}}$, and $\mathrm{P}^{\mathrm{QMA}}$ denote the set of languages decidable by a polynomial-time deterministic Turing machine with access to, respectively, logarithmically many adaptive queries to a QMA oracle, polynomially many parallel queries to a QMA oracle, and polynomially many adaptive queries to a QMA oracle, respectively. It is known that $\mathrm{P}^{\mathrm{QMA}[\log]} = \mathrm{P}^{||\mathrm{QMA}}$ [GPY20]. The classes $\mathsf{EXP}^{||\mathsf{QMA}}$ and $\mathrm{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$ are defined analogously, except the former has an exponential-time deterministic Turing machine as its base (and hence can make exponentially many parallel queries), and where the oracle is for $\mathsf{QMA}_{\mathsf{EXP}}$, respectively.

6.3 Encoding Computation into Measurement Problems on Low Energy Spaces

6.3.1 Overview and Circuit-to-Hamiltonian Mappings

We point the reader to chapter 1 for an introduction to history states and circuit-to-Hamiltonian mappings. We consider an abstract route. More precisely, we will offload the question of proving $\mathsf{EXP}^{||\mathsf{QMA}}$-hardness of TI-APX-SIM to a quantum

	universal for	scaling of $T(N)$	symmetry	interaction graph	locality	local dimension	instance ℓ encoded in
[KSV02]	QMA	$\frac{1}{\text{poly}}$	none	arbitrary	5	2	local terms
[KR03b]				arbitrary	3	2	
[KKR06]				arbitrary	2	2	
[OT08]				2D planar	2	2	
[Aha+07]				1D line	12	2	
[BBT06]	StoqMA	$\frac{1}{\text{poly}}$	none	arbitrary	2	2	local terms
[GI09]	QMA_{EXP}	$\frac{1}{\text{poly}}$	translational	1D line	2	huge	
[BCO17]				1D line	2	42	system size
[BP17b]				3D fcc lattice	4	4	
[FL16]†	PreciseQMA	$\frac{1}{\text{exp}}$	none	arbitrary	3	2	local terms
[Koh+20]‡			translational	1D line	2	42	system size

Table 6.1: History state Hamiltonians from existing literature that satisfy definition 6.4 with varying characteristics. Shown the complexity class for which they encode a witness in the ground state energy, size T of history state in terms of the system size N, further properties of the Hamiltonian, as well as their dependence on the problem instances ℓ.
† references [KR03b] as the underlying construction; ‡ references [BCO17].

circuit that simulates the oracle calls. Starting from a rigorous definition of this type of computation we wish to encode, we will then require two mild assumptions on the type of circuit-to-Hamiltonian mapping used to translate the circuit to a ground state (the ability to access two outputs of the computation locally). Given the mapping has these two properties, $\text{EXP}^{\|\text{QMA}}$-hardness will follow.

6.3.2 Hamiltonians with a Universal Ground State

The specifics of the Hamiltonian used to encode computation is, for our purposes, irrelevant: whether it takes the shape of a history state Hamiltonian, features a more complicated clock construction such as in [Aha+07; GI09; BT14; BCO17], or is something completely different is not important.

More concretely, what we require of the circuit-to-Hamiltonian mapping is the ability to single out a low-energy subspace that encodes valid computations. This valid low-energy subspace needs to be separated from the rest of the spectrum, which may include states which do not encode a valid computation. Beyond this

fundamental requirement, and in order to translate a circuit that simulates oracle queries into a ground state energy problem, the circuit-to-Hamiltonian mapping needs to have two high-level properties (formal requirements in definition 6.4):

1. The possibility to inflict an energy penalty onto the output of a computation, which means that we can break the low-energy subspace of valid computations into a space V_{YES} and V_{NO} that correspond to computations that evolve correctly and incorrectly, respectively; and such that the largest eigenvalue of V_{YES} is below the smallest eigenvalue of V_{NO}.

2. Locality in the encoding, in the sense that neighbouring qubits in the circuit at some time-step need to map to neighbouring spins in the many-body system.

Both these assumptions are mild, and readers familiar with circuit-to-Hamiltonian mappings will immediately recognise that both of these are generally satisfied.

Definition 6.3 (Conformity)**.** *Let H be a Hamiltonian with some well-defined structure S (such as k-local interactions, all constraints drawn from a fixed finite family, with a fixed geometry such as 1D, translational invariance, etc). We say a Hermitian operator P* conforms *to H if $H+P$ also has structure S.*

For example, if H is a 1D translationally invariant Hamiltonian on qubits, then P conforms to H if $H+P$ is also 1D translationally invariant.

We now define an abstract notion of local circuit-to-Hamiltonian mappings; intuition given subsequently. The combination of an input penalty and history state Hamiltonian we denote with $H_{\text{w}} \coloneqq H_{\text{prop}} + H_{in}$.

Definition 6.4 (Local Circuit-to-Hamiltonian Mapping)**.** *Let $\mathcal{X} = (\mathbb{C}^2)^{\otimes m}$ and $\mathcal{Y} = (\mathbb{C}^2)^{\otimes n}$. A map $H_{\text{w}} : \mathrm{U}(\mathcal{X}) \mapsto \mathrm{Herm}(\mathcal{Y})$ is a* local circuit-to-Hamiltonian mapping *if, for any $T > 0$ and any sequence of 2-qubit unitary gates $U = U_T U_{T-1} \cdots U_1$, the following hold:*

1. *(Overall structure) $H_{\text{w}}(U) \geq 0$ has a non-trivial null space, i.e. $\mathrm{Null}(H_{\text{w}}(U)) \neq \emptyset$. This null space is spanned by (some appropriate notion of) "correctly initialized computation history states", i.e. with ancillae qubits set "correctly" and gates in U "applied" sequentially.*

2. *(Local penalization and measurement) Let q_1 and q_2 be the first two output wires of U (each a single qubit), respectively. Let $S_{\text{pre}} \subseteq \mathcal{X}$ and $S_{\text{post}} \subseteq \mathcal{Y}$ denote the sets of input states to U satisfying the structure enforced by $H_{\text{w}}(U)$ (e.g. ancillae initialized to zeroes), and null states of $H_{\text{w}}(U)$, respectively. Then, there exist projectors M_1 and P_T, projector M_2 conforming to $H_{\text{w}}(U)$, and a bijection $f : S_{\text{pre}} \mapsto S_{\text{post}}$, such that for all $i \in \{1,2\}$ and $|\phi\rangle \in S_{\text{pre}}$, the state $|\psi\rangle = f(|\phi\rangle)$ satisfies*

$$\text{Tr}\left(|0\rangle\langle 0|_i \left(U_T U_{T-1} \dots U_1\right) |\phi\rangle\langle\phi| \left(U_T U_{T-1} \dots U_1\right)^\dagger\right) = \text{Tr}\left(|\psi_T\rangle\langle\psi_T| M_i\right), \tag{6.2}$$

where $|\psi_T\rangle = P_T |\psi\rangle / \||P_T |\psi\rangle\|_2$ is $|\psi\rangle$ postselected on measurement outcome P_T (we require $P_T |\psi\rangle \neq 0$). Moreover, there exists a function $g : \mathbb{N}\times\mathbb{N} \mapsto \mathbb{R}$ such that

$$\||P_T |\psi\rangle\|_2^2 = g(m,T) \text{ for all } |\psi\rangle \in \text{Null}\left(H_{\text{w}}(U)\right), \tag{6.3}$$

$$M_i = P_T M_i P_T. \tag{6.4}$$

The map H_{w}, and all operators/functions above (M_1, M_2, P_T, f, g) are computable given U.

Intuition. We stress the following about definition 6.4:

1. It places no restrictions on the efficiency of computing H_{w}, M_1, M_2, P_T, f, g. Any such resource-restriction will later be application-dependent.

2. The term "local" in "local circuit-to-Hamiltonian mapping" is not referring to the locality of $H_{\text{w}}(U)$. Rather, it refers to the fact that local measurements on the output qubits of U can be simulated via local measurements on the ground space of $H_{\text{w}}(U)$ (up to postselection) via bijection f and eq. (6.2). Also, there is no restriction *a priori* on $g(m,T)$, other than $g(m,T) \neq 0$ for all $m, T \geq 0$.

3. For our applications, we only require simulation of local measurements on output qubits 1 and 2 of U; hence the phrasing of Point 2 in definition 6.4

(which, we note, makes this notion of local simulation milder than that used in universality results such as [CMP18; Koh+20; PB20]). Intuitively, the first qubit will encode whether U accepts or rejects (i.e. outputs 1 or 0, respectively). In the setting of APX-SIM, M_1 will play the role of observable A from definition 6.1; as such, M_1 need not necessarily conform to $H_{\mathrm{w}}(U)$. In contrast, M_2 will be used to penalize a certain "flag qubit" in our construction of lemma 6.3, and will be part of the Hamiltonian H from definition 6.1; as such, we require M_2 to conform to $H_{\mathrm{w}}(U)$. Finally, there is nothing particular about the choice of $|0\rangle\langle 0|_i$ in eq. (6.2); any fixed single-qubit projector would suffice.

4. Eq. (6.3) says all null states of $H_{\mathrm{w}}(U)$ have the same weight on the final time step, T. This is used, for example, in lemma 6.7 when we wish to exchange one $|\phi'\rangle \in S_{\mathrm{post}}$ with another state $|\phi\rangle \in S_{\mathrm{post}}$, and say something meaningful about the computation encoded in $|\phi'\rangle$ versus $|\phi\rangle$.

5. For the case of Kitaev's circuit-to-Hamiltonian construction [KSV02], P_T projects the clock register down to $|T\rangle$, and is 3-local since there the clock is encoded in unary. In all history state Hamiltonians, $g(m,T) = \|P_T|\psi\rangle\|_2^2 = 1/(T+1)$ for $|\psi\rangle$ a uniform history state.[2] Finally, eq. (6.4) captures the fact that M_i has a clock register projecting onto $|T\rangle\langle T|$, and so is supported solely on the Hilbert space corresponding to time T (i.e. projected onto by P_T).

6.3.3 Oracle Queries as Subroutines

Reducing $\mathrm{P}^{\|\mathsf{QMA}}$ to a single quantum verification circuit. We now give a generic construction for embedding an arbitrary $\mathrm{P}^{\|\mathsf{QMA}}$ circuit into a *single* quantum verification circuit (i.e. an "un-sound QMA circuit"). In doing so, as explained in [GY19], we must allow for the possibility that QMA oracle queries may be *invalid*, in that they might violate the QMA promise gap. This entails accounting for two potential obstacles: first, from an "oracle query" perspective, a QMA oracle fed an invalid QMA query may output 0 or 1 arbitrarily. Second, from a "QMA verification circuit"

[2] Modified history state Hamiltonians with non-uniform superpositions over the time steps such as in [BC18b] scale accordingly with the weight on the last timestep.

perspective, we cannot assume anything about the acceptance probability of the verifier when fed the optimal proof for an invalid instance, even after standard error reduction. Specifically, for a valid YES (resp., valid NO) QMA instance, standard error reduction implies the optimal proof is accepted with probability exponentially close to 1 (resp., 0); for an invalid instance, this optimal probability may still be $1/2$. The first of these obstacles requires one to define a valid $\mathsf{P}^{\|\mathsf{QMA}}$ machine's final output bit to be independent of how invalid queries are answered [Gol06] (otherwise, the output of the $\mathsf{P}^{\|\mathsf{QMA}}$ machine is not necessarily well-defined on a given input).

We now give the construction in lemma 6.3. Since the aim of this paper is *generic* reductions for lifting hardness results for one class of problems to another, we state the following lemma rather abstractly. For this, we first require two definitions, the first of which is standard.

Definition 6.5. *(Deterministic class) A set C of languages is a* deterministic class *if, for any language $L \in C$, there exists a deterministic Turing machine M which can decide L under the resource constraints specified by C. Formally, given any input $x \in \{0,1\}^n$, M halts after using $R(n)$ resources (where R may specify bounds on time or space), and accepts if $x \in L$ or rejects if $x \notin L$.*

Standard examples of deterministic classes include P, PSPACE, and EXP.

Definition 6.6. *(Existentially quantified quantum verification class (QVClass)) A set C of promise problems is an* existentially quantified quantum verification class *if any promise problem $A = (A_{\text{yes}}, A_{\text{no}}, A_{\text{inv}})$ in C satisfies the following. There exist computable functions $f, g, h : \mathbb{N} \mapsto \mathbb{N}$, as well as a deterministic Turing machine M which, for any input $x \in \{0,1\}^n$, uses $R(n)$ resources to produce a quantum verification circuit V (consisting of 1- and 2-qubit gates) and thresholds $c, s \in \mathbb{R}^+$ such that $c - s > 1/h(n)$. Here, $R(x)$ refers to resources such as time, space, etc, as required by C. The circuit V takes in a quantum proof $|\psi\rangle$ on $f(n)$ qubits, $g(n)$ ancilla qubits initialized to all zeroes, and has a designated output qubit, such that:*

- *(YES case) If $x \in A_{\text{yes}}$, there exists a quantum proof $|\psi\rangle$ on $f(n)$ qubits such that measuring the output qubit of $V|\psi\rangle|0\cdots0\rangle$ in the standard basis yields 1 with probability at least c.*

- *(NO case) If $x \in A_{\text{no}}$, for all quantum proofs $|\psi\rangle$ on $f(n)$ qubits, measuring the output qubit of $V|\psi\rangle|0\cdots0\rangle$ in the standard basis yields 1 with probability at most s.*

Without loss of generality, we assume the output qubit of V is the first wire exiting V.

For example, for QMA, $R(n)$ denotes a polynomial bound (with respect to n) on the number of time and space steps taken by M, while f and g are fixed polynomials, $c = 2/3$ and $s = 1/3$. In this way, classes such as NP, NEXP, QCMA, QMA, and so forth are examples of a QVClass.

We are now ready to state the following abstract lemma. As a concrete guiding example, consider D = P and Q = QMA below.

Lemma 6.3. *Let $x \in \{0,1\}^n$ be an instance of a problem in $\mathrm{D}^{\|\mathrm{Q}}$, where D is a deterministic class (such as P) and Q is a QVClass (such as QMA). Let U be a $\mathrm{D}^{\|\mathrm{Q}}$ machine (we will typically think of U as a circuit with access to an oracle for Q) making m parallel queries to a Q-oracle to decide x. Then, there exists a quantum circuit V with the following properties:*

1. *Given x and U, V can be computed in time polynomial in the sizes of U and the verifier for Q (both of which may be viewed as quantum circuits consisting of 1- and 2-qubit gates).*

2. *V takes as input $m+2$ registers: n input register A containing $x \in \{0,1\}^n$, m proof registers B_i containing a joint quantum proof $|w_{1\cdots m}\rangle$ (where ideally $|w_{1\cdots m}\rangle = |w_1\rangle|w_2\rangle\cdots|w_m\rangle$), with register B_i to be verified by a Q-circuit V_i (see fig. 6.1), and an ancilla register C which is assumed to be initialized to the all-zeroes state. Without loss of generality, we assume each verifier V_i has the same completeness and soundness parameters c and s, respectively.*

3. *V has two designated output wires: q_{out} is supposed to encode the output of U, and q_{flag} the number of Q queries made by U which were YES instances. (Without loss of generality, these are the first and second wires exiting V, respectively.) Formally, suppose V is fed the joint proof $|w_{1\cdots m}\rangle$, and let $|\psi\rangle$ denote the output state of V.*

Let sets S_0 and S_1 partition $\{0,1\}^m$ such that the D *machine underlying U rejects (resp. accepts) when given a string of query responses $y \in S_0$ (resp. $y \in S_1$). Define*

$$p_{y,w} := \Pr\left(\bigwedge_{i=1}^{m} V_i \text{ outputs } y_i \,\middle|\, |w_{1\cdots m}\rangle \right).$$

Note that in the ideal case $|w_{1\cdots m}\rangle = |w_1\rangle \cdots |w_m\rangle$, $p_{y,w}$ simplifies to

$$p_{y,w} = \prod_{i=1}^{m} \Pr(V_i \text{ outputs } y_i \mid |w_i\rangle).$$

In both cases,

$$\mathrm{tr}\left(|\psi\rangle\langle\psi| \cdot |1\rangle\langle 1|_{q_{\text{out}}} \right) = \sum_{y \in S_1} p_{y,w} \tag{6.5}$$

$$\mathrm{tr}\left(|\psi\rangle\langle\psi| \cdot |1\rangle\langle 1|_{q_{\text{flag}}} \right) = \sum_{y \in \{0,1\}^m} p_{y,w} \cdot \sin^2\left(\frac{\sqrt{3}}{2m} \cdot \mathrm{HW}(y) \right). \tag{6.6}$$

where $HW(y)$ is the Hamming weight of y.

Proof. As depicted in fig. 6.1, V is constructed by translating the D machine underlying U into a quantum circuit U', and then "simulating" the m (parallel) oracle calls U makes as sub-routines, in the sense of executing their Q-verification circuits V_i on the relevant subset of $|w_{1\cdots m}\rangle$ given by $|w_i\rangle\langle w_i| = \mathrm{Tr}_{B_j \neq B_i}[|w_{1\cdots m}\rangle\langle w_{1\cdots m}|]$. Note that U' is diagonal in the standard basis, i.e. is a classical circuit, and computes the inputs to the Q-verification circuits V_i on-the-fly; the lanes $|q_i\rangle$ in fig. 6.1 indicating the inputs to the respective verification subroutines are product states diagonal in the computational basis, i.e. $|q_1\rangle |q_2\rangle \cdots |q_m\rangle$.

The gate $R(\theta)$ in fig. 6.1 denotes the rotation matrix

$$R(\theta) = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix}.$$

To formally state the action of the overall circuit V, let X, Y, Z denote the input

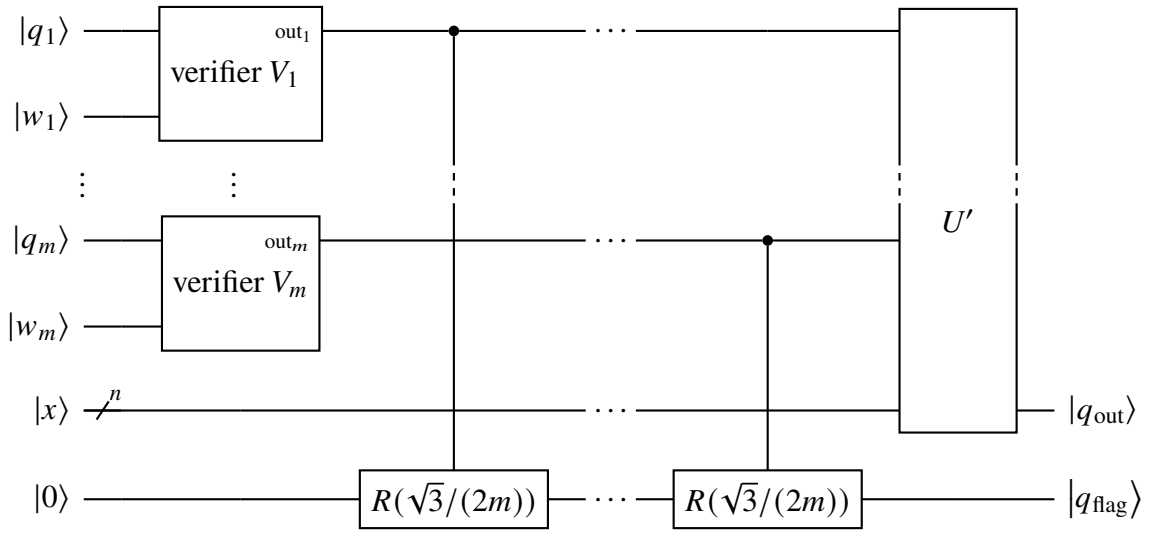

Figure 6.1: The circuit V constructed in lemma 6.3. The V_i are the Q-verifiers, each taking input $|q_i\rangle$ and proof/witness $|w_i\rangle$. (In principle, states $|w_i\rangle$ can be entangled as one joint state $|w_{1\cdots m}\rangle$; this is dealt with in the proof of lemma 6.7.) U' denotes the host postprocessing circuit in the original $\mathrm{D}^{\|\mathrm{Q}}$ circuit U, which takes the Q-query responses and outputs U's final answer. The gates $R(\sqrt{3}/(2m))$ denote a rotation in the standard basis of angle $\sqrt{3}/(2m)$. For simplicity, we have not depicted any preprocessing needed by U to compute the inputs $|q_i\rangle$ to the Q verifiers V_i, nor have we depicted the ancilla register C. For clarity, as a black box, the circuit V takes in the input to the U circuit, the joint proof $|w_{1\cdots m}\rangle$, and the ancilla register C.

registers to U' holding input $x \in \{0,1\}^m$, query response string $y = y_1 \cdots y_m$, and ancilla (initialized to all zeroes), respectively. Since U' is a classical circuit, without loss of generality it maps any

$$|x\rangle_X |y\rangle_Y |0\cdots 0\rangle_Z \mapsto |x\rangle_X |y\rangle_Y |0\cdots 0 f(y)\rangle_Z ,$$

where $f(y)$ is the output of U' (i.e. the D machine) given query response string y. If we now let F denote the flag qubit register, the output $|\psi\rangle$ of V is given by

$$|\psi\rangle = \sum_{y\in\{0,1\}^m} \alpha_y |x\rangle_X |y\rangle_Y |0\cdots 0 f(y)\rangle_Z \left(\cos\left(\frac{\sqrt{3}}{2m}\cdot \mathrm{HW}(y)\right)|0\rangle + \sin\left(\frac{\sqrt{3}}{2m}\cdot \mathrm{HW}(y)\right)|1\rangle\right)_F , \tag{6.7}$$

where for succinctness we omit registers such as those containing proof $|w_{1\cdots m}\rangle$, since U' does not act on these registers. Here, $\mathrm{HW}(y)$ is the Hamming weight of y,

and where

$$|\alpha_y|^2 = \Pr\left(\bigwedge_{i=1}^{m} V_i \text{ outputs } y_i \,\middle|\, |w_{1\cdots m}\rangle\right).$$

This immediately yields Equations (6.5) and (6.6). □

Remarks. We now make several remarks regarding lemma 6.3, including for the case when Q is a class of *promise* problems. For concreteness, in our discussion here we set D = P and Q = QMA, in which case the construction of lemma 6.3 runs in polynomial-time in n.

1. The construction of lemma 6.3 and fig. 6.1 says nothing about valid versus invalid QMA queries (i.e. when q_i is an invalid QMA query string in fig. 6.1). This will be dealt with in section 6.3.4.

2. When we later use lemma 6.3, we will penalize the flag qubit register F carefully so as to force all *valid* queries to be answered correctly. This is, in a nutshell, the purpose of the flag qubit.

6.3.4 Generic Hardness Constructions via a Lifting Lemma

We now proceed by encoding a family of $\mathrm{D}^{\|\mathrm{Q}}$ instances into a local circuit-to-Hamiltonian construction and penalise the *flag* qubit (not the output qubit!) to encourage the ground space of $H_{\mathrm{w}}(V)$ to encode correct query answers. We now comment on outstanding issues following lemma 6.3.

How to fix soundness for lemma 6.3. The reduction of lemma 6.3 is not *sound*, meaning a NO instance of $\mathrm{P}^{\|\mathsf{QMA}}$ was not necessarily mapped to a "NO QMA circuit" V. This is because, intuitively, each of the verifiers V_i in fig. 6.1 has a potentially different implicit quantifier for its proof register $|w_i\rangle$ (namely, $\exists$ for YES queries $|q_i\rangle$ and $\forall$ for NO queries $|q_i\rangle$).

By penalising the flag qubit, any YES query $|q_i\rangle$ which incorrectly has V_i outputting $|0\rangle$ is assigned an additional "unnecessary" energy penalty, lifting any such history state above the true ground state energy. (Crucially, we have no knowledge of the actual ground state energy itself, and this value encodes the number of YES and NO queries.) In short, adding the flag penalty will have the effect of

ensuring that the ground space will be spanned by states encoding computation for which as many "good" witnesses $|w_i\rangle$ as possible are fed into V.[3]

The Lifting Lemma. The main lemma of this section which encompasses all of the open points raised is as follows.

Lemma 6.4 (Lifting Lemma for APX-SIM)**.** *Let $x \in \{0,1\}^n$ be an instance of an arbitrary $\mathrm{D}^{\|\mathrm{Q}}$ problem, U a $\mathrm{D}^{\|\mathrm{Q}}$ machine deciding x, and V the verification circuit output by lemma 6.3. Fix a local circuit-to-Hamiltonian mapping H_{w}, and assume the notation in definition 6.4. Suppose, there exists a computable function $\alpha : \mathbb{N} \mapsto \mathbb{N}$, such that, for any ϵ satisfying*

$$0 \leq \epsilon \leq \frac{1}{\alpha}\left(\frac{1}{\alpha} + \frac{12\|M_2\|^2}{\Delta(H_{\mathrm{w}}(V))}\right)\left(\frac{8m^2}{3g(m,T)}\right), \tag{6.8}$$

Then the Hamiltonian $H \coloneqq \alpha(n)H_{\mathrm{w}}(V) + M_2$ satisfies:

- *If x is a YES instance, then for all $|\psi\rangle$ with $\langle\psi|H|\psi\rangle \leq \lambda_{\min}(H) + \frac{1}{\alpha^2}$,*

$$\langle\psi|M_1|\psi\rangle \leq g(m,T)\cdot m \cdot \max(1-c+\epsilon, s) + \frac{12\|M_2\|}{\alpha\Delta}.$$

- *If x is a NO instance, then for all $|\psi\rangle$ with $\langle\psi|H|\psi\rangle \leq \lambda_{\min}(H) + \frac{1}{\alpha^2}$,*

$$\langle\psi|M_1|\psi\rangle \geq g(m,T)\,(1 - m\cdot\max(1-c+\epsilon, s)) - \frac{12\|M_2\|}{\alpha\Delta}.$$

Remarks.

1. The Lifting Lemma's sole degree of freedom is the function α. All other quantities appearing, $\|M_2\|$ (flag penalty, lemma 6.3), Δ (spectral gap of $H_{\mathrm{w}}(V)$), $g(m,T)$ (weight of final time step T in history state, eq. (6.3)), c (completeness) and s (soundness), m (number of Q-queries), are fixed functions stemming from the choice of circuit-to-Hamiltonian construction (definition 6.4) and classes D and Q.

[3] While we generally call states of the form definition 2.14 history states, adding a penalty will result in a ground state as superposition of the same vectors, but with weights biased away from the location of the penalty in the time register.

2. If we think of M_1 as playing the role of observable A from definition 6.1 (where we follow the notation from definition 6.4), the Lifting Lemma brings us very close to obtaining $\mathrm{D}^{\|\mathrm{Q}}$-hardness of APX-SIM for the class of Hamiltonians produced by H_{w}. To formally "instantiate" such a hardness result for a fixed class $\mathrm{D}^{\|\mathrm{Q}}$, it remains to choose an appropriate function $\alpha(n)$ (n the input size) so that ϵ is sufficiently small so as to create an appropriate promise gap between the YES and NO thresholds.

3. Note that for $m \in \omega(1)$, the standard completeness/soundness parameters of $2/3$ and $1/3$ will not suffice to create a promise gap in the Lifting Lemma. Thus, error reduction for the QVClass Q to completeness $1 - 1/\mathrm{poly}(n)$ versus soundness $1/\mathrm{poly}(n)$ appears necessary. This, in particular, means our construction cannot *a priori* be applied with Q = StoqMA, since the latter is not known to have error reduction.

Proof of the Lifting Lemma. *Step i. Low energy states of H are close to uniform history states.* We first show that for a suitably-chosen function α, any low energy state with respect to H has a ground state of $H_{\mathrm{w}}(V)$ which is nearby. (For clarity, we define the spectral gap as difference between the two smallest *distinct* eigenvalues of A.)

Lemma 6.5. *For brevity, define shorthand Δ for $\Delta(H_{\mathrm{w}}(V))$. Fix any function $\alpha : \mathbb{N} \mapsto \mathbb{N}$ such that*

$$\alpha > \max\left(\frac{4\|M_2\|}{\Delta}, \frac{\Delta}{3\|M_2\|^2}, 1\right), \tag{6.9}$$

and any $\delta \leq 1/\alpha^2$. Then, for any $|\psi\rangle$ such that $\langle\psi| H |\psi\rangle \leq \lambda_{\min}(H) + \delta$, there exists a uniform history state $|\phi\rangle \in \mathrm{Null}\,(H_{\mathrm{w}}(V))$ such that

$$\| |\psi\rangle\langle\psi| - |\phi\rangle\langle\phi| \|_{\mathrm{tr}} \leq \frac{12\|M_2\|}{\alpha\Delta} \tag{6.10}$$

and where $|\phi\rangle$ has energy

$$\langle\phi| H |\phi\rangle \leq \lambda_{\min}(H) + \delta + \frac{12\|M_2\|^2}{\alpha\Delta}. \tag{6.11}$$

Proof. The ground state deviation bound of the Extended Projection Lemma (lemma 6.1), and Equations (6.1) and (6.9) imply the first claim via

$$\| |\psi\rangle\langle\psi| - |\phi\rangle\langle\phi| \|_{\mathrm{tr}} \leq 2 \frac{\left(\|M_2\| + \sqrt{\|M_2\|^2 + \frac{\Delta}{\alpha}} \right)}{\alpha\Delta - 2\|M_2\|} \leq \frac{6\|M_2\|}{\alpha\Delta - 2\|M_2\|} \leq \frac{12\|M_2\|}{\alpha\Delta}.$$

Here we have used eq. (6.9) to go from the second last to last expression in this chain of inequalities.

A similar calculation using the third claim of lemma 6.1 yields the second claim of this lemma. □

Step ii. All valid queries answered correctly. So far we have shown that all low-energy states with respect to H are close to unbiased history states, i.e. the states in the kernel of $H_{\mathrm{w}}(V)$. Yet M_2 penalizes the flag qubit, as formalized in lemma 6.3. What does this imply for the witness $|w_{1\cdots m}\rangle$? We now show that given a sufficiently small δ, the history state close to $|\psi\rangle$ encodes a series of queries such that as many of the valid queries as possible are answered as Yes.

Lemma 6.6. *Assume the notation of lemma 6.3, which showed*

$$\mathrm{tr}\left(|\psi\rangle\langle\psi| \cdot |0\rangle\langle 0|_{q_{\mathrm{flag}}} \right) = \sum_{y \in \{0,1\}^m} p_{y,w} \cdot \cos^2\left(\mathrm{HW}(y)\sqrt{3}/(2m) \right),$$

where $|\psi\rangle$ denoted the output of V given joint proof $|w_{1\cdots m}\rangle$, and $p_{y,w} = \Pr\left(\bigwedge_{i=1}^m V_i \text{ outputs } y_i \mid |w_{1\cdots m}\rangle \right)$. Suppose there exists an $i \in \{1,\ldots,m\}$ and $\epsilon > 0$ such that $|w_{1\cdots m}\rangle$ is "ϵ-suboptimal on proof i", meaning there exists a local proof $|w_i'\rangle$ such that

$$\Pr(V_i \text{ outputs } 1 \mid |w_{1\cdots m}\rangle) = \Pr\big(V_i \text{ outputs } 1 \mid |w_i'\rangle\big) - \epsilon. \tag{6.12}$$

Then there exists a proof $|w'_{1\cdots m}\rangle = |w_1'\rangle \otimes \cdots \otimes |w_m'\rangle$ which causes V to output $|\psi'\rangle$ satisfying

$$\mathrm{tr}\left(|\psi\rangle\langle\psi| \cdot |0\rangle\langle 0|_{q_{\mathrm{flag}}} \right) \geq \mathrm{tr}\left(|\psi'\rangle\langle\psi'| \cdot |0\rangle\langle 0|_{q_{\mathrm{flag}}} \right) + \frac{3}{8m^2}\epsilon.$$

Proof. The proof of this lemma is long as was primarily proved by Sevag Gharibian. As a result we omit it from this thesis and instead refer the reader to version in the paper [WBG20]. □

We now translate lemma 6.6, which held for circuits, to the following main lemma of Step ii, which gives us the corresponding desired statement for Hamiltonians (i.e. after the circuit-to-Hamiltonian construction is applied).

Lemma 6.7. *Suppose history state* $|\phi\rangle \in \mathrm{Null}\,(H_{\mathrm{w}}(V))$ *has preimage* $|\psi_{\mathrm{in}}\rangle = f^{-1}(|\phi\rangle)$ *(for bijection* f *from definition 6.4), where* $|\psi_{\mathrm{in}}\rangle$ *has proof* $|w_{1\cdots m}\rangle$*. If there exists* $\epsilon \geq 0$ *and* $i \in [m]$ *such that* $|w_{1\cdots m}\rangle$ *is* ϵ*-suboptimal on proof* i *(as defined in lemma 6.6), then*

$$\langle\phi|\, H\, |\phi\rangle \geq \lambda_{\min}(H) + \frac{3g(m,T)}{8m^2}\epsilon, \tag{6.13}$$

for $g(m,T)$ *defined in definition 6.4.*

Proof. Let $|\psi_{\mathrm{out}}\rangle = V|\psi_{\mathrm{in}}\rangle$ and $\left|\psi'_{\mathrm{out}}\right\rangle$ denote the output of V (fig. 6.1) when all proofs are proofs are optimal (i.e. set to $\left|w'_i\right\rangle$ in the terminology of lemma 6.6). Recall eq. (6.2) in the definition 6.4 of a local circuit-to-Hamiltonian mapping, which said M_2 simulates the projector $|0\rangle\langle 0|_{q_{\mathrm{flag}}}$ via

$$\mathrm{Tr}\left(|0\rangle\langle 0|_2 \left(U_T U_{T-1}\ldots U_1\right)|\psi_{\mathrm{in}}\rangle\langle\psi_{\mathrm{in}}|\left(U_T U_{T-1}\ldots U_1\right)^\dagger\right) = \mathrm{Tr}\left(|\phi_T\rangle\langle\phi_T|\, M_i\right), \tag{6.14}$$

(since we assumed in lemma 6.3 that the second output qubit of V is the flag qubit), where $|\phi_T\rangle$ is the history state $|\phi\rangle$ projected down onto time step T, which succeeds

with probability $g(m,T)$ (definition 6.4). We thus have

$$\begin{aligned}
\langle\phi|H|\phi\rangle &= \langle\phi|M_2|\phi\rangle && (6.15)\\
&= g(m,T)\langle\phi_T|M_2|\phi_T\rangle && (6.16)\\
&= g(m,T)\,\mathrm{tr}\left(|\psi_{\mathrm{out}}\rangle\langle\psi_{\mathrm{out}}|\cdot|0\rangle\langle 0|_{q_{\mathrm{flag}}}\right) && (6.17)\\
&\geq g(m,T)\left(\mathrm{tr}\left(\left|\psi'_{\mathrm{out}}\right\rangle\left\langle\psi'_{\mathrm{out}}\right|\cdot|0\rangle\langle 0|_{q_{\mathrm{flag}}}\right)+\frac{3}{8m^2}\epsilon\right) && (6.18)\\
&= \langle\phi'|M_2|\phi'\rangle + g(m,T)\left(\frac{3}{8m^2}\epsilon\right) && (6.19)\\
&\geq \lambda_{\min}(H)+\frac{3g(m,T)}{8m^2}\epsilon, && (6.20)
\end{aligned}$$

where the second statement follows from Equation eq. (6.4), the third from eq. (6.14), fourth from lemma 6.6, fifth from eq. (6.14) and defining $|\phi'\rangle = f(V^\dagger|\psi'_{\mathrm{out}}\rangle)$, and the last since $|\phi'\rangle \in \mathrm{Null}(H_{\mathrm{w}}(V))$ by the definition of f in definition 6.4 being a bijection. □

Step iii. Sufficiently high overlap with computation. We now argue that any low energy state of H must correctly encode the original $\mathrm{D}^{\|\mathrm{Q}}$ computation, and hence an appropriate measurement of a ground state will read off the $\mathrm{D}^{\|\mathrm{Q}}$ computation's answer.

Lemma 6.8. *Consider any* $|\psi\rangle$ *satisfying* $\langle\psi|H|\psi\rangle \leq \lambda_{\min}+\delta$. *If* $\delta \leq 1/\alpha^2$ *and*

$$\epsilon < \left(\delta+\frac{12\|M_2\|^2}{\alpha\Delta}\right)\left(\frac{8m^2}{3g(m,T)}\right), \tag{6.21}$$

then

- *if* x *is a YES instance for* $\mathrm{D}^{\|\mathrm{Q}}$, *then*

$$\mathrm{Tr}(|\psi\rangle\langle\psi|M_1) \leq g(m,T)\cdot m\cdot\max(1-c+\epsilon,s)+\frac{12\|M_2\|}{\alpha\Delta}. \tag{6.22}$$

- *if* x *is a NO instance for* $\mathrm{D}^{\|\mathrm{Q}}$, *then*

$$\mathrm{Tr}(|\psi\rangle\langle\psi|M_1) \geq g(m,T)\,(1-m\cdot\max(1-c+\epsilon,s))-\frac{12\|M_2\|}{\alpha\Delta}. \tag{6.23}$$

Proof. We first use lemma 6.5 to map, assuming $\delta \leq 1/\alpha^2$, $|\psi\rangle$ to a history state $|\phi\rangle \in \mathrm{Null}\,(H_{\mathrm{w}}(V))$ such that

$$\| |\psi\rangle\langle\psi| - |\phi\rangle\langle\phi| \|_{\mathrm{tr}} \leq \frac{12\|M_2\|}{\alpha\Delta} \quad \text{and} \quad \langle\phi|\,H\,|\phi\rangle \leq \lambda_{\min}(H) + \delta + \frac{12\|M_2\|^2}{\alpha\Delta}. \tag{6.24}$$

We next use lemma 6.7 to obtain that, for any ϵ satisfying

$$\epsilon < \left(\delta + \frac{12\|M_2\|^2}{\alpha\Delta}\right)\left(\frac{8m^2}{3g(m,T)}\right),$$

the preimage $|\phi_{\mathrm{in}}\rangle = f^{-1}(|\phi\rangle)$ contains proof $|w_{1\cdots m}\rangle$ which is not ϵ-suboptimal on any proof $i \in [m]$. In words, for any $i \in [m]$, if V_i has optimal acceptance probability p_i^*, then

$$\Pr(V_i \text{ outputs } 1 \mid |\phi_{\mathrm{in}}\rangle) \geq p_i^* - \epsilon.$$

Now, by lemma 6.3, if q_i is a YES query, $p_i^* \geq c$, if q_i is a NO query, $p_i^* \leq s$, and if q_i is invalid, then p_i^* can only be said to satisfy the trivial bounds $0 \leq p_i^* \leq 1$. Letting $S_{\mathrm{yes}}, S_{\mathrm{no}}, S_{\mathrm{inv}}$ denote the partition of $[m]$ corresponding to YES, NO, and INVALID queries, it follows via the commutative quantum union bound (lemma 6.2) that

$$\Pr\left(\bigwedge_{i\in[m]} V_i \text{ outputs correct answer} \mid |\phi_{\mathrm{in}}\rangle\right) \geq 1 - \left(\left|S_{\mathrm{yes}}\right|(1-c+\epsilon) + |S_{\mathrm{no}}|s\right) =: p_{\mathrm{good}}$$

where note $|S_{\mathrm{inv}}|$ does not appear since any answer to an invalid query is considered correct.

Finally, since by definition 6.5, U' in fig. 6.1 is a deterministic computation, it follows that V correctly accepts (respectively, correctly rejects) with probability at least p_{good}, given $|\phi_{\mathrm{in}}\rangle$, when the $\mathrm{D}^{\|\mathrm{Q}}$ instance x is a YES instance (respectively, NO instance). Recall from eq. (6.2) that measuring M_1 on a history state $|\phi\rangle = f(|\phi_{\mathrm{in}}\rangle)$ simulates a measurement on the output qubit of V (fig. 6.1) via

$$\mathrm{Tr}\left(|0\rangle\langle 0|_1\, V\,|\phi_{\mathrm{in}}\rangle\langle\phi_{\mathrm{in}}|\, V^\dagger\right) = \mathrm{Tr}\left(|\phi_T\rangle\langle\phi_T|\, M_1\right), \tag{6.25}$$

for $|\phi_T\rangle = P_T|\phi\rangle / \|P_T|\phi\rangle\|_2$. Thus, if x is a YES instance,

$$\begin{aligned} p_{\text{good}} &\leq \operatorname{tr}\left(V|\phi_{\text{in}}\rangle\langle\phi_{\text{in}}|V^\dagger \cdot |1\rangle\langle 1|_{q_{\text{out}}}\right) \\ &= 1 - \operatorname{tr}\left(V|\phi_{\text{in}}\rangle\langle\phi_{\text{in}}|V^\dagger \cdot |0\rangle\langle 0|_{q_{\text{out}}}\right) \\ &= 1 - \operatorname{Tr}(|\phi_T\rangle\langle\phi_T| M_1) \\ &= 1 - \frac{1}{g(m,T)} \operatorname{Tr}(|\phi\rangle\langle\phi| M_1), \end{aligned}$$

where the third/fourth statements follow from definition 6.4 and Equations (6.3) and (6.4). Via an analogous argument for the NO case, we conclude:

- If x is a YES instance for $\mathrm{D}^{\|\mathrm{Q}}$, then

$$\operatorname{Tr}(|\phi\rangle\langle\phi| M_1) \leq g(m,T)(1 - p_{\text{good}}) \leq g(m,T) \cdot m \cdot \max(1 - c + \epsilon, s).$$

- If x is a NO instance for $\mathrm{D}^{\|\mathrm{Q}}$, then

$$\operatorname{Tr}(|\phi\rangle\langle\phi| M_1) \geq g(m,T)\left(1 - m \cdot \max(1 - c + \epsilon, s)\right).$$

This was for history state $|\phi\rangle$. The claim now follows for the original state $|\psi\rangle$ in the claim by combining eq. (6.24) with Hölder's inequality. □

The Lifting Lemma now follows; we restate it below for convenience.

Lemma 6.4 (Lifting Lemma)**.** *Let $x \in \{0,1\}^n$ be an instance of an arbitrary $\mathrm{D}^{\|\mathrm{Q}}$ problem, U a $\mathrm{D}^{\|\mathrm{Q}}$ machine deciding x, and V the verification circuit output by lemma 6.3. Fix a local circuit-to-Hamiltonian mapping H_{w}, and assume the notation in definition 6.4. Then, there exists a computable function $\alpha : \mathbb{N} \mapsto \mathbb{N}$ such that, for any ϵ satisfying*

$$0 \leq \epsilon \leq \frac{1}{\alpha}\left(\frac{1}{\alpha} + \frac{12\|M_2\|^2}{\Delta(H_{\text{w}}(V))}\right)\left(\frac{8m^2}{3g(m,T)}\right), \tag{6.26}$$

the Hamiltonian $H \coloneqq \alpha(n)H_{\text{w}}(V) + M_2$ satisfies:

- *If x is a YES instance, then for all $|\psi\rangle$ with $\langle\psi|H|\psi\rangle \leq \lambda_{\min}(H) + \frac{1}{\alpha^2}$,*

$$\langle\psi|M_1|\psi\rangle \leq g(m,T)\cdot m\cdot \max(1-c+\epsilon, s) + \frac{12\|M_2\|}{\alpha\Delta}.$$

- *If x is a NO instance, then for all $|\psi\rangle$ with $\langle\psi|H|\psi\rangle \leq \lambda_{\min}(H) + \frac{1}{\alpha^2}$,*

$$\langle\psi|M_1|\psi\rangle \geq g(m,T)\,(1-m\cdot \max(1-c+\epsilon, s)) - \frac{12\|M_2\|}{\alpha\Delta}.$$

Proof. Follows immediately from lemma 6.8 by setting $\delta = 1/\alpha^2$. □

As mentioned earlier, to apply lemma 6.4 to particular classes $\mathrm{D}^{\|\mathrm{Q}}$, it remains to select α appropriately so that lemma 6.4 creates the desired promise gap between YES and NO cases.

6.3.5 Applying the Lifting Lemma

We now employ the Lifting Lemma (lemma 6.4) to obtain hardness results for various complexity classes and types of circuit-to-Hamiltonian constructions. For clarity, the lemma applies for any local circuit-to-Hamiltonian construction (definition 6.4) and class $\mathrm{D}^{\|\mathrm{Q}}$ for deterministic class D (definition 6.5) and QVClass Q (definition 6.6). What is required to apply it is to set α (relative to the other fixed quantities, which are $\|M_2\|$ (flag penalty, lemma 6.3), Δ (spectral gap of $H_{\mathrm{w}}(V)$), $g(m,T)$ (weight of final time step T in history state, eq. (6.3)), c (completeness) and s (soundness), m (number of Q-queries)) so that the desired problem gap is obtained in lemma 6.4. For a given QVClass we must be able to choose the completeness and soundness bounds must satisfy

$$m\cdot \max(1-c+\epsilon, s) = O\left(\frac{1}{\mathrm{poly}(n)}\right). \tag{6.27}$$

Application 1: Translationally-invariant systems.

Corollary 6.1. *$\forall$-TI-APX-SIM (and hence APX-SIM) is $\mathsf{EXP}^{\|\mathsf{QMA}}$-hard for a nearest neighbour, translationally invariant Hamiltonian on qudits of local dimension 44, for a 1-local observable A, and $\delta = 1/\exp(n)$, $b-a = \Omega(1/\exp(n))$.*

circuit-to-Ham. construction	interaction topology	measurement precision	APX-SIM variant	hardness
[OT08] [Aha+07]	2D planar, n-n, local dim 2 1D line, n-n, local dim 12	$\frac{1}{\text{poly}}$	APX-SIM	$\mathsf{P}^{\parallel\mathsf{QMA}}$
[FL16]	3-local, local dim 2	$1/\text{exp}$	APX-SIM	PSPACE
[BCO17] [BP17b]	t-i, 1D line, n-n, local dim 44† t-i, 3D fcc lattice, 4-local, local dim 4	$\frac{1}{\text{poly}}$	TI-APX-SIM	$\mathsf{EXP}^{\parallel\mathsf{QMA}}$ $\mathsf{EXP}^{\parallel\mathsf{QMA}}$
[Koh+20]	t-i, 1D line, n-n, local dim 42	$1/\text{exp}$	TI-APX-SIM	PSPACE

Table 6.2: Hardness of APX-SIM variants for various families of many-body systems. Measure precision is in terms of the system (Hamiltonian) size, not the input size. n-n stands for 2-local nearest-neighbour interactions, and t-i abbreviates translationally-invariant systems. Since PreciseQMA = PSPACE, and the latter is low for itself, $\mathsf{P}^{\parallel\mathsf{PreciseQMA}} = \mathsf{PSPACE}$.
† By corollary 6.1, two extra symbols are necessary as compared to the raw construction in table 6.1, increasing the local dimension from 42 to 44.

Proof sketch. Let $x \in \{0,1\}^n$ be the input given to the EXP machine.

- First note that the circuit V constructed in lemma 6.3 takes in as input x, and proof $|w_{1\cdots m}\rangle$. Since the EXP machine can produce QMA queries $|q_i\rangle$ of length *exponential* in n, this means that although x is n bits, the proof $|w_{1\cdots m}\rangle$ is on $\exp(n)$ qubits and the verifiers V_i are exponential-time. This exactly matches what is expected for a "$\mathsf{QMA}_{\mathsf{EXP}}$"-setup.

- Apply [GI09] as our local circuit-to Hamiltonian construction (where the local terms for this construction are describable in $\mathrm{O}(1)$ bits). However, if we were to write out the full circuit of lemma 6.3 during the reduction, it would take $\exp(n)$ time (since recall V has $\exp(n)$ size). This is not an issue: [GI09] implements the action of V from lemma 6.3, not as an explicit circuit, but via a *Quantum Turing Machines (QTM)* (i.e. each V_i is given via a QTM, and these are called as subroutines by a global QTM taking in x and $|w_{1\cdots m}\rangle$).

As per [GI09], $A = M_1$ and M_2 can now be chosen as constant-norm 2-local observables. Since, however, the length of the computations in EXP and verifiers V_i are exponential, we now require $\Delta, g(m,T)$ scaling as $1/\exp(n)$ (since m and T are in $\mathrm{O}(\exp(n))$). Similarly, since $m \in \exp(n)$, we use standard error reduction on the

verifiers V_i to ensure $c \geq 1 - 1/\text{poly}(N)$ (i.e. $1 - 1/\exp(n)$) and $s \leq 1/\text{poly}(N)$ (i.e. $s \leq 1/\exp(n)$). Setting α to be a sufficiently large fixed polynomial in N (i.e. a fixed exponential in n), we obtain a $1/\text{poly}(N)$ (i.e. $1/\exp(n)$) promise gap in lemma 6.4. Finally, due to our use of QTMs rather than explicit circuits, the reduction runs in time polynomial in n.

The reduced local dimension can be obtained by using [BCO17]'s translationally-invariant construction instead of [GI09]; the rest of the argument is the same. We obtain a 1-local observable A by incrementing the local dimension slightly (from 42 to 44; two extra letters for the quantum Thue system suffice to signal that the computation halted, as can be seen easily). □

Results for higher-dimensional lattices can be trivially obtained by repeating the Hamiltonian in a translationally-invariant fashion; a qualitatively different result for a 3D crystal lattice is the following:

Corollary 6.2. *∀-TI-APX-SIM (and hence APX-SIM) is* $\mathsf{EXP}^{\|\mathsf{QMA}}$*-hard for a* 4*-local, translationally invariant Hamiltonian on qudits of local dimension* 4 *on a 3D fcc lattice, for a 1-local observable A,* $\delta = 1/\exp(n)$*,* $b - a = \Omega(1/\exp(n))$*.*

Proof idea. Follows as in corollary 6.1, but using [BP17b]. □

Application 2: Exponential Precision.

Corollary 6.3. *∀-TI-APX-SIM (and hence APX-SIM) is* PSPACE*-hard for a nearest-neighbour translationally invariant Hamiltonian on qudits of local dimension* 44*, for a 1-local observable A,* $\delta = 1/\exp(n)$*,* $b - a = \Omega(1/\exp(n))$*.*

Proof idea. Use [Koh+20] (translationally invariant) as the circuit-to-Hamiltonian construction. As PreciseQMA = PSPACE by [FL16], we know that $\mathsf{P}^{\|\text{PreciseQMA}} = \mathsf{P}^{\|\text{PSPACE}} = \text{PSPACE} = \text{PreciseQMA}$. We encode the circuit from lemma 6.3, but modify it so that only a single query to a PreciseQMA oracle is made and have it such that the flag qubit and the output qubit are the same. We note that while we cannot generally amplify a PreciseQMA completeness-soundness gap polynomially close to 1 and 0 in polynomial time (otherwise PreciseQMA = QMA), respectively,

we *can* attain such amplification in *exponential* time in "circuit-world". Namely, we can apply the Marriott-Watrous strong error reduction technique [MW05], which will blow up the runtime of the circuit to exponential, but crucially keep the space usage polynomial. We can thus augment lemma 6.3 such that it performs such an amplification manually, so we can assume that the single oracle subquery has an output which, for valid queries, gives a probability polynomially close to 1 and 0.

We add a projector P which penalises the output/flag qubit. Finally, let the observable we wish to measure in APX-SIM also be the observable P. Then equivalents of lemma 6.5, lemma 6.5, and lemma 6.7 can be proven where the $1/\text{poly}$ factors are replaced by $1/\text{exp}$, and in lemma 6.7 is applies to just a single query. It is then possible to prove lemma 6.8 where the measurement $M_1 = P$. Since PreciseQMA has an exponentially small promise gap, and is a polynomial time computation, then measuring P with promise gap $b - a = \Omega(1/\exp(n))$ is PreciseQMA-hard. □

6.3.6 Containment Results

In this section, we speak of the $1/\text{poly}$-gap and $1/\text{exp}$-gap regimes to mean the settings in which all promise gaps in LH and APX-SIM are at least inverse polynomial and at least inverse exponential, respectively, in the input size, n.

Lemma 6.9. *Let $\mathcal{F}$ be a family Hamiltonians and Q a QVClass (definition 6.6) such that* $\text{LH}(g)$ *for $\mathcal{F}$ is in Q. In the* $1/\text{poly}$*-gap regime,* $\text{APX-SIM} \in \text{P}^{\text{Q}[\log]}$*, and in the* $1/\text{exp}$*-gap regime,* $\text{APX-SIM} \in \text{P}^{\text{Q}}$.

Proof. This is a straightforward application of the algorithm given in Section A.2 of [Amb14]. The algorithm first uses the Q-oracle to conduct a binary search in order to obtain an additive error η estimate of the ground state energy of H. (This is where the containment of $\text{LH}(g, N)$ in Q is used, for example.) With η in hand, we run one final Q-query to check if there is a low-energy state (relative to η) of H whose expected value for the observable A satisfies the YES-criteria for APX-SIM.

The only question is the level precision required, i.e. how should η scale? In the $1/\text{poly}$-gap regime, it suffices to use logarithmically many adaptive queries to the Q-oracle to obtain $\eta \in \text{O}(1/\text{poly}(n))$, as required to distinguish inverse polynomial

promise gaps. In the $1/\exp$-gap regime, we need polynomially many adaptive queries to obtain $\eta \in \mathrm{O}(1/\exp(n))$. This yields the claim. □

A similar, but not quite identical result is

Corollary 6.4. *APX-SIM where Q = PreciseQMA in corollary 6.3 is contained in* PSPACE.

Proof. The proof follows as for lemma 6.9 in the $1/\exp$-gap regime; we need $1/\mathrm{poly}\, N$ many queries to obtain $\eta \in \mathrm{O}(1/\exp(N))$. As $\mathrm{P}^{\mathrm{PreciseQMA}} = \mathrm{PSPACE}$ the claim follows. □

Finally, corollary 6.1 showed ∀-TI-APX-SIM (and hence APX-SIM) is $\mathsf{EXP}^{\|\mathsf{QMA}}$-hard (in the $1/\exp$-gap regime). This is a somewhat odd-looking complexity class; the following theorem shows it is actually equal to something "closer to $\mathrm{P}^{\mathsf{QMA}}$ in appearance".

Theorem 6.2. $\mathsf{EXP}^{\|\mathsf{QMA}} = \mathrm{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$.

Proof. For $\mathsf{EXP}^{\|\mathsf{QMA}} \subseteq \mathrm{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$, corollary 6.1 showed ∀-TI-APX-SIM (and hence APX-SIM) is $\mathsf{EXP}^{\|\mathsf{QMA}}$-hard in the $1/\exp$-gap regime. But since the family of Hamiltonians in TI-APX-SIM can be verified in the $1/\exp$-gap regime in $\mathsf{QMA}_{\mathsf{EXP}}$, lemma 6.9 says APX-SIM $\in \mathrm{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$.

The reverse containment, $\mathrm{P}^{\mathsf{QMA}_{\mathsf{EXP}}} \subseteq \mathsf{EXP}^{\|\mathsf{QMA}}$ is similar to Beigel's original proof [Bei91] that $\mathrm{P}^{\mathsf{NP}[\log]} \subseteq \mathrm{P}^{\|\mathsf{NP}}$: the EXP machine makes all possible queries within the decision tree of the P machine up front and in parallel, then simulates the P machine, along the way selectively using the results of whichever queries the adaptive nature of the P machine would have needed. The only catch is that normally, Hamiltonians verified in $\mathsf{QMA}_{\mathsf{EXP}}$ act on $\exp(n)$ qubits, whereas those in QMA act on $\mathrm{poly}(n)$ qubits. However, since an EXP Turing Machine can write exponentially long queries to the oracle, it can simply take the TI Hamiltonian H the P machine would have fed to the $\mathsf{QMA}_{\mathsf{EXP}}$ oracle (here we are using that the TI-local Hamiltonian problem is $\mathsf{QMA}_{\mathsf{EXP}}$-complete [GI09]), and explicitly write it out fully (i.e. write out all $\exp(n)$ local terms H_i), and feed it to the QMA oracle (which, being fed an $\exp(n)$-size input, is allowed an $\exp(n)$-size proof by definition). □

Corollary 6.5. *∀-TI-APX-SIM and TI-APX-SIM are both* $\mathrm{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$*-complete for families of Hamiltonians for which the Local Hamiltonian problem is* $\mathsf{QMA}_{\mathsf{EXP}}$*-complete.*

Proof. Follows from corollary 6.1, lemma 6.9, and theorem 6.2. □

Corollary 6.6. *∀-TI-APX-SIM is* $\mathrm{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$*-complete for a 1D, nearest neighbour, translationally invariant Hamiltonian for a 2-local observable A, and* $\delta = 1/\exp(n)$*,* $b - a = \Omega(1/\exp(n))$*.*

Proof. From corollary 6.5. □

6.4 Discussion

A higher-level goal is to "unify" the complexities of other computational problems with the hardness of the Local Hamiltonian problem, such that the completeness of the latter for a family of Hamiltonians (relative to, e.g., QMA, QCMA etc) immediately implies completeness of APX-SIM, GSCON, and potentially other low energy properties for the appropriate related complexity classes. Furthermore, the relation between the complexity of determining low energy properties and universality properties for classes Hamiltonians should be investigated.

Finally, a minor point to be addressed is how to make the lifting construction generalise to complexity classes for which we are unable to run error reduction: e.g. StoqMA (see section 6.3.4 for a discussion). The complexity of APX-SIM for Hamiltonian for which LH is StoqMA-complete is known to be $\mathrm{P}^{\mathrm{StoqMA}[\log]}$-complete, but we cannot simply include this at present as a consequence of our lifting construction.

Chapter 7

Complexity of Finding Critical Points

7.1 Introduction

In chapter 3 and chapter 5 we emphasised the importance of quantum phase transitions to modern physics, and in chapter 3 we demonstrated that determining the phase diagram of a Hamiltonian is uncomputable in general. However, it is clearly the case that there are subsets of materials for which we can compute the phase diagram including the classic example of the 1D quantum Ising model and frustration free, qubit Hamiltonians on a line [BG15]. Indeed one might expect that most Hamiltonians which are not perversely complicated have computable phase diagrams.

A necessary condition for a phase transition to take place is the closing of the spectral gap. A key tool in addressing the question of whether a system is gapped in the infinite-size limit that is the so-called "Knabe bound" [Kna88]. Loosely speaking, the Knabe bound is a "finite-size criterion" saying that, given a frustration-free Hamiltonian, if the spectral gap of the Hamiltonian decays slowly enough with the system size, then it is necessarily gapped in the thermodynamic limit. Over time, multiple improvements and variations of the Knabe bound have been derived [BG15; GM16; LM19; Lem20; Ans20]. Indeed, the Knabe bound has been used extensively to determine phases and spectral gaps of frustration-free systems, including variants of the AKLT model [AR+20; LSW20], to characterise the phases of translationally invariant Hamiltonians on 1D chains of qubits [BG15; GM16], and characterising gaps for Product Vacua with Boundary States (PVBS) models [LN19].

Finite-size criteria similar to this are often (implicitly) assumed in condensed matter physics when performing numerical studies: the idea that the phase (or gap) at finite lattice sizes allows us to extrapolate to the phase (or gap) in the thermodynamic limit [FS93; Tot+95; Yam97; GSMW13]. This is a very natural property — the idea that once a system that is "large enough" to reflect the thermodynamic limit, that the phase at that system size should reflect the phase in the thermodynamic limit.

Beyond the Knabe method, there exist other techniques for determining whether spectral gaps close or remain open, including the commonly-used Martingale Method [Nac96] and variants thereof [SS03; KL18]. The Martingale Method says that given an absorbing sequence of increasingly large sections of a lattice which tend towards the full lattice, there are three extra conditions placed on the local terms of the Hamiltonian which must be uniformly satisfied along the sequence. If so, there is a lower bound on the spectral gap. Furthermore, several numerical algorithms have been developed to compute spectral gaps, including variational algorithms [HWB19; Jon+19] and using density matrix renormalisation group techniques [CM17]

However, as shown in [CPGW15a; Bau+18b] and in chapter 3, we cannot resolve the phase/spectral gap question in general. Nonetheless, these undecidability and uncomputablity results only imply that it is impossible to study the phase diagrams or spectral gap of models in full generality; no claims are made regarding more restricted sub-types, which may well be solvable. By considering restricted families of Hamiltonians (e.g. frustration-free $1D$ qubit Hamiltonians as above), we can still hope to determine useful properties about these families.

Here we focus our attention on either one-parameter or two-parameter continuous family of Hamiltonian which satisfy a computable "finite-size criterion" similar in spirit to the Knabe bound, which gives an avenue to extrapolate to the infinite-size limit, given conditions on how the local gap scales. Furthermore, the family of many-body systems we include are promised to have a *single* critical boundary; either a single critical point for the one-parameter case or a critical line in the two-parameter setting. The two phases we delineate are a gapped and gapless phase, but can, in principle, be any other type of phase of interest (e.g. topological).

The finite size criterion assumes that if the gap decays sufficiently slowly (resp. quickly) for increasing lattice sizes, then the overall Hamiltonian must be gapped (resp. gapless) in the thermodynamic limit. This finite-size criterion immediately implies that the system cannot have an uncomputable phase in the thermodynamic limit, as solving for its phase on some finite-sized region is always a computationally bounded task (if intractable, and where we exclude pathological cases of uncomputable matrix entries or threshold sizes).

As a first contribution, we show that these finite size conditions place an upper bound on the computational complexity of approximating critical lines, by placing them within the class $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$—i.e. solvable by a polynomial-time Turing Machine with a $\mathsf{QMA}_{\mathsf{EXP}}$ oracle—and which is deemed "slightly harder" than just $\mathsf{QMA}_{\mathsf{EXP}}$ [Amb14].

For these restricted sets of one or two-parameter Hamiltonians, we show that the problem of determining the critical boundary in parameter space to even constant precision is $\mathsf{QMA}_{\mathsf{EXP}}$ (for the one-parameter case) resp. $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$-hard (for the two-parameter case).

Loosely speaking, we prove the two following theorems (the rigorous versions are given in section 7.3).

Theorem 7.1 (Informal)**.** *Let $\{H_N\}$ be a family of local translationally invariant Hamiltonians on an infinite lattice, indexed by $N \in \mathbb{N}$, and such that $H_N = H_N(\varphi)$ depends on an external parameter $\varphi \in [0,1]$. Suppose the following promises hold for Hamiltonians in this family: i. For all φ, the phase in the thermodynamic limit can be extrapolated from the order parameter obtained from a finite lattice size; and ii. There exists precisely one critical parameter φ^* such that for $\varphi < \varphi^*$, the system is in phase A; otherwise in phase B. Then the problem of determining φ^* to even constant precision is $\mathsf{QMA}_{\mathsf{EXP}}$-hard; and determining φ^* up to polynomial precision (in N) is contained in $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$.*

And the two-parameter case reads as follows.

Theorem 7.2 (Informal)**.** *Let the setup be as in theorem 7.1, but such that now $H_N = H_N(\varphi, \theta)$ depends on two parameters $\varphi \in [0, \mathrm{poly}\, N]$ and $\theta \in [0,1]$. Suppose*

the following promises hold for all Hamiltonians in this family: i. For all φ and θ, the phase in the thermodynamic limit can be extrapolated from the order parameter obtained from a finite lattice size; and ii. There exists precisely one critical boundary $\varphi^*(\theta)$ such that on one side the system is in phase A; on the other side in phase B. Then determining the critical boundary $\varphi^*(\theta)$ to constant precision is $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$-complete.

We emphasise that in theorem 7.2, the precision to which $\varphi^*(\theta)$ is to be resolved in the φ-direction is constant, but over a poly N parameter range, in contrast to the θ-direction, and the theorem 7.1 case. The two phase diagram cases that are hard to differentiate between are shown in fig. 7.4 for the 1-parameter case and fig. 7.1 for the 2-parameter case.

This means that even for systems satisfying the finite size criterion and a promise of a single phase transition within the parameter range, there is unlikely to be an efficient analytic method or efficient computational algorithm to determine gapped/gaplessness—or, more generally, phase diagrams (unless $\mathsf{QMA}_{\mathsf{EXP}}$ is efficiently solvable). In other words, even if we restrict to the set of Hamiltonians which are known to be gapped/gapless in the thermodynamic limit based on some finite-size scaling criteria, and even if we know that the regions within which the system is gapped or gapless are delineated in a smooth manner, determining the critical boundary between these regions is generally computationally intractable. Alternatively, any efficiently-computable condition cannot fully characterise the set of all many-body phase diagrams, even if the phase diagrams are simple, and even if the many-body systems satisfy finite-size scaling criteria.

Furthermore, this implies that for nearest-neighbour, translationally invariant Hamiltonians which have exactly two phases and a single critical boundary, determining the phase diagram to $\mathrm{O}(1)$ precision is computationally difficult. In particular, we show that there exists a Hamiltonian with two parameters (φ, θ) with a phase diagram which looks like one of the cases illustrated in fig. 7.1; however determining which one of the two is $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$-complete. We contrast this with the undecidability/uncomputability results which require an infinite number of phases, and which only satisfy a local-global promise for an uncomputably large system.

Overview of the Chapter. This chapter is organised as follows. In section 7.2 we set up necessary notation, give basic definitions and summarise some results derived elsewhere that our construction relies on. Furthermore, we formally define the local-global conditions (definitions 7.3 and 7.4), as well as the critical parameter problems 1-CRT-PRM and 2-CRT-PRM in definitions 7.5 and 7.6. In section 7.3, we summarise the two main results, leading up to the containment proof in section 7.4, and the two explicit hardness constructions in sections 7.5 and 7.6. In section 7.7, we finally give a fairly extensive discussion of further implications and open questions.

7.2 Preliminaries

7.2.1 Notation

We denote the evaluation of logical formula as with square brackets. For example, for $x, y \in \mathbb{R}$, $[x > y]$ is equal to 1 if $x > y$ and 0 if $x \leq y$. For $z \in \{0, 1\}$, then $[[x > y] \wedge z]$ is then the logically AND of $[x > y]$ and z. We further abbreviate the integer set $[n] \coloneqq \{0, \ldots, n-1\}$ for $n \in \mathbb{N}$ (context should make clear what the square brackets mean). For sums in large expressions we will often write $\sum_{y<x}$ to represent a sum over all y and x such that $y < x$, instead of $\sum_x \sum_{y<x}$. The domain and running variables in these cases will be clear from the context. Finally, in this chapter we will also be using the definitions of gapped and gapless used in definition 3.1 and definition 3.2.

7.2.2 Quantum Phase Transitions

We formally define one- and two- critical parameter problems rigorously in this section; to this end, we first need to rigorously introduce what we mean by the term *critical parameter* for a local Hamiltonian in the thermodynamic limit. We use the definition of quantum phase transition given in definition 2.1 (which is the definition given by Sachdev).

Definition 7.1 (Critical Parameter/Critical Point, φ^*)**.** *Let $H(\varphi, \{\theta_i\}_i)$ be a Hamiltonian defined in the thermodynamic limit which undergoes a phase transition as a function of φ. Then a* critical parameter *is a point at which the phase transition happens when all other parameters are held constant. We denote this point with φ^*.*

7.2.3 Finite Systems and Relation to the Thermodynamic Limit

For experimentalists with access to only finite-sized systems and finite precision measurements, it may not be possible to determine where non-analyticities occur in the ground state energy, and thus to determine where the phase transitions take place. Instead of pinpointing the critical point itself, measurements that indicate which phase one is currently in are often used instead; multiple such observations then allow bounding the exact location where a phase transition is likely to occur.

Order Parameters. Commonly an *order parameter* is an observable which characterizes a certain property of a phase and is used to distinguish phases for a system. One possible definition is given as follows.

Definition 7.2 (Efficiently Computable Local Order Parameter)**.** *Assume a translationally invariant Hamiltonian* $H^{\Lambda(L)}(\varphi) = \sum_{\langle i,j\rangle} h_{i,j}(\varphi)$ *has two phases A and B, and let* $h_{i,j}(\varphi)$ *be describable in n bits. A* local order parameter $O_{A/B}$ *is a projector* $O_{A/B}$ *that is* $O(1)$*-local and computable in* $\mathrm{poly}(n)$ *time, if* $\langle O_{A/B}\rangle$ *undergoes a non-analytic change between phases A and B at a critical parameter* φ^**.*

All order parameters we will use will trivially fall into this category, as can be easily verified. We note that compared to our definitions of order parameters used in previous chapters, here we have a constraint on how easily the order parameter is computable. For the most part, this is just a technicality.

For a finite system, examining how the order parameter of the systems changes could yield insight over the phase that the system is in when approaching the thermodynamic limit, assuming that for a particular set of parameter values $\varphi, \theta, \ldots$, and a sufficiently large systems size the system is in the same phase as in the limit. While this is not generically the case [CPW15; Bau+18b; BCW19], many Hamiltonians studied in nature *do have* this property: for sufficiently large but finite sizes, the order parameter indicates the same phase as it would in the thermodynamic limit [PT81; BZJ85; HB88].

We formalise this notion of being able to extrapolate from a finite-size system to the thermodynamic limit with the following definitions.

Definition 7.3 (Local-Global Phase Condition)**.** *A translationally invariant Hamiltonian* $H^{\Lambda(L)}(\varphi) = \sum_{\langle i,j\rangle} h_{i,j}(\varphi)$ *defined on a square lattice* $\Lambda(L)$ *is defined to have the* locally-globally phase property *if in the thermodynamic limit it has two phases A and B distinguished by an order parameter* $O_{A/B}$*. If* $h_{i,j}(\varphi)$ *is describable in n bits, then there exists an integer* $L_0 = \mathrm{O}(\exp(\mathrm{poly}(n)))$*, an* $\omega = \Omega(1/\mathrm{poly}(L))$ *and polynomials* p, q *with* $1/p(L) - 1/q(L) = \Omega(1/\mathrm{poly}(L))$ *such that for the states* $|\psi\rangle$ *satisfying* $\langle\psi| H^{\Lambda(L)}(\varphi) |\psi\rangle \leq \lambda_{\min}(H^{\Lambda(L)}(\varphi)) + \omega$ *the following holds:*

Phase A: if for $L = L_0$ *and* $|\varphi - \varphi^*| \geq 1/\mathrm{poly}(L)$*, such that* $\langle\psi| O_{A/B} |\psi\rangle \leq 1/p(L)$*, then* $H^{\Lambda(L)}(\varphi)$ *is in phase A for all* $L \geq L_0$ *and in the thermodynamic limit.*

Phase B: if for all $L = L_0$ *and* $|\varphi - \varphi^*| \geq 1/\mathrm{poly}(L)$*, such that* $\langle\psi| O_{A/B} |\psi\rangle \geq 1/q(L)$*, then* $H^{\Lambda(L)}(\varphi)$ *is in phase B for all* $L \geq L_0$ *and in the thermodynamic limit.*

Furthermore, L_0 *is independent of* φ *and is computable in time* $\mathrm{poly}(n)$.

We note that this condition is fulfilled by many Hamiltonians. For example, for the transverse quantum Ising model in 1D, a phase transition occurs in terms of the ratio of magnetic field strength to coupling strength. In this case, we see that magnetization along the Z direction acts as an order parameter and satisfies the above condition [Sac11].

The Spectral Gap. Treating the Local-Global phase condition in a more generic sense, i.e. as a property that holds locally and lets you deduce thermodynamic properties, unveils other familiar conditions that allow a similar picture. One such set of conditions are Knabe bounds [Kna88]. In brief, Knabe bounds treat the case where one phase has is gapped while the other is gapless, and the gapped case can be discriminated by the gap behaviour on finite-sized systems.

More concretely, the condition is such that the spectral gap remains open—i.e. lower-bounded by a constant—in the thermodynamic limit if it closes sufficiently slowly at finite lattice sizes. In other words, if we take a Hamiltonian restricted to a finite but growing region of the interaction graph, and if the spectral gap on this restricted graph closes slowly enough then the system is guaranteed to be gapped in

the thermodynamic limit. Such a Local-Global gap bound has been shown to exist for the unfrustrated case [Kna88].

In particular, Knabe proved that if the local gap on N spins is larger than the threshold value $1/(N-1)$ for some $N > 2$ the system is gapped in the thermodynamic limit [Kna88]. Recently improvements to this bound have been made such that the threshold value is $6/N(N+1)$ which is known to be asymptotically optimal [GM16]. Another well-known method of bounding spectral gaps is given by Nachtergaele [Nac96], successively improved by [SS03; KL18]. Based on relations between ground space projectors, it says that if a model is gapless then the spectral gap cannot decay too slowly with system size.

We define a "finite-size criterion" which states that provided the spectral gap decays sufficiently quickly/slowly as the lattice size increases above some computable lattice size L_0, then the Hamiltonian is gapless/gapped in the thermodynamic limit. Given this criterion, we examine the complexity of Hamiltonians which satisfy it:

Definition 7.4 (Local-Global Gap Condition)**.** *A translationally invariant Hamiltonian $H^{\Lambda(L)}(\varphi) = \sum_{\langle i,j\rangle} h_{i,j}(\varphi)$ defined on a square lattice $\Lambda(L)$ is defined to be* locally-globally gapped *if there exist polynomials p, q with $1/p(L) - 1/q(L) = \Omega(1/\mathrm{poly}(L))$ and $L_0 \in \mathbb{N}$ such that*

Gapped: if $\Delta(H^{\Lambda(L)}(\varphi)) \geq 1/p(L)$ for some $L = L_0$, then there exists a constant $c > 0$ such that in the thermodynamic limit the spectral gap $\lim_{L\to\infty} \Delta(H^{\Lambda(L)}(\varphi)) \geq c$.

Gapless: if $\Delta(H^{\Lambda(L)}(\varphi)) \leq 1/q(L)$ for some $L = L_0$, then the Hamiltonian is gapless in the thermodynamic limit.

Furthermore, L_0 is independent of φ. Let $h_{i,j}(\varphi)$ be describable in n bits, then L_0 is computable in time $\mathrm{poly}(n)$ and $L_0 = \mathrm{O}(\exp(\mathrm{poly}\, n))$.

This is in part motivated by Knabe's bound: if we consider the gapped case as providing a lower bound (and ignoring the gapless case with its $1/q(L)$ bound) then we get a version of Knabe's bound where it is explicitly promised that the system size for which the "gappedness" can be verified at is poly-time computable.

Knabe's bound states that provided for some system size $N > 2$ if the spectral gap is greater than $1/(N-1)$ then the Hamiltonian is gapped in the thermodynamic limit. However, the point at which this happens can be at *any* N. As such this does not necessarily conflict with the possibility of undecidability results for frustration-free Hamiltonians (similar to those proved in [CPW15]). Instead, the size of the systems for which the Knabe bound holds would be uncomputable (and thus actually implementing the bound is uncomputable). The Local-Global gap we have assumed to hold in definition 7.4 is explicitly computable and L_0 is assumed to be computable in polynomial time, thus forcing the problem of determining the spectral gap to be decidable for systems satisfying definition 7.4.

As such our promise is strictly stronger than the regular Knabe bound. It is further worth noting that the polynomials that bound the gaps will vary between different classes of Hamiltonians and that by standard padding arguments the exact scaling of these polynomials is not overly relevant: if, for instance, the gapped/gapless property of the many-body system is determined by patches distributed across the system, where the density of the patches goes as $\sim 1/\sqrt{L}$ in the system size L, the polynomials $p(L)$ and $q(L)$ can effectively be replaced by $p(\sqrt{L})$ and $q(\sqrt{L})$. This argument is the same as for the local Hamiltonian problem, where the $1/\text{poly}$ promise gap can be magnified with the same technique—naturally without any interesting implications regarding the problems' complexity, or the resulting phenomenology exhibited by the system.

In this work, we shall consider different classes of Hamiltonians that have distinct polynomial scalings; the exact form of which is irrelevant, but crucially all of them obey the finite size criteria as per definition 7.4.

7.2.4 Critical Parameter Problem Definitions

With the notion of phase transitions and Local-Global properties clarified within the last two sections, we can now define the Critical Parameter Problem (1-CRT-PRM) as follows.

Definition 7.5 (1-CRT-PRM$_f$)**.**

Input: *$N \in \mathbb{N}$. A constant-size set of k-local interaction terms $h^{(l)}(\varphi) \in \mathcal{B}((\mathbb{C}^d)^{\otimes S_l})$, for $l \in I$, and such that $S_l \subset \Lambda$, and $|S_l| \leq k$ $\forall l$. Two positive numbers $\alpha < \beta$ which satisfy $\beta - \alpha = \Omega(1/f(N))$. α, β and the matrix entries of each of the $\{h^{(l)}(\varphi)\}$ are specified to $|N| \coloneqq \lceil \log_2 N \rceil$ bits of precision.*

Promise: *Let $S + (i,j) \coloneqq \{(a+i, b+j) : (a,b) \in S\}$. Define the translationally-invariant Hamiltonian on Λ via*

$$H \coloneqq \sum_{l \in I} \sum_{(i,j) \in \Lambda} h_{S_l + (i,j)}(\varphi). \tag{7.1}$$

H has two phases A and B, and satisfies a Local-Global property as per definition 7.3 or definition 7.4, for some $L_0 = \text{poly}\, N$, independent of φ. There is precisely one critical parameter $\varphi^* \in [0,1] \setminus [\alpha, \beta]$ as per definition 7.1. If $\varphi < \varphi^*$ the system is in phase A, and for $\varphi > \varphi^*$ it is in phase B.

Output: YES *if $\varphi^* \leq \alpha$.*
NO *if $\varphi^* \geq \beta$.*

This problem characterises the hardness of estimating the point at which there is a phase transition (e.g. from gapped to gapless), for a one-parameter translationally-invariant Hamiltonian which is promised to have precisely one such transition; the input is the specification of the Hamiltonian to precision N, and the input size (given in, say, binary) is thus linear in $|N|$.

We kept the *precision* to which the critical point has to be estimated as a parameter, denoted by the function f subscript to the problem definition; in brief, 1-CRT-PRM$_{\text{poly}}$ denotes the case where we wish to approximate it to precision $\Omega(1/\text{poly})$, and we define the "precise" version of the problem to be Precise-1-CRT-PRM = 1-CRT-PRM$_{\exp\text{poly}}$. For ease of notation, we let 1-CRT-PRM = 1-CRT-PRM$_{\Theta(1)}$ be the case where the phase transition has to be estimated to constant precision.

We emphasise that while the subscript determines the precision to which we wish to compute the critical point, the Local-Global property as per definitions 7.3

and 7.4 is still required with polynomial precision throughout. Natural extensions of definition 7.5 where either scalings are given as parameters are of course possible.

A two-parameter variant of 1-CRT-PRM can be defined analogously as follows.

Definition 7.6 (2-CRT-PRM$_f$)**.**

Input: $N \in \mathbb{N}$. *A finite set of k-local interactions* $h^{(l)}(\theta,\varphi) \in \mathcal{B}((\mathbb{C}^d)^{\otimes S_l})$, *for* $l \in I$, *and such that* $S_l \subset \Lambda$, *and* $|S_l| \leq k\ \forall l$. *Four positive numbers* $\alpha_1 < \beta_1$ *and* $\alpha_2 < \beta_2$, *such that the rectangle* $[\alpha_1,\beta_1]\times[\alpha_2,\beta_2]$ *covers an* $\Omega(1/f(N))$ *area.* $\alpha_1,\beta_1,\alpha_2,\beta_2$ *and the matrix entries of each of the* $\{h^{(l)}(\theta,\varphi)\}$ *are specified to* $\mathrm{poly}(|N|)$ *bits of precision.*

Promise: H *is defined as in eq.* (7.1)*, and satisfies a Local-Global property for two phases A and B, as per definition 7.3 or definition 7.4 for* $L_0 = \mathrm{poly}\, N$*, independent of* θ *and* φ. *The critical line* $\theta^*(\varphi)$ *is a function of* φ*—i.e. for each* φ*, there exists exactly one critical parameter* θ^* *as per definition 7.1 such that* $H(\theta,\varphi)$ *is in phase A if* $\theta < \theta^*$*, and in phase B if* $\theta > \theta^*$. *The critical line* $\theta^*(\varphi)$ *is either such that the rectangle* $R \coloneqq [\alpha_1,\beta_1+y]\times[\alpha_2,\beta_2+y]$ *lies completely in phase A, or completely in phase B, where*

$$y \coloneqq \max\{\varphi^* : H(\theta,\varphi) \text{ is in phase } B\ \forall\varphi < \varphi^*, \forall\theta\}.$$

is promised to be well-defined, and that $H(0,\varphi)$ *is in phase A for all* $\varphi > y$*, and in phase B for all* $\varphi < y$.

Output: Yes *if the critical line is such that rectangle R lies in phase A.*
No *otherwise.* [1]

We emphasise that this definition is a direct analogue of the 1-CRT-PRM case, where the rectangle's role was taken by the one-dimensional interval $[\alpha,\beta]$; the offset y is necessary to obtain a well-defined problem definition, and is analogous to how APX-SIM (definition 6.1) is defined with respect to a "natural reference point",

[1] This order is switched compared to the 1-CRT-PRM case, matching the bounds for APX-SIM, whereas the bounds for 1-CRT-PRM match those of the Local Hamiltonian problem.

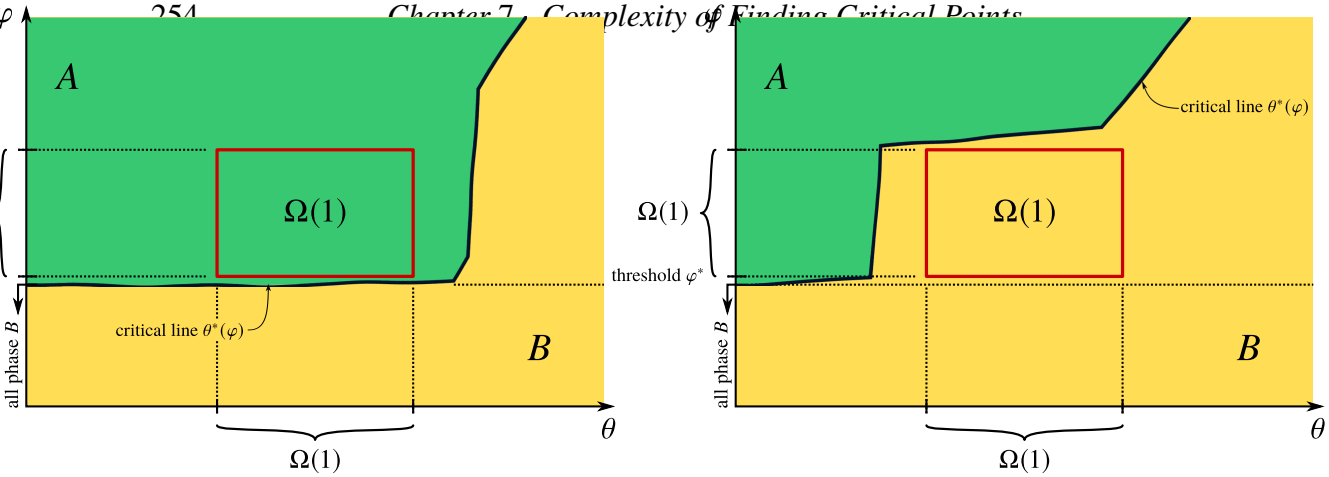

Figure 7.1: The two possible phase diagrams for the family of 2-parameter Hamiltonians we construct. The critical line $\varphi^*(\theta)$ is continuous, and there is a $\Omega(1)$-sized area which in the first case is guaranteed to be completely in phase A (e.g. a gapped phase), and in the second case completely in phase B (e.g. a gapless phase), for parameters $\theta \times \varphi \in [0,1] \times [0, \text{poly}\, N]$. We prove that determining which of the two cases holds is a $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$-complete problem.

i.e. the Hamiltonian's ground state energy. The "natural reference point" for phase diagrams we choose is simply a point along one of the parameter axes below which the system is completely in one of the two phases, as shown in fig. 7.1.

Other equivalently well-motivated definitions can of course be given. We give the following variant of 2-CRT-PRM, which reads more akin to the way 1-CRT-PRM is formulated, but which is otherwise identical in meaning to definition 7.6.

Definition 7.7 (2-CRT-PRM$_f$, alternative formulation)**.**

Input: *$N \in \mathbb{N}$. A finite set of k-local interactions $h^{(l)}(\theta,\varphi) \in \mathcal{B}((\mathbb{C}^d)^{\otimes S_l})$, for $l \in I$, and such that $S_l \subset \Lambda$, and $|S_l| \leq k$ $\forall l$. Positive numbers $\varphi, \alpha_1, \beta_1$, such that $\beta_1 - \alpha_1 = \Omega(1)$. $\alpha_1, \beta_1, \varphi$ and the matrix entries of each of the $\{h^{(l)}(\theta,\varphi)\}$ are specified to* $\text{poly}(|N|)$ *bits of precision.*

Promise: *H is defined as in eq. (7.1), and satisfies a Local-Global property for two phases A and B, as per definition 7.3 or definition 7.4 for $L_0 = \text{poly}\, N$, independent of θ and φ. The critical line $\theta^*(\varphi)$ is a function of φ—i.e. for each φ, there exists exactly one critical parameter θ^* as per definition 7.1 such that $H(\theta,\varphi)$*

is in phase A if $\theta < \theta^*$, and in phase B if $\theta > \theta^*$. It is promised that

$$y \coloneqq \max\{\varphi^* : H(\theta,\varphi) \text{ is in phase } B\ \forall \varphi < \varphi^*, \forall \theta\}$$

is well-defined, and that $H(0,\varphi)$ is in phase A for all $\varphi > y$, and in phase B otherwise. Furthermore, there is an interval $S_\kappa \coloneqq [y+\kappa, y+2\kappa]$ for $\kappa = \Omega(1/f(N))$, such that if $\varphi \in S_\kappa$, then for all $\varphi \in S_\kappa$ either $\theta^*(\varphi) > \beta_1$ or $\theta^*(\varphi) < \alpha_1$.

Output: YES $\theta^*(\varphi) > \beta_1$ *for all* $\varphi \in S_\kappa$.
NO $\theta^*(\varphi) < \alpha_1$ *for all* $\varphi \in S_\kappa$. [2]

While definitions 7.6 and 7.7 sound somewhat contrived, we emphasise that the core idea behind it is analogous to how APX-SIM is defined. The latter asks: if I take a low-energy state, what's the expectation value with respect to a given observable? It is natural, in this context, to imply the ground state within the problem definition, and not to demand it to be given as input in first place. In a similar fashion, to approximate a critical line in a phase diagram to some desired precision it is conceivable that one knows a region below which the phase diagram is entirely in one phase; and to draw the critical line from this reference point onward to some desired precision.

Just as the precision to which we wish to approximate the ground state energy dictates how *hard* a problem it will be, the Local-Global properties are in place to ensure we can prove *containment* of the problems, i.e. to place them within a complexity class that solely depends on how hard it is to solve the Local-Global property. Example variants—and the ones we will focus on in this paper—are when phase A and B are gapped vs. gapless states; the Local-Global property is then given by definition 7.4.

When proving complexity results, we will be interested in the class $\mathsf{QMA}_{\mathsf{EXP}}$, which is to QMA what NEXP is to NP; its use over QMA is a technicality based

[2] This order is switched compared to the 1-CRT-PRM case, matching the bounds for APX-SIM, whereas the bounds for 1-CRT-PRM match those of the Local Hamiltonian problem.

on how the input for a translationally-invariant system is specified (i.e., as a single interaction term repeated over the lattice).

7.2.5 APX-SIM

To prove the hardness result for the two-parameter case, we will prove a reduction to the hardness of simulating measurements on local Hamiltonians.

As we aim to establish the hardness of 2-CRT-PRM for translationally-invariant Hamiltonians, we will consider the following variant of APX-SIM discussed in chapter 6. We put it in a form more comparable to the definitions of CRT-PRM, but it is essentially the same as definition 6.1:

Definition 7.8 ($\forall$-TI-APX-SIM(H, A, a, b, δ) [GPY20])**.**

Input: *Local term h of a translationally invariant Hamiltonian $H = \sum h$, and a local observable A, and real numbers a, b such that $b - a \geq N^{-c'}$, for N the number of qubits H acts on and $c, c' > 0$ some constants.*

Promise: *Let S_δ be the set of all states $|\psi\rangle$ satisfying $\langle\psi| H |\psi\rangle \leq \lambda(H) + \delta$ for $\delta \geq N^{-c}$. For any $|\psi\rangle \in S_\delta$, we have either $\langle\psi| A |\psi\rangle \geq b$ or $\langle\psi| A |\psi\rangle \leq a$.*

Output: YES $\langle\psi| A |\psi\rangle \geq b$ *for all* $|\psi\rangle \in S_\delta$.
NO $\langle\psi| A |\psi\rangle \leq a$ *for all* $|\psi\rangle \in S_\delta$.

Thus the problem promises that all states within energy δ of the ground state (i.e. low-energy states) have similar expectation values when measured relative to A. (Note we have switched the yes and no instances from those defined in [GPY20], however, this choice is arbitrary and does not change any results). We will also make use of corollary 6.1.

7.2.6 The Gottesman-Irani Hamiltonian and the Local Hamiltonian Problem

We will make direct use of Gottesman and Irani's constructive hardness proof showing that the translationally-invariant Local Hamiltonian problem is $\mathsf{QMA}_{\mathsf{EXP}}$-complete,

in the sense of utilising the ability to encode quantum computation into the ground state of a Hamiltonian, and summarize their result in the following statement.

Theorem 7.3 (Gottesman and Irani [GI09])**.** *A Gottesman-Irani Hamiltonian* $G_N \coloneqq \sum_{i=1}^{N} h_i$ *is a translationally-invariant nearest neighbour Hamiltonian on a one-dimensional spin chain with finite local dimension* $d \in \mathbb{N}$*, and with open boundary conditions;* G_N *has the following properties:*

1. *For all* $N \geq 10$*, the ground state of* G_N *is a history state which encodes a binary counter with output* $N-2$ *in binary; and then takes this as input to a universal quantum Turing machine* $\mathcal{M}$*. Part of the input to* $\mathcal{M}$ *remains unconstrained.*

2. *If* N *describes a* QMA *verifier and a valid problem instance for it—as e.g. shown in [BCO17, Fig. 11]—then there exist two polynomials* p, q *such that* $1/p(N) - 1/q(N) = \Omega(1/\mathrm{poly}\, N)$*, and*

$$\lambda_{\min}(G_N) \begin{cases} \geq 1/p(N) & \text{if } \mathcal{M} \text{ outputs No} \\ \leq 1/q(N) & \text{if } \mathcal{M} \text{ outputs Yes.} \end{cases}$$

   Determining which case occurs is $\mathsf{QMA}_{\mathsf{EXP}}$*-complete.*

The existence and construction of G_N is a by-now standard technique; aside from Gottesman and Irani's original construction there also exists a lower local dimension variant with local dimension $d = 42$ [BCO17]. Other nondeterministic computations can be encoded in a similar fashion, and we will make use of this fact in section 7.6 (as seen in chapter 6).

7.3 Main Results

Now that we have introduced the necessary technical background, we can give rigorous statements of our two main results, theorems 7.1 and 7.2.

Theorem 7.4. 1*-CRT-PRM is* $\mathsf{QMA}_{\mathsf{EXP}}$*-hard and contained in* $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$ *for Hamiltonians satisfying either definition 7.3 or definition 7.4.*

In fact, we prove a slight bit more than this; we show that 1-CRT-PRM is $\mathsf{QMA}_{\mathsf{EXP}}$-hard, even if the parameter is only to be inferred to constant precision (as stated in the informal theorem 7.1).

Theorem 7.5. 2-*CRT-PRM is* $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$*-complete for Hamiltonians satisfying either definition 7.3 or definition 7.4.*

Containment within some complexity class for the two cases hinges, as aforementioned, on how hard it is to answer the respective local-global properties; the corresponding reductions are proven in section 7.4; for us, we will constrain the threshold system size for the local-global promises to be polynomial in the input size, which in both cases will result in a containment within $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$. Whether containment for the 1-CRT-PRM case can be made tighter (e.g. prove containment within $\mathsf{QMA}_{\mathsf{EXP}}$) is an open question (see the extended discussion in section 7.7). We prove the two hardness results in sections 7.5 and 7.6, respectively, by explicitly constructing translationally-invariant nearest neighbour families of Hamiltonians which encode the answer of a $\mathsf{QMA}_{\mathsf{EXP}}$ resp. $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$-hard problem within the decision problem of whether the system is in phase A or B.

Each Hamiltonian will be defined on a lattice Λ, $\{H_N(\varphi)\}_{N\in\mathbb{N}}$, where $H_N(\varphi) = \sum_{i\sim j} h_N^{(i,j)}(\varphi) + \sum_{i\in\Lambda} h_N^{(i)}(\varphi)$ for neighbouring sites $i \sim j$, $i, j \in \Lambda$. All local terms $h^{(i)} \coloneqq h^1_{\{i\}} \otimes \mathbb{1}_{\Lambda\setminus\{i\}}$, and analogously for $h^{(i,j)}$ constructed from a constant two-site matrix h^2. Furthermore, all the matrix entries of h^1 and h^2 will be specified to $O(|N|) = O(\log_2 N)$ bits of precision; and naturally we allow the local terms to depend on the family parameter N in a trivially-computable fashion (i.e. we demand the matrix entries have to be computable classically from N in time $\mathrm{poly}\,|N|$), where we do however require a constant-bounded interaction strength $\|h^1\|, \|h^2\| \leq 1$. As aforementioned, containment in $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$is a corollary of the local-global phase/gappedness promise imposed on the systems.

7.4 Containment of 1- and 2-CRT-PRM in $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$

Lemma 7.1. *Consider a Hamiltonian $H^{\Lambda(L)}(\varphi)$ such that the local terms are describable in n bits, and that satisfies the global-local* ***gap*** *condition (definition 7.4)*

. *Then for* $n = O(\log(L))$ *determining whether* $\Delta(H^{\Lambda(L)}(\varphi)) \geq 1/q(L)$ *or* $\leq 1/p(L)$ *for any point in parameter space is in* $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$.

Proof. The algorithm showing containment of SPECTRAL GAP (as defined in [Amb14]) in $\mathsf{P}^{\mathsf{QMA[log]}}$ can be used. However, now the Hamiltonian is promised to be translationally invariant, we need exponentially less information to input the Hamiltonian; the only input is now L which only requires $n = \mathrm{O}(\log(L))$ bits to express. As the required precision, on the other hand, is still $1/\mathrm{poly}\, L$, the relevant complexity class is now $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$. □

Lemma 7.2. *Consider a Hamiltonian* $H^{\Lambda(L)}(\varphi)$ *such that the local terms are describable in n bits, and that satisfies the global-local* ***phase*** *condition (definition 7.3). Then for* $n = O(\log(L))$ *and all states* $\langle\psi| H^{\Lambda(L)}(\varphi) |\psi\rangle \leq \lambda_{\min}(H^{\Lambda(L)}(\varphi)) + \omega$ *determining whether* $\langle\langle\psi| O_{A/B} |\psi\rangle\rangle \geq 1/q(L)$ *or* $\leq 1/p(L)$ *for any point in parameter space such that* $|\varphi - \varphi^*| \geq 1/\mathrm{poly}(L)$ *is in* $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$.

Proof. Follows directly from corollary 6.5 in the following way. Set δ from the definition of ∀-TI-APX-SIM in definition 7.8 to be equal to the energy parameter ω from definition 7.3. Let the order parameter $O_{A/B}$ be the local observable to be measured (i.e. A in the definition of ∀-TI-APX-SIM), and let the polynomials p, q correspond to the bounds a, b. Then finding $\langle\psi| O_{A/B} |\psi\rangle$ for a state with energy $\langle\psi| H^{\Lambda(L)}(\varphi) |\psi\rangle \leq \lambda_{\min}(H^{\Lambda(L)}) + \omega$ is just an instance of ∀-TI-APX-SIM. □

Theorem 7.6. *Let* H *be a translationally invariant Hamiltonian with local terms describable in n bits, and on a lattice of size* $L = \exp(\mathrm{poly}\, n))$. *Further let* H *satisfy either the global-local gap property in definition 7.4 or the global-local phase property in definition 7.3, and have a critical point at* φ^*. *Then the 1-CRT-PRM and 2-CRT-PRM are contained in* $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$.

Proof. We start with the more general case.

Local-Global Phase Assumption. For 1-CRT-PRM, we must show it is possible to find an approximation $\tilde{\varphi}^*$ such that for $n = \mathrm{O}(\log(L))$

$$|\tilde{\varphi}^* - \varphi^*| < \mathrm{O}(1/\mathrm{poly}\, L) = \mathrm{O}(1/\exp n).$$

with an algorithm in $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$.

The algorithm is as follows:

- Calculate L_0. By definition 7.3, this can be calculated in $\text{poly}(n)$ time.
- Take a $L_0 \times L_0$ region of the lattice. Using the algorithm in lemma 7.1, determine $\langle O_{A/B} \rangle = \langle \psi | H | \psi \rangle$ for states satisfying $\langle \psi | H | \psi \rangle \leq \lambda_{\min}(H) + \omega$ to precision $1/r(L_))$, where $r(L) \gg q(L) > p(L)$. As per lemma 7.2, this can be done using $\text{poly}(n)$ many calls to a $\mathsf{QMA}_{\mathsf{EXP}}$ oracle.
- If $\langle O_{A/B} \rangle < 1/p(L) + 1/r(L)$ then by the global-local phase condition definition 7.3, the Hamiltonian must be in phase A. If $\langle O_{A/B} \rangle > 1/q(L) - 1/r(L)$, then we know it must globally be in phase B. Due to the earlier promise, we are guaranteed that it satisfies one of these conditions.
- Perform a binary search through the parameter space of φ. Using $\text{O}(\text{poly}(n)) = \text{O}(\log(L))$ runs of the algorithm we can identify the point at which the order parameter from $\langle O_{A/B} \rangle < 1/p(L) + 1/r(L)$ to $\langle O_{A/B} \rangle > 1/q(L) - 1/r(L)$ or visa-versa to within $O(1/\text{poly}(L))$ precision. The interval in which the expectation $\langle O_{A/B} \rangle$ changes must contain the critical parameter φ^*.
- The result is we get an estimate $\tilde{\varphi}^*$ such that

$$|\tilde{\varphi}^* - \varphi^*| < \text{O}(1/\text{poly}\, L) = \text{O}(1/\exp n).$$

Finally we note that running the algorithm to compute the $\langle O_{A/B} \rangle$ takes $\text{poly}(n) = \text{O}(\log(L))$ $\mathsf{QMA}_{\mathsf{EXP}}$ queries at each point in parameter space. To do the binary search procedure, we choose $\text{poly}(n)$ points, each of which runs this algorithm. Thus, overall we make $\text{poly}(n) = \text{O}(\log(L))$ queries to the $\mathsf{QMA}_{\mathsf{EXP}}$ oracle and hence 1-CRT-PRM is in $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$.

For the 2-CRT-PRM case, the extra ingredient is the offset $y = \varphi^*$ along the $\theta = 0$ axis in definition 7.6. As we are promised that there is a unique such critical point, we can use binary search for the special case $H(0, \varphi)$ to approximate y to precision

$O(1/\exp n)$ which takes at most $\text{poly}(n)$ oracle queries. Using definition 7.7 as our definition of 2-CRT-PRM, we choose some $\varphi \in [y+\kappa, y+2\kappa]$ and query the order parameter for that point as above. This can all be done in $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$.

Spectral Gap Assumption. The proof for containment in the case that the Hamiltonian satisfies the global-local spectral gap condition is almost identical. This is because the algorithm to determine the spectral gap to precision $O(1/\exp n) = 1/\text{poly}\, L$ precision is also contained in $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$, as shown in lemma 7.1. □

7.5 $\mathsf{QMA}_{\mathsf{EXP}}$ Hardness of 1-CRT-PRM

In order to prove $\mathsf{QMA}_{\mathsf{EXP}}$-hardness of 1-CRT-PRM, we explicitly construct a 1-parameter family of Hamiltonians $H_N(\varphi)$ on a qudit lattice Λ, where $N \in \mathbb{N}$ and $\varphi \in [0,1]$ are encoded into phases of a local term, and with all matrix entries specified to bit precision of at most $|N|$.

Proof Outline. The local terms $h^N(\varphi)$ will be a function of two parameters: N and φ (both encoded into the phase of a local term). Here N encodes the problem instance and should be thought of as changing the form of the Hamiltonian. This point is subtle, but important: in the thermodynamic limit, where the lattice size is infinite, the only term that encodes the instance is the local coupling terms $h^N(\varphi)$. For every such h^N, we now ask: is the critical point φ^* above or below some threshold? This is entirely analogous to the local Hamiltonian problem, where say the size of the spin chain N determines the instance, and we ask whether the ground state is above or below some threshold. So what does N encode? It is an integer that encodes the hard problem that we reduce to the phase decision; as such, not all N might be valid (as the hard classes we use in the reduction are promise problems; as such, there might be invalid N too for which we cannot say anything). When making reductions between promise problems we need only show that if the promise of the initial problem is satisfied, then the promise of the problem we are mapping it to much also be satisfied. We will not be concerned with the case where N corresponds to an invalid instance.

In contrast to the LH Problem, or the question of whether a system is gapped or not in the thermodynamic limit, we *know* by construction that the system will be in

phase A or B for various choices of φ. This means φ is the parameter of interest: we ask whether the critical point φ^* is above or below some threshold—just as we could have asked for the spectral gap to be larger or smaller than some threshold, or the ground state energy for that manner—or the colour of the resulting material. This, in turn, means that *every instance, indexed by N, is itself a family of Hamiltonians, parameterised by φ, and 1-CRT-PRM asks questions about families of Hamiltonians.* The complexity scaling will then be in terms of $|N|$, i.e. the number of bits required to encode N, as the input size, or precision to which we describe the matrix elements of the local terms. This is entirely natural: we would expect the problem to be ever harder the more precise we specify the local terms; and the precision to which we want to resolve a critical point to be related to the precision to which the local Hamiltonian is specified as well.

Proof Outline. $H_N(\varphi)$ is constructed so that its ground state partitions the lattice into checker board grids of varying side length (motivated by the idea in [BCW19]). Within each square of the checker board there is a QTM-to-Hamiltonian mapping, which means that its ground state is a so-called "history state"; for the particular Turing machine we choose to encode, the ground state represents the following procedure:

1. Perform QPE to extract N from local terms.
2. Perform a phase comparator QPE on the unitary encoding φ and $\exp(\mathrm{i}tG_N)$, where G_N is a translationally-invariant local spin Hamiltonian with $\mathsf{QMA}_{\mathsf{EXP}}$-hard local Hamiltonian problem (on a spin chain of length N). This computation is performed with an unconstrained input state. Assuming said input state is an eigenstate of G_N with eigenvalue λ, the phase comparator QPE extracts the *difference* $\lambda - \varphi$ to bit precision $\sim |N|$.
3. If $\varphi < \lambda$ we set an output flag to $|0\rangle$. If $\varphi > \lambda$ we set it to $|1\rangle$.

We can then give an energy penalty to the $|0\rangle$ state of the flag qubit at the output of this computation, which ensures that the hitherto unconstrained input state to the phase comparator QPE assumes G_N's ground state; we can thus assume that

$\lambda = \lambda_{\min}(G_N)$. Adding an unconditional bonus on the square size that the history computation runs on (using a 2D Marker Hamiltonian), this combination of history state and Marker Hamiltonian results in a ground space energy

$$\lambda_{\min}(H_N(\varphi)) = \begin{cases} < 0 & \text{if } \lambda_{\min}(G_N) < \varphi, \text{ and} \\ \geq 0 & \text{otherwise.} \end{cases}$$

As G_N satisfies a promise gap, there will also be a promise gap on $\lambda_{\min}(H_N(\varphi))$ with respect to φ.

We will describe this construction outlined above in rigorous detail in the following sections; we start with the construction of the quantum Turing machine performing the listed operations in section 7.5.1, and under the assumption of having access to all necessary gates directly. In section 7.5.2, we do away with this assumption and replace the execution of $\exp(\mathrm{i}tG_N)$ with a variant based on performing Hamiltonian simulation; as well as using Solovay-Kitaev to replace all remaining gates by a fixed gate set.

We then embed the resulting quantum Turing machine into a history state Hamiltonian in section 7.5.3, combine this construction with a 2D Marker Tiling in section 7.5.4, lift the phase comparison ground state energy to a phase transition in section 7.5.5, prove that our construction has a unique phase transition in section 7.5.6, and show the reduction $\mathsf{QMA}_{\mathsf{EXP}} \longrightarrow$ 1-CRT-QTM in section 7.5.7. Finally, in section 7.5.8, we prove that this constructed family of Hamiltonians indeed satisfies the local-global gap property from definition 7.4.

7.5.1 A Phase Comparator Quantum Turing Machine

In this section we consider a multi-tape QTM which will take as input two different numbers $N \in \mathbb{N}$ and $\varphi \in [0,1]$[3] as well as an input state $|\nu\rangle \in (\mathbb{C}^d)^{\otimes N}$. Rather than straightforwardly inputting N, φ on the tape, we will give the QTM access to particular gates such that when it performs quantum phase estimation on these gates, the resulting string will be an encoding of N and φ.

[3]Some of the previous theorems are stated for $\varphi \in [0, \mathrm{poly}(N)]$. We will actually prove a result for $\varphi \in [0,1]$ first and then extend it to a larger interval.

As we require both numbers to be extracted from the complex phase of a unitary, we need to encode them into the fractional part of the phase in some fashion. To this end, we devise the following encoding map for N:

Definition 7.9 (Encoding)**.** *Let* $(\cdot\,,\cdot)$ *represent the string concatenation operation, let* k *be a fixed integer, and let* $n = |N|$*. Let a string* $w \in [4]^{2n} \equiv ([4]^n, [4]^n)$ *be* valid *if* $w \in ([2]^n, 2^{\times n})$*, and denote the set of all valid strings of length* $2n$ *as* V_{2n}*, where further* $V \coloneqq \bigcup_{n=1}^{\infty} V_{2n}$*. For* $N \in \mathbb{N}$*, we define*

$$\mathrm{enc} : \mathbb{N} \longrightarrow V \quad \textit{where} \quad \mathrm{enc}(N) = (N_1 N_2 \cdots N_n, \overbrace{22\cdots2}^{n \text{ times}})$$

where N *has binary expansion* $N = N_1 N_2 \cdots N_n$*. For* $z \in \mathbb{N}$*, we set* $\mathrm{enc}^{-1}(z)$ *to be the number* z *truncated to the first half (rounded down) of the base-4 digits of* z*.*

We remark that the encoding is in base 4 and hence the number of **bits** required for $\mathrm{enc}(N)$ is $|\mathrm{enc}(N)| = 4|N|$. We note that for the inverse map in definition 7.9, we have that $\mathrm{enc}^{-1}\mathrm{enc}(N) = N$ for all $N \in \mathbb{N}$, and such that if z has m base-4 digits (i.e. $m \leq 4^m$), it always holds that $\mathrm{enc}^{-1}(z) \leq 2^m$.

In this section we loosely speak of performing QPE to base four; this is of course simply a shorthand for performing base 2 QPE with twice the number of bits; it is straightforward to verify that the following derivation is well-defined in this context. In order to assess how well-suited said encoding is to get a precise handle on the QPE error, we formulate the following technical lemma:

Lemma 7.3 (Encoded QPE Extraction)**.** *Denote with* V *the set of all valid strings from definition 7.9. Consider the map*

$$V \longrightarrow \mathbb{R} \quad \textit{where} \quad V \ni y \longmapsto 0.y \quad \textit{(in base 4)}$$

and set $U_y \coloneqq \mathrm{diag}(\exp(\mathrm{i}\pi 0.y), 1)$*. Denote with* $|y|$ *the length of* y *in base 4. Let the output of performing quantum phase estimation on the unitary* U_y *(wrt. to the eigenstate* $|0\rangle$*) with a perfect gate set (i.e. all gates are performed without error) on* $t \in 2\mathbb{N}$*,* $t \geq 2$ *qudits be* $\sum_{m\in[4]^t} \beta_m |m\rangle$*. Then if* $t \geq |y|$*,* $\beta_y = 1$ *and all other* $\beta_i = 0$*; if*

$t < |y|$ the total probability amplitude on valid strings is bounded as

$$\sum_{m \in V_t} |\beta_m|^2 < \frac{1}{2^{t/2}}.$$

Proof. We consider QPE on t qudits for two cases: $t < |y|$, $t \geq |y|$. The fact that we work with base 4 numbers bears no significance since we indirectly treat the setup as a base-2 QPE of twice the length, and it is useful to view the following proof through this lens.

Case t ≥ |y|. In this case the quantum phase estimation can be done exactly; this means the probability amplitude on states which are not y is precisely zero [NC10, Sec. 5.2].

Case t < |y|. In this case the QPE is not performed exactly and the output is some superposition clustered around the best t digit estimates of y; the following analysis closely follows the QPE error analysis in [NC10, Sec. 5.2]. As QPE is done in little Endian order, if we denote with $b \in [4]^t$ the string such that $b/4^t$ is the best t digit approximation to y *less* than y, we know that b is simply y truncated on the right hand side to t digits. We note that $b \notin V_t$ as the truncation means that $b_{t/2} \neq 2$ (remember that we assumed t even, so $t/2 \in \mathbb{N}$).

Denote with $b' \in V_t$ the closest *valid* string to y, in the same sense as b (i.e. such that $b'/4^t$ is closest to y amongst all $b' \in V_t$). Since $b'_{t/2} = 2$, we clearly have $|b - b'| \geq 4^{t/2}$, as the two strings have to differ by at least 1 at the $(t/2)^{\text{th}}$ position.

By [NC10, Eq. 5.27] we see that the probability of measuring an outcome m further than $e \in \mathbb{N}$ away from b is bounded by

$$p(|b - m| > e) \leq \frac{1}{2(e-1)}. \tag{7.2}$$

As b' was the closest valid string to y, we know that all other valid strings $m \in V_t$ also satisfy $|b - m| \geq 4^{t/2}$. This means

$$\sum_{m \in V_t} |b_m|^2 \leq \sum_{m \in V_t} \frac{1}{2(4^{t/2}-1)} \leq \frac{2^{t/2}}{2(4^{t/2}-1)} \leq \frac{1}{2^{t/2}}$$

as $|V_t| = 2^{t/2}$ by definition 7.9. The claim follows. □

QPE extracts the phase λ of some unitary U with respect to an eigenstate $|u\rangle$—i.e. such that $U|u\rangle = e^{i\lambda}|u\rangle$. This means the algorithm assumes that there exists a "black box" capable of preparing the register $|u\rangle$ to be the correct eigenstate. More often than not we do not have such a state: obtaining an eigenstate for an operator is generally at least as hard as estimating the associated eigenvalue.[4] On the other hand, this allows us to form a nondeterministic variant of QPE, in the sense that if we leave $|u\rangle$ unconstrained, we can ask questions like "does there exist a state $|u\rangle$ for which λ_u is less than a certain quantity?" This notion of a nondeterministic QPE has been employed in different contexts before, e.g. in the context of Hamiltonian simulation [Koh+20].

In the following lemma we will thus assume that the eigenstate $|u\rangle$ is an external quantity to be supplied to the procedure, where we keep in mind that we translate the algorithm to a history state Hamiltonian in due course. As history state Hamiltonians allow for an unconstrained section (which can later be filtered by a suitable penalty addition), the above decision problem of existence of a state $|u\rangle$ for which λ_u is below or above some threshold then maps naturally to the eigenspectrum of the Hamiltonian. More concretely, as the particular unitary we are interested in performing nondeterministic QPE on is itself a local Hamiltonian with a hard ground state energy problem, the notion of existence/non-existence of an eigenstate $|u\rangle$ with a phase λ_u below a certain threshold can be analysed precisely as in the case of encoding a (nondeterministic) QMA verifier in hardness proofs of the local Hamiltonian problem [KSV02].

In the remainder of this section, we will leave the number of bits to which QPE is performed (generally called t in the following) vs. the length of the chain on which the Hamiltonian-to-be-analysed sits (generally called N, or an integer $z \leq N$ depending on the context) independent; yet in order to prove hardness, we will later on require the number of bits large enough to resolve the promise gap of the encoded Hamiltonian to high enough precision.

[4] Generally, it is as "cheap" to calculate $\langle u|U|u\rangle$ as it is to write out $|u\rangle$, modulo polynomial overhead.

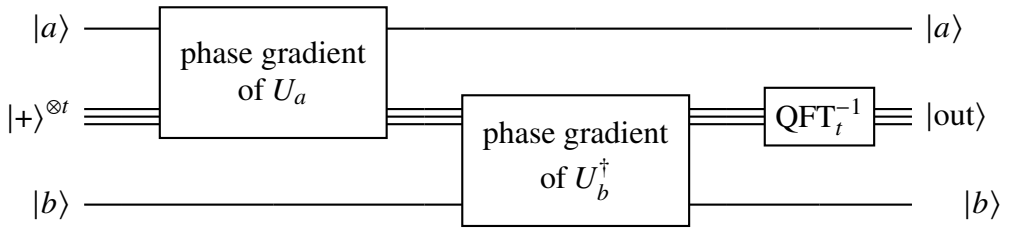

Figure 7.2: Phase comparator circuit. For two unitaries U_a and U_b with eigenstates $|a\rangle$, $|b\rangle$ and associated eigenvalues λ_a, λ_b, the output register $|\mathrm{out}\rangle$ contains a t-bit approximation to the phase $\lambda_a - \lambda_b$, as can be seen by writing out the phase gradient operations as given in [NC10, Fig. 5.2].

In order to compare two phases extracted via QPE, there is two options: extract each phase individually and perform a binary comparison, or perform QPE on the first unitary and the second unitary's inverse and compare against 0. We opt for the latter, as it will be easier to prove that a *single* critical point exists. Details in due course.

We now describe the QTM we will utilise for the rest of the 1-CRT-PRM hardness proof.

Lemma 7.4 (Multi-QPE QTM)**.** *Let G_z be a Gottesman-Irani Hamiltonian on a chain of length z, defined in theorem 7.3. Let $N \in \mathbb{N}$. Denote by $U_\varphi, U_N \in SU(2)$ and U_{G_z} the unitaries*

$$U_{G_z} = e^{\mathrm{i}\pi G_z} \quad U_\varphi = \begin{pmatrix} e^{\mathrm{i}\pi\varphi} & 0 \\ 0 & 1 \end{pmatrix} \quad U_N = \begin{pmatrix} e^{\mathrm{i}\pi 0.\,\mathrm{enc}(N)} & 0 \\ 0 & 1 \end{pmatrix}$$

where $\mathrm{enc}(N)$ *is given in definition 7.9. Let $|v\rangle \in (\mathbb{C}^d)^{\otimes N}$.*

Then there exists a quantum Turing Machine, denoted $\mathcal{M}(N, \varphi, t, |v\rangle)$, with access to the unitary gates U_φ, U_N, and U_{G_z} for all $z \in [N]$, all powers $2, 4, \ldots, 2^t$ of the U_{G_z} gates, as well as $R_k := \mathrm{diag}(1, e^{2\pi\mathrm{i}/2^k})$ for all $k = 1, \ldots, t$ in addition to the standard gate set; such that $\mathcal{M}$ acts on a Hilbert space of $2+4t+t$ qubits and N qudits, plus a slack space of size at most $\mathrm{poly}\, t$ *(left implicit in the following); and such that $\mathcal{M}$ performs the following operations:*

1. *Initialise the first $2+4t+t$ registers to zero, and assume the last N qudit registers are in state $|v\rangle$.*

2. *Execute QPE on U_N on $4t$ qubits to get a state*

$$|\chi'\rangle = |00\rangle_f \otimes \left(\sum_{z\in[4]^t} \gamma_z |z\rangle \right) \otimes |0\rangle^{\otimes t} \otimes |\nu\rangle$$

where the γ_z are the amplitudes from quantum phase estimation of U_N.

3. *For any basis state $|z\rangle$ in $|\chi'\rangle$ that is invalid, $\mathcal{M}$ places a marker on the first qubit of the tape (the flag space, labelled with subscript f), such that the first qubit is flipped to 1 if $z \in V_t$. This gives*

$$|\chi''\rangle = \left(\sum_{z\in V_t} |10\rangle_f \gamma_z |z\rangle + \sum_{z\notin V_t} |00\rangle_f \gamma_z |z\rangle \right) \otimes |0\rangle^{\otimes t} \otimes |\nu\rangle$$

4. *Let $z' := \min\{N, \mathrm{enc}^{-1}(z)\}$ as per definition 7.9, such that always $z' \leq 2^t$. On $|\nu\rangle$, the QTM performs a phase comparator QPE as shown in fig. 7.2 with the two unitaries $U_{G_{z'}}$ (on the $|\nu\rangle$ register) and $U_\varphi^\dagger$ (on a $|0\rangle$ ancilla register). More concretely, for the Hamiltonian $G_{z'}$, let its eigenstates be $\{|g_{z'}\rangle\}$ and let $|\nu\rangle = \sum_g \kappa_g(z') |g_{z'}\rangle \left|\xi_{z',g}\right\rangle$ be a decomposition with respect to a bipartition into z' and $N - z'$ qubits; as we chose the basis of the first subsystem, the $\left|\xi_{z',g}\right\rangle$ are not necessarily orthogonal, but can be assumed normalised. Then the input to this step can be written as*

$$|\chi''\rangle = \sum_{z\in[4]^t} \gamma_z |[z \in V_t]0\rangle_f |z\rangle \otimes |0\rangle^{\otimes t} \otimes \sum_g \kappa_g(z') |g_{z'}\rangle \left|\xi_{z',g}\right\rangle,$$

where $[z \in V_t]$ is equal to 1 iff $z \in V_t$ and is otherwise 0 (cf. section 7.2.1 for notation). The output of the total QTM after this stage is then

$$|\chi'''\rangle = \sum_{z\in[4]^t} \gamma_z |[z \in V_t]0\rangle_f |z\rangle \sum_{x\in[2]^t} \sum_g \alpha_x(z', g)\kappa_g(z') |x\rangle |g_{z'}\rangle |\xi_{z',g}\rangle$$

The $\alpha_x(z', g)$ are the coefficients obtained from the comparator QPE routine for operator G_z and U_φ on eigenstate $|g_{z'}\rangle$.

5. *The flag qubit is updated via* $|b0\rangle_f \longmapsto |b[x \leq 0]]\rangle_f$*, which corresponds to the comparison* $\lambda_{\min}(G_z) \leq \varphi$ *to* t *digits of precision. The resulting state is*

$$\begin{aligned}|\chi\rangle = &\sum_{z\in V_t}\sum_{x\leq 0}\sum_{g}\alpha_x(z',g)\gamma_z\kappa_g(z')\,|11\rangle_f\,|z\rangle\,|x\rangle\,|g_{z'}\rangle\left|\xi_{z',g}\right\rangle + \\ &\sum_{z\in V_t}\sum_{x>0}\sum_{g}\alpha_x(z',g)\gamma_z\kappa_g(z')\,|10\rangle_f\,|z\rangle\,|x\rangle\,|g_{z'}\rangle\left|\xi_{z',g}\right\rangle + \\ &\sum_{z\notin V_t}\sum_{x\leq 0}\sum_{g}\alpha_x(z',g)\gamma_z\kappa_g(z')\,|01\rangle_f\,|z\rangle\,|x\rangle\,|g_{z'}\rangle\left|\xi_{z',g}\right\rangle + \\ &\sum_{z\notin V_t}\sum_{x>0}\sum_{g}\alpha_x(z',g)\gamma_z\kappa_g(z')\,|00\rangle_f\,|z\rangle\,|x\rangle\,|g_{z'}\rangle\left|\xi_{z',g}\right\rangle.\end{aligned} \tag{7.3}$$

The QTM $\mathcal{M}$ *runs for time* $T = \mathrm{O}(2^{4t})$.

Proof. This QTM can be implemented by dovetailing a set of QTMs which perform QPE on $4t$ qubits for U_N, as well as t qubits for $U_{G_{z'}}$ and $U_\varphi^\dagger$, where the last one implementing the conditional QPE for $G_{z'}$ can be trivially implemented by adding additional control lanes. QPE without gate approximation to t bits of precision normally takes $\mathrm{O}(t^2)$ calls to $U_{G_{z'}}$ (which we assumed to have access to as a single gate for now) [NC10, Sec. 5.2]. Implementing the up to 2^{4t}-powers of the three phase gates U_N, $U_{G_{z'}}$ and $U_\varphi^\dagger$ takes time $\propto 2^{4t}$. All other operations take time $\operatorname{poly} t$, and it is clear that the QTM does not need in excess of $\operatorname{poly} t$ of slack work space. Hence we have an overall runtime of $\mathrm{O}(2^{4t})$. □

For the next set of lemmas we define the following quantity: for the output state $|\chi\rangle$ of $\mathcal{M}(N,\varphi,t)$ from lemma 7.4, we set

$$\begin{aligned}\eta(N,\varphi,t,|\nu\rangle) &:= \operatorname{Tr}\Big((|11\rangle\langle 11|_f \otimes \mathbb{1})\,|\chi\rangle\langle\chi|\Big) \\ &= \sum_{z\in V_t}\sum_{x\leq 0}\sum_{g}|\alpha_x(z',g)|^2|\gamma_z|^2|\kappa_g(z')|^2\end{aligned} \tag{7.4}$$

where as in lemma 7.4 we have $z' = \min\{N, \mathrm{enc}^{-1}(z)\}$. This is the total probability that $|\chi\rangle$ will have an accepted flag—i.e. the first qubit will be in the $|11\rangle_f$ state as given in eq. (7.3). It will later be shown that when the above QTM is encoded in a circuit-to-Hamiltonian mapping, then the ground state energy depends on η.

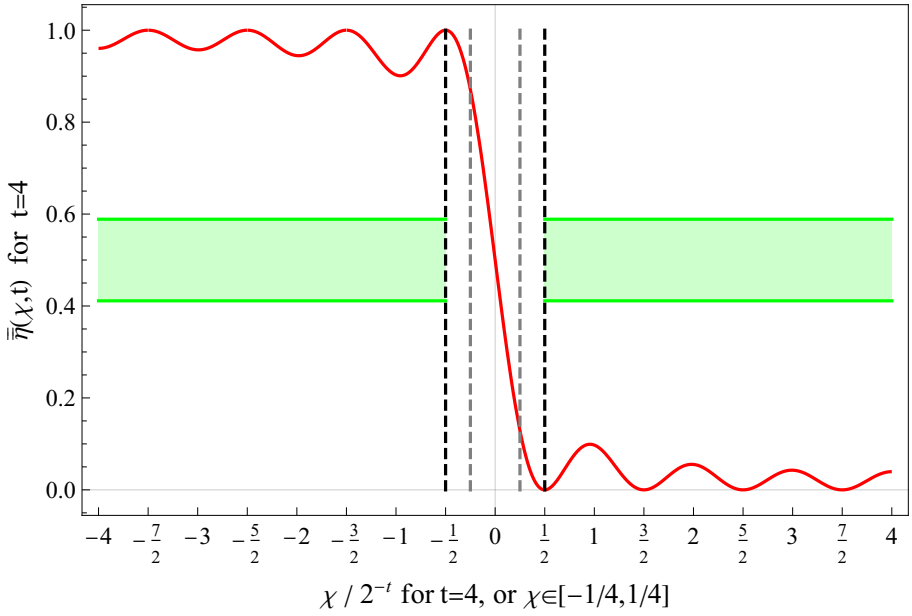

Figure 7.3: $\bar{\bar{\eta}}(\chi,t)$ (red line) for $t=4$ from lemma 7.5, vs. $\chi \in (-1/4,1/4)$. The interval between the black dashed vertical lines denote the region within which we prove $\bar{\bar{\eta}}$ to be strictly monotonously falling, and hence η from eq. (7.4) to be strictly monotonously increasing; within the grey dashed areas the slope of the red line is ≥ 1, as shown in lemma 7.5. The shaded green region marks the interval $[1/3,2/3]$, which lemma 7.5 proves $\bar{\bar{\eta}}(\chi,t)$ to be bounded away from.

As a first step, we present a monotonicity argument for the value of η around the point where φ equals the eigenvalue associated to an eigenstate $|\nu\rangle$ of G_N.

Lemma 7.5. *Let $|\nu\rangle$ be an eigenvector of G_N with eigenvalue λ, $t \geq |N|$, and let $\eta(N,\varphi,t,|\nu\rangle)$ be defined as in eq. (7.4). Then for $t \geq |N|$ and $\varphi \in (\lambda - 2^{-t} + 2^{-3t/2}, \lambda - 2^{-3t/2})$*

$$\frac{\partial \eta(N,\varphi,t,|\nu\rangle)}{\partial \varphi} \geq 1.$$

Furthermore, for all $\varphi < \lambda - 2^{-t} + 2^{-3t/2}$, $\eta \leq \pi^2/24$ and $\varphi > \lambda - 2^{-3t/2}$, $\eta \geq 1 - \pi^2/24$.

Proof. We guide the reader to fig. 7.3 to aid in an intuitive understanding of the proof. As a first step, we abbreviate $\xi := \lambda - \varphi$. By [NC10, Eq. 5.26] and eq. (7.4), and relabelling $L = x$ to follow the notation in [NC10] closely, we can write

$$\eta(N,\varphi,t,|\nu\rangle) = \sum_{z\in V_t} \sum_{x\leq 0} \sum_{g} |\alpha_x(z',g)|^2 |\gamma_z|^2 |\kappa_g(z')|^2 \overset{*}{=} \sum_{L=-2^{t-1}+1}^{0} |\alpha_L|^2$$
$$\text{for} \quad \alpha_L := 2^{-t} \frac{1-\exp\left(2\pi i(2^t\xi - L)\right)}{1-\exp(2\pi i(\xi - 2^{-t}L))} = 2^{-t} \frac{\sin\left(2^t\pi\xi\right)}{\sin(\pi(\xi - 2^{-t}L))} \tag{7.5}$$

where in the step marked with $\stackrel{*}{=}$ we have used the fact that for $t \geq |N|$, precisely one of the γ_z and $\kappa_g(z')$ equal 1, and all others are zero; so $\alpha_L \equiv \alpha_x(z', g)$ for those z and g for which $\gamma_z = \kappa_g(z') = 1$. We set $\bar{\eta}(\xi, t) \coloneqq \eta(N, \varphi, t, |\nu\rangle)$. As

$$\bar{\eta}(\xi, t) = \sum_{L \leq 0} |\alpha_L|^2 = 1 - \sum_{L > 0} |\alpha_L|^2, \tag{7.6}$$

we can calculate the midpoint where $\bar{\eta}(\xi, t) = 1/2$; this happens at $\xi = 2^{-t-1}$, as can be confirmed by explicit calculation.

Set $\chi \coloneqq \xi - 2^{-t-1}$ and define $\bar{\bar{\eta}}(\chi, t) \coloneqq \bar{\eta}(\chi + 2^{-t-1}, t)$, such that $\bar{\bar{\eta}}(0, t) = 1/2$, and

$$\begin{aligned}
\bar{\bar{\eta}}(-\chi, t) = \bar{\eta}(-\chi + 2^{-t-1}, t) &= \sum_{L=-2^{t-1}+1}^{0} 2^{-2t} \frac{\sin^2(2^t \pi(-\chi + 2^{-t-1}))}{\sin^2(\pi((-\chi + 2^{-t-1}) - 2^{-t} L))} \\
&\stackrel{*}{=} \sum_{L=-2^{t-1}+1}^{0} 2^{-2t} \frac{\sin^2(2^t \pi(-\chi - 2^{-t-1} + 2^{-t}))}{\sin^2(\pi(-\chi - 2^{-t-1} - 2^{-t}(L-1)))} \\
&= \sum_{L=-2^{t-1}+1}^{0} 2^{-2t} \frac{\sin^2(2^t \pi(\chi + 2^{-t-1}) - \pi)}{\sin^2(\pi(\chi + 2^{-t-1} - 2^{-t}(1-L)))} \\
&= \sum_{L=1}^{2^{t-1}} 2^{-2t} \frac{\sin^2(2^t \pi(\chi + 2^{-t-1}))}{\sin^2(\pi(\chi + 2^{-t-1} - 2^{-t} L))} \\
&\stackrel{**}{=} 1 - \bar{\bar{\eta}}(\chi, t),
\end{aligned} \tag{7.7}$$

where in the line with $\stackrel{*}{=}$ we added $2^{-t-1} - 2^{-t-1} = 0$ in the enumerator and denominator, and in the last step $\stackrel{**}{=}$ we made use of eq. (7.6).

Gradient Bound: Note that we can now write:

$$|\alpha_L|^2 = 2^{-2t} \frac{\cos^2(2^t \pi \chi)}{\sin^2(\pi(\chi - 2^{-t-1}(2L-1)))}.$$

From the above we see that:

$$\begin{aligned}\bar{\bar{\eta}}(\chi,t) &= \sum_{L=-2^{t-1}+1}^{0} 2^{-2t}\frac{\cos^2(2^t\pi\chi)}{\sin^2(\pi(\chi-2^{-t-1}(2L-1)))} \\ &= \sum_{L=0}^{2^{t-1}-1} 2^{-2t}\frac{\cos^2(2^t\pi\chi)}{\sin^2(\pi(\chi+2^{-t-1}(2L+1)))} = \sum_{L=0}^{2^{t-1}-1}|\alpha_{-L}|^2,\end{aligned}$$

where we have just relabelled $L \to -L$. Differentiating this expression wrt. χ gives

$$\begin{aligned}2^{2t}\frac{\partial\bar{\bar{\eta}}(\chi,t)}{\partial\chi} = \sum_{L=0}^{2^{t-1}-1}\Bigg[&-\frac{2^{t+1}\pi\sin\left(2^t\pi\chi\right)\cos\left(2^t\pi\chi\right)}{\sin^2(\pi(\chi+2^{-t-1}(2L+1)))} \\ &-2\pi\cos\Big(\pi(\chi+2^{-t-1}(2L+1))\Big)\frac{\cos^2(2^t\pi\chi)}{\sin^3(\pi(\chi+2^{-t-1}(2L+1)))}\Bigg].\end{aligned}$$

Now note that for $t \geq 1$, $\chi \in (0, 2^{-t-1})$ and $0 \leq L \leq 2^{t-1}-1$ we have that $2^t\pi\chi \leq \pi/2$ and $0 \leq \pi(\chi+2^{-t-1}(2L+1)) \leq \pi/2$. Thus all of the sine and cosine terms in the above expression are individually positive, and hence both the terms in the summand are individually negative, for all $0 \leq L \leq 2^{t-1}-1$.

We now focus on the 0^{th} coefficient, i.e. α_0. For $0 \leq x \leq \pi/2$ the following holds: $1/\sin(x) \geq 1/x$, thus giving:

$$2^{2t}\left|\frac{\partial|\alpha_0|^2}{\partial\chi}\right| \geq \frac{2^{t+1}\pi\sin\left(2^t\pi\chi\right)\cos\left(2^t\pi\chi\right)}{(\pi(\chi+2^{-t-1}))^2} \tag{7.8}$$

$$+\frac{2\pi\cos^2(2^t\pi\chi)}{|\pi(\chi+2^{-t-1})|^3}\cos\Big(\pi(\chi+2^{-t-1})\Big) \tag{7.9}$$

Now we consider the interval $\chi \in (0, 2^{-t-1}-2^{-3t/2})$. Within this interval, and for $t \geq 1$, $\cos\left(2^t\pi\chi\right)$ has its minimum value at the rightmost limit, $\cos^2(\pi/2-\pi2^{-t/2}) = \sin^2(\pi2^{-t/2}) \geq 2^{-t}$. Similarly, we have that

$$2^{-3t/2}\pi \leq \pi\left|\chi+2^{-t-1}\right| \leq 2^{-t-1}\pi,$$

and hence $\cos\big(\pi(\chi+2^{-t-1})\big) \geq \cos(\pi/4) \geq 1/\sqrt{2}$. Together with eq. (7.9) and dropping the term with $\sin\left(2^t\pi\chi\right)$ in the denominator (which vanishes for $\chi \to 0$), the 0^{th} coefficient α_0 thus satisfies

$$2^{2t}\left|\frac{\partial|\alpha_0|^2}{\partial\chi}\right| \geq \frac{2\pi 2^{-t}}{(\pi 2^{-t-1})^3} \times \frac{1}{\sqrt{2}} = \frac{8\sqrt{2}}{\pi^2} \times 2^{2t} \geq 2^{2t}, \tag{7.10}$$

Hence $\partial|\alpha_0|^2/\partial\chi \leq -1$. Since $\partial|\alpha_L|^2/\partial\chi < 0$ for all $0 \leq L \leq 2^{t-1}-1$, then $\partial|\alpha_0|^2/\partial\chi > \partial\bar{\bar{\eta}}/\partial\chi$ for $\chi \in (0, 2^{-t-1} - 2^{-3t/2})$.

Thus

$$\frac{\partial\bar{\bar{\eta}}(\chi,t)}{\partial\chi} \leq -1.$$

Using the antisymmetry of the function $\bar{\bar{\eta}}(\chi,t) - 1/2$ around the point $\chi = 0$, we see the same bounds on the derivative hold for the whole interval $\chi \in (-2^{-t-1} + 2^{-3t/2}, 2^{t-1} - 2^{-3t/2})$. Finally, noting that $\partial\chi/\partial\varphi = -1$

$$\frac{\partial\eta(N,t,\varphi,|\nu\rangle)}{\partial\varphi} \geq 1,$$

for the interval $\varphi \in (\lambda - 2^{-t} + 2^{-3t/2}, \lambda - 2^{-3t/2})$.

Bounds Outside Interval. To address the bounds when χ is outside of the monotonicity interval $\chi \in [-1/4, 1/4] \setminus (-2^{-t}/2 + 2^{-3t/2}, 2^{-t}/2 - 2^{-3t/2})$, we again by eq. (7.7) we only have to consider the right half of the interval. There we have

$$\begin{aligned}
\bar{\bar{\eta}}(\chi,t) = \sum_{L\leq 0}|\alpha_L|^2 &\overset{*}{=} \sum_{L=-2^{t-1}+1}^{0} 2^{-2t}\frac{\cos^2(2^t\pi\chi)}{\sin^2(\pi(\chi + 2^{-t-1} - 2^{-t}(L+1)))} \\
&\leq \sum_{L=0}^{2^{t-1}-1}\frac{2^{-2t}}{\sin^2(\pi(\chi + 2^{-t-1} + 2^{-t}(L-1)))} \\
&= \sum_{L=1}^{2^{t-1}}\frac{2^{-2t}}{\sin^2(\pi(\chi + 2^{-t-1} + 2^{-t}L))}.
\end{aligned} \tag{7.11}$$

where in the step $\overset{*}{=}$ we again added and subtracted 2^{-t-1} in the denominator. For $\chi \in [2^{-t-1} - 2^{-3t/2}, 1/4]$ and $L = 1, \ldots, 2^{t-1}$, we can bound

$$\pi(\chi + 2^{-t-1} + 2^{-t}L) \geq \pi 2^{-t}(L+1) - \pi 2^{-3t/2}$$

and thus

$$\sin\Big(\pi(\chi+2^{-t-1}+2^{-t}L)\Big) \geq \sin\Big(\pi(2^{-t}(L+1)-2^{-3t/2})\Big) \geq \frac{1}{2}\times\pi(2^{-t}(L+1)-2^{-3t/2}). \tag{7.12}$$

Combining eqs. (7.11) and (7.12), we get

$$\begin{aligned}
\bar{\bar{\eta}}(\chi,t) \leq & \frac{4}{\pi^2}\sum_{L=1}^{2^{t-1}}\frac{1}{(L+1-2^{-t/2})^2} \\
\leq & \frac{4}{\pi^2}\sum_{L=2}^{\infty}\frac{1}{L^2}\left(1-\frac{2^{-t/2}}{L}\right)^{-2} \\
= & \frac{4}{\pi^2}\sum_{L=2}^{\infty}\frac{1}{L^2}\left(1+\frac{2\times 2^{-t/2}}{L}+\mathrm{O}\left(\frac{2^{-t}}{L^2}\right)\right) && (7.13)\\
= & \frac{4}{\pi^2}\sum_{L=2}^{\infty}\frac{1}{L^2}+\frac{8}{\pi^2}\sum_{L=2}^{\infty}\frac{2^{-t/2}}{L^3}+\mathrm{O}\left(2^{-t}\right) && (7.14)\\
= & \frac{4}{\pi^2}\times\left(\frac{\pi^2}{6}-1\right)+\mathrm{O}(2^{-t/2}) && (7.15)\\
\leq & \frac{\pi^2}{24}.
\end{aligned}$$

Here for eq. (7.13) we have used a binomial expansion and for eq. (7.15) we have used the well known identity $\sum_{n=1}^{\infty} n^{-2} = \pi^2/6$. The bound for negative χ follows by eq. (7.7). □

Lemma 7.5 puts bounds on η for a specific input state $|v\rangle$. Here we consider the maximum value η can take: this corresponds to the maximum acceptance probability that $\mathcal{M}$ can have when the input state $|v\rangle$ is unconstrained.

Corollary 7.1. *Let η be as defined in eq. (7.4), and $|N|$ be the number of base-2 digits[5] of N.*

- *If $t \geq |N|$ and if $\varphi \leq \lambda_{\min}(G_N) - 2^{-t} + 2^{-3t/2}$, we have*

$$\max_{|v\rangle}\eta(N,\varphi,t,|v\rangle) \leq \frac{\pi^2}{24}.$$

[5] N is expressed in binary, but $\mathrm{enc}(N)$ in quaternary, hence $|\mathrm{enc}(N)| = 4|N|$; note that we expand to $4t$ bits in lemma 7.4, hence the statement "$t \geq |N|$" implies we expanded enough digits to see all of $\mathrm{enc}(N)$ exactly.

- *If $t \geq |N|$ and $\varphi \geq \lambda_{\min}(G_N) - 2^{-3t/2}$, we get*

$$\max_{|\nu\rangle} \eta(N, \varphi, t, |\nu\rangle) \geq 1 - \frac{\pi^2}{24}.$$

- *If $t < |N|$, then irrespective of the value of φ,*

$$\max_{|\nu\rangle} \eta(N, \varphi, t, |\nu\rangle) = \mathrm{O}\left(\frac{1}{2^{t/2}}\right).$$

Proof. We address each case individually.

Case t < |N|. We do not expand enough bits to expand $\mathrm{enc}(N)$ in full: by lemma 7.3, the probability mass on valid strings (none of which are $\mathrm{enc}(N)$) is $\leq 1/2^{t/2}$.

Case t ≥ |N|. Again by lemma 7.3 we know that $\gamma_{\mathrm{enc}(N)} = 1$ and all other $\gamma_z = 0$. If $\varphi \geq \lambda_{\min}(G_N) - 2^{-3t/2}$, then choose input state $|\psi_0\rangle = |g_{\min}\rangle |\xi_0\rangle$ where $|g_{\min}\rangle$ is the ground state of G_N (and $|\xi_0\rangle$ is just the state resulting from the bipartition of $|\psi_0\rangle$'s state space in the proof of lemma 7.4). By applying lemma 7.5, we see that

$$\eta(N, \varphi, t, |\psi_0\rangle) \geq 1 - \frac{\pi^2}{24}.$$

The result for $\varphi \geq \lambda_{\min}(G_N) - 2^{-3t/2}$ follows.

If $\varphi \leq \lambda_{\min}(G_N) - 2^{-t} + 2^{-3t/2}$, then consider any eigenstate $|\psi_i\rangle$ of G_N. We see that if $\varphi \leq \lambda_{\min}(G_N) - 2^{-t} + 2^{-3t/2}$, then for any $\left|\psi_g\right\rangle = |g\rangle \left|\xi_g\right\rangle$, where $|g\rangle$ is an eigenstate of G_N with corresponding eigenvalue λ_g, $\varphi \leq \lambda_g - 2^{-t} + 2^{-3t/2}$. As a result, for any energy eigenstate $\left|\psi_g\right\rangle$, by lemma 7.5,

$$\eta(N, \varphi, t, |\psi_i\rangle) \leq \frac{\pi^2}{24}.$$

Any state $|\nu\rangle \in (\mathbb{C}^d)^{\otimes N}$ can be written as $|\nu\rangle = \sum_g \kappa_g(N) |g_N\rangle \left|\xi_{N,g}\right\rangle$; hence by convexity of eq. (7.4) in the coefficients of $|\nu\rangle$ we have

$$\max_{|\nu\rangle} \eta(N, \varphi, t, |\nu\rangle) \leq \frac{\pi^2}{24}.$$

□

7.5.2 An Approximate Phase Comparator QTM

So far we have assumed that the QTMs can implement the algorithms without error, by providing them with all the necessary gates required. In this section we relax these assumptions and show that the error in the output is bounded sufficiently small for our purposes, even if the QTM only has access to a fixed universal gate set.[6]

If we wish our QTM to have a fixed predetermined number of gates available for an arbitrary track length t and arbitrary length inputs N, φ, the gate powers of $U_{G_{z'}}$, as well as the controlled rotations R_k necessary for the Fourier transform subroutine for QPE in lemma 7.4 cannot be given explicitly; we need to approximate them. The $R_k = \mathrm{diag}(1, e^{-\mathrm{i}\pi 2^k})$ gates can be approximated via the Solovay-Kitaev algorithm to the necessary precision. For $U_{G_{z'}}$, however, such a simple compilation argument does not work: we need to perform Hamiltonian simulation in order to implement $U_{G_{z'}}$ itself, for any given spin chain length $z' \in [N]$; to this end, we include the following result.

Lemma 7.6 (Hamiltonian Simulation QTM)**.** *Let $S := \{h^{(l)}\}_{l \in I}$ be a constant and finite set of local interactions of a translationally-invariant Hamiltonian $H = \sum_{i=1}^{z} h_i$, $h_i \in S$, defined on a spin chain of length $z \in \mathbb{N}$, and let $\epsilon > 0$. Then there exists a QTM which, on input S and z, simulates the time evolution $U(T) := \exp(\mathrm{i}HT)$ as a circuit $\tilde{U}(T)$ to precision $\|\tilde{U}(T) - U(T)\| \leq \epsilon$ in spectral norm, in time $\tilde{\mathrm{O}}(T^2 z^2/\epsilon)$.*[7]

Proof. This is a straightforward application of a second order Trotter formula (see e.g. [CS19]). We first assume that we can implement the local time evolution operators $U_i(T) = \exp(\mathrm{i}T h_i)$ for $H = \sum_i h_i$ exactly. For time $\delta > 0$, a second order Trotter formula (e.g. [Ber+06]) breaks up :

$$\tilde{U}(\delta) = \prod_{i=1}^{z} U_i(\delta/2) \prod_{i=z}^{1} U_i(\delta/2),$$

[6] The motivation for this is that when the QTM is encoded in a Hamiltonian, in order to have a fixed local Hilbert space dimension for all N, φ, and t, the QTM must have a gate set which does not depend on N, φ or t.

[7] As per convention, $\tilde{\mathrm{O}}$ hides polylogarithmic factors in the argument.

which requires $O(z)$ gates, and has an error bound

$$\|\tilde{U}(\delta) - U(\delta)\| = O\left(z\delta^2\right).$$

As we require a simulation to time T, the overall error will be

$$\|\tilde{U}(T) - U(T)\| \leq \frac{T}{\delta}\|\tilde{U}(\delta) - U(\delta)\| = O\left(T\delta z\right),$$

requiring $O(zT/\delta)$ gates. Demanding $T\delta z \leq \epsilon$ means $\delta = \Theta(\epsilon/Tz)$; the Trotter simulation thus requires $\Theta(z^2T^2/\epsilon)$ gates overall.

In case we cannot implement the local Trotter operators $U_i(\delta)$ exactly, we have to approximate them with a sequence of elementary gates $\tilde{U}_i(\delta)$, e.g. using Solovay-Kitaev [DN05b]: as errors in a circuit accumulate at most linearly, the approximation has to be precise to

$$\|\tilde{U}_i(\delta) - U_i(\delta)\| = \epsilon/\Theta(z^2T^2/\epsilon) = \Theta(\epsilon^2/z^2T^2) =: \epsilon'$$

where ϵ' is now the precision we must approximate each $\tilde{U}_i(\delta)$. Now we know that Solovay-Kitaev can approximate any unitary operation to within precision ϵ' within $O(\log^4 1/\epsilon')$ many steps. The claim follows. □

Since we will be encoding our QTMs in a Hamiltonian where part of the ground state is chosen in a nondeterministic manner, approximating gates leads not only to slight errors in the gates, but also since the gates no longer have the same eigenvalues, it may lead to differences between which eigenvalue is nondeterministically chosen when the unconstrained input state $|\nu\rangle$ is nondeterministically chosen. In the following two lemmas we characterise this error.

Lemma 7.7. *Let $V(s) = e^{iHs}$ for a Hamiltonian H such $\|H\|_\infty \leq \pi/4s$, and let $\tilde{V}(s)$ satisfy $\|V(s) - \tilde{V}(s)\|_\infty \leq \epsilon$. Then there exist an effective Hamiltonian H' such that $\tilde{V}(s) = e^{iH's}$, and a constant $\kappa = O(1)$ such that*

$$\|H' - H\|_\infty \leq \frac{\kappa\epsilon}{s}.$$

Proof. See supplementary information of [PW09]. □

The following lemma shows that even when the Hamiltonian simulation regime is used (to sufficient accuracy) the new value of η maximised over all input states, has similar bounds to the value obtained without Hamiltonian simulation.

Lemma 7.8 (Hamiltonian Simulation Error)**.** *Let $\mathcal{M}(N,\varphi,t,|\nu\rangle)$ be the QTM described in lemma 7.4 with all gates done without error. Then there exists a QTM $\mathcal{M}'(N,\varphi,t,|\nu\rangle)$ performing the same algorithm, except where the phase estimation for $U_{G_{z'}}$ is instead performed by a Hamiltonian simulation algorithm in lemma 7.6, and such that $\mathcal{M}'$ satisfies the following:*

1. *Let $\eta'(N,\varphi,t,|\nu\rangle)$ be defined in the same way as $\eta(N,\varphi,t,|\nu\rangle)$ from eq. (7.4), but corresponding to the output of $\mathcal{M}'$. The following bounds are satisfied:*

   - *If $t \geq |N|$ and if $\varphi \leq \lambda_{\min}(G_N) - 2^{-t} + \mathrm{O}(2^{-3t/2})$, we have*

   $$\max_{|\nu\rangle} \eta'(N,\varphi,t,|\nu\rangle) \leq \frac{\pi^2}{24}.$$

   - *If $t \geq |N|$ and $\varphi \geq \lambda_{\min}(G_N) - \mathrm{O}(2^{-3t/2})$, we get*

   $$\max_{|\nu\rangle} \eta'(N,\varphi,t,|\nu\rangle) \geq 1 - \frac{\pi^2}{24}.$$

   - *If $t < |N|$, then irrespective of the value of φ,*

   $$\max_{|\nu\rangle} \eta'(N,\varphi,t,|\nu\rangle) = \mathrm{O}\left(\frac{1}{2^{t/2}}\right).$$

2. *The runtime overhead relative to $\mathcal{M}$ is at most* $\mathrm{poly}(N, 2^t)$.

Proof. We first note that the largest power of $e^{\mathrm{i}\pi G_N}$ to be performed in the QPE in lemma 7.4 is $T = 2^t$. Writing $G'_N := 4G_N/\pi N$, we have

$$e^{\mathrm{i}\pi G_N T} = e^{\mathrm{i}G'_N \pi^2 NT/4} = \left(e^{\mathrm{i}G'_N}\right)^{\pi^2 NT/4}.$$

We set $V(s) \coloneqq e^{\mathrm{i}G'_N s}$, and note that $V(s)$ satisfies the conditions of lemma 7.7 for $s \leq 1$. Let H' be the effective Hamiltonian generated in the Hamiltonian simulation scheme given in lemma 7.6 for the short pulse $V(1)$, such that $\tilde{V}(s) = e^{\mathrm{i}H's}$ and $\|V(s) - \tilde{V}(s)\|_\infty \leq \epsilon'$ is the precision to which we want to perform Hamiltonian simulation of the small time step $V(1)$; we leave ϵ' implicit for now, and determine its scaling in due course. By lemma 7.7 it holds that $\|H' - G'_N\|_\infty \leq \kappa\epsilon'$ for some constant κ. Set $H'' \coloneqq \frac{\pi N}{4} H'$, then

$$\|H'' - G_N\|_\infty = \left\| \frac{\pi N}{4} H' - \frac{\pi N}{4} G'_N \right\| \leq \frac{\pi N}{4} \kappa\epsilon' \leq \frac{\kappa\epsilon}{\pi T}$$

where we have chosen

$$\epsilon' \leq \frac{4\epsilon}{\pi^2 NT},$$

and consequently

$$\begin{aligned} |\lambda_{\min}(\pi G_N) - \lambda_{\min}(\pi H'')| &\leq \kappa\epsilon/T \\ |\lambda_{\min}(2\pi G_N) - \lambda_{\min}(2\pi H'')| &\leq 2\kappa\epsilon/T \\ &\vdots \\ |\lambda_{\min}(\pi T G_N) - \lambda_{\min}(\pi T H'')| &\leq \kappa\epsilon. \end{aligned}$$

This immediately implies that the deviation for QPE even in the highest Endian bit is upper-bounded by $\kappa\epsilon$.

Let $\tilde{U}(T) \coloneqq \tilde{V}(1)^{\pi^2 NT/4}$. Then by an iterative expansion we have

$$\begin{aligned} \|U(T) - \tilde{U}(T)\|_\infty &= \|V(1)^{\pi^2 NT/4} - \tilde{V}(1)^{\pi^2 NT/N}\|_\infty \\ &\leq \frac{\pi^2 NT}{4} \|V(1) - \tilde{V}(1)\|_\infty \\ &\leq \frac{\pi^2 NT}{4} \epsilon' = \epsilon. \end{aligned}$$

Now we would like this deviation of QPE to be *less* than the smallest digit of precision of the exact QPE, which is satisfied for $\epsilon = \mathrm{o}(2^{-2t})$. We thus know that if $\varphi > \lambda_{\min}(G_{z'}) - 2^{-3t/2}$ or $\varphi < \lambda_{\min}(G_{z'}) - 2^{-t} + 2^{-3t/2}$, then $\varphi > \lambda_{\min}(H'') -$

$\kappa\epsilon/\pi$ or $\varphi < \lambda_{\min}(G_{z'}) - 2^{-t} + 2^{-3t/2} + \kappa\epsilon/\pi$, respectively. Thus, by lemma 7.5 and corollary 7.1, the same bounds as in corollary 7.1 hold if we choose $\epsilon = \mathrm{O}(2^{-2t})$.

The runtime overhead is then determined by an outer loop of applying $V(s)$ $\pi^2 NT/4 = \mathrm{poly}\, 2^t$ times, and by the cost of approximating $V(s)$ to precision ϵ' in spectral norm, which by lemma 7.6 takes $\tilde{\mathrm{O}}(N^2/\epsilon') = \mathrm{O}(N^3 T/\epsilon) = \mathrm{O}(N^3 2^{3t} t)$, which is $\mathrm{O}(\mathrm{poly}(N, 2^t)$. The claim follows. □

We now fully characterise the output of the QTM when non-determinism and approximate gate sets are taken into account. That is, the QTM only has access to a standard universal gate set and the gates U_φ, U_N.

Lemma 7.9 (Gate Approximation Error)**.** *Let $\mathcal{M}'(N, \varphi, t, |\nu\rangle)$ be the QTM described in lemma 7.8 with the Hamiltonian simulation subroutine performed, but all other gates still done exactly. Then there exists a QTM $\mathcal{M}''(N, \varphi, t, |\nu\rangle)$ that satisfies the following.*

1. *$\mathcal{M}''$ only has access to a fixed universal gate set and the gates U_φ, U_N.*
2. *Let $\eta''(N, \varphi, t, |\nu\rangle)$ be defined in the same way as $\eta(N, \varphi, t, |\nu\rangle)$ from eq. (7.4), but corresponding to the output of $\mathcal{M}''$. Then $\max_{|\nu\rangle} \eta''(N, \varphi, t, |\nu\rangle)$ satisfies the same bounds as $\max_{|\nu\rangle} \eta'(N, \varphi, t, |\nu\rangle)$ in lemma 7.8.*
3. *The additional runtime overhead relative to $\mathcal{M}'$ is at most a factor $\mathrm{poly}\log(N, 2^t)$.*

Proof. By lemma 7.6, we already know that the Hamiltonian simulation subroutine utilises a fixed gate set. We approximate all other gates (apart from U_N and U_φ, which are given explicitly)—of which there are at most $\#g := \mathrm{poly}(N, 2^t)$ many by combining the runtime of $\mathcal{M}$ from lemma 7.4 and runtime overhead of $\mathcal{M}'$ from lemma 7.8. We know that in order to approximate a const-local gate U to precision ϵ using Soloay-Kitaev, an overhead $\mathrm{O}(\log^4 1/\epsilon)$ is introduced. Choosing $\epsilon = 1/\mathrm{poly}(N, 2^{2t})$ small enough such that $\#g\epsilon = 2^{-2t}/t$ suffices to satisfy both claims. □

We need one final property of η'': we will need to show that it is monotonically increasing within a certain region. This doesn't follow straightforwardly from the

monotonicity of η proved in lemma 7.5 as the approximation using Solovay-Kitaev and the Hamiltonian simulation routine may change the gradient $\partial\eta''/\partial\varphi$ to negative at some point. We show this is not the case for a sufficiently large precision.

Lemma 7.10 (Approximate η'' Monotonicity)**.** *Let $|v\rangle$ be an eigenvector of G_N with eigenvalue λ, and let $\eta''(N,\varphi,t,|v\rangle)$ be defined as in lemma 7.9. Then, provided the gates are approximated to precision $\epsilon \leq 2^{-2t}$, $\eta''(N,\varphi,t,|v\rangle)$ is monotonically increasing for $\varphi \in \left(\lambda - 2^{-t} + \mathrm{O}(2^{-3t/2}), \lambda - \mathrm{O}(2^{-3t/2})\right)$.*

Proof. We can express $\eta = \langle\psi| U(\varphi) |\phi\rangle$ for some initial states $|\phi\rangle$, and some appropriate state $|\psi\rangle$, respectively, and $U(\varphi) = U_1 U_{\mathrm{PG}}(\varphi) U_2$, where $U_{\mathrm{PG}}(\varphi)$ denotes the phase gradient part dependent on φ, and U_1 and U_2 collect all the unitary operations before and after (and are independent of φ). We then have

$$\frac{\partial\eta}{\partial\varphi} = \langle\psi| U_1 \left(\frac{\partial U_{\mathrm{PG}}(\varphi)}{\partial\varphi}\right) U_2 |\phi\rangle .$$

This means the deviation is entirely dependent on the gradient of the phase gradient matrix (as expected).

Denote with $\tilde{U}_1, \tilde{U}_2$ the circuits U_1 and U_2 circuits with the Solovay-Kitaev approximation used for any gates which cannot be performed exactly, which we assume to be implemented to accuracy ϵ, i.e. such that $\|U_i - \tilde{U}_i\|_\infty \leq \epsilon$ for $i = 1, 2$. We emphasise our assumption that no gate in U_{PG} has to be approximated. Analogously to before, we then have

$$\frac{\partial\tilde{\eta}}{\partial\varphi} = \langle\psi| \tilde{U}_1 \left(\frac{\partial U_{\mathrm{PG}}(\varphi)}{\partial\varphi}\right) \tilde{U}_2 |\phi\rangle .$$

Consequently,

$$\begin{aligned} \left|\frac{\partial\tilde{\eta}}{\partial\varphi} - \frac{\partial\eta}{\partial\varphi}\right| &\leq \left|\mathrm{Tr}\left(\frac{\partial U_{\mathrm{PG}}(\varphi)}{\partial\varphi}\left(\tilde{U}_1 |\phi\rangle\langle\psi| \tilde{U}_2 - U_1 |\phi\rangle\langle\psi| U_2\right)\right)\right| \\ &\leq \left\|\frac{\partial U_{\mathrm{PG}}(\varphi)}{\partial\varphi}\right\|_\infty \left\|\tilde{U}_1\tilde{U}_2 - U_1 U_2\right\|_\infty . \end{aligned} \tag{7.16}$$

The dependence of $U_{\mathrm{PG}}(\varphi)$ only comes from the controlled U_φ operations, and

$U_{\mathrm{PG}}(\varphi) = cU_\varphi U_\varphi^2 \dots cU_\varphi^{2^t-1}$ ($U_{\mathrm{PG}}(\varphi)$ actually has a set of t Hadamards, however, we can absorb these into U_1 for convenience). All controlled gate powers within the phase gradient circuit are of the form

$$\mathrm{diag}\left(1,1,1,\exp\left(2\pi\mathrm{i}\varphi 2^k\right)\right) \quad \text{with derivative} \quad 2\pi\mathrm{i}\times 2^k\,\mathrm{diag}\left(0,0,0,\exp\left(2\pi\mathrm{i}\varphi 2^k\right)\right).$$

Thus using the product formula we can write

$$\begin{aligned}\left\|\frac{\partial U_{\mathrm{PG}}(\varphi)}{\partial\varphi}\right\|_\infty \leq{}& \left\|\frac{\partial cU_\varphi}{\partial\varphi} cU_\varphi^2 \dots cU_\varphi^{2^t-1}\right\|_\infty \\ &+ \left\|cU_\varphi \frac{\partial cU_\varphi^2}{\partial\varphi} cU_\varphi^{2^2} \dots cU_\varphi^{2^t-1}\right\|_\infty \\ &\vdots \\ &+ \left\|cU_\varphi cU_\varphi^2 \dots \frac{\partial cU_\varphi^{2^t-1}}{\partial\varphi}\right\|_\infty \\ \leq t &\left\|\frac{\partial cU_\varphi^{2^t-1}}{\partial\varphi}\right\|_\infty\end{aligned}$$

From this standard product formula for the derivative of a sequence of gates and using eq. (7.16) this means

$$\left|\frac{\partial\tilde{\eta}}{\partial\varphi} - \frac{\partial\eta}{\partial\varphi}\right| \leq 2\pi\times t2^t\times 2\epsilon.$$

By lemma 7.5, $\partial\eta/\partial\varphi$ within the interval $\varphi \in (\lambda - 2^t + \mathrm{O}(2^{3t/2}), \lambda - \mathrm{O}(2^{3t/2}))$ is ≥ 1; it thus suffices to demand $2\pi\times t2^t\times 2\epsilon \leq 2^{-t}$; a choice of $\epsilon = 2^{-2t}$ proves the claim. □

The results within this section then culminate in a result proving that important properties which hold for η hold for the approximated version η''; in particular montonicity in a particular interval and bounds on the energy outside of this interval.

Theorem 7.7 (Phase Comparator QTM)**.** *Let $N \in \mathbb{N}$, $\varphi \in [0,1]$. For any Gottesman-Irani Hamiltonian G_N there exists a quantum Turing machine $\tilde{\mathcal{M}}(N,\varphi,t,|\nu\rangle)$ with access to special gates U_N and U_φ as in lemma 7.4 which, on input $t \in \mathbb{N}$, $|\nu\rangle \in (\mathbb{C}^d)^{\otimes N}$,*

and in time $\mathrm{poly}(2^t, N)$ *produces an output state as* $\mathcal{M}''$ *in lemma 7.9. Abbreviating*

$$\tilde{\eta}(N,\varphi,t) \coloneqq \max_{|\nu\rangle} \eta''(N,\varphi,t,|\nu\rangle), \tag{7.17}$$

where $\eta''(N,\varphi,t,|\nu\rangle)$ *is defined in in lemma 7.9, we have that*

$$\tilde{\eta}(N,\varphi,t) \begin{cases} \geq 1 - \frac{\pi^2}{24} & t \geq |N| \text{ and } \varphi \geq \lambda_{\min}(G_N) - \mathrm{O}(2^{-3t/2}) \\ \leq \frac{\pi^2}{24} & t \geq |N| \text{ and } \varphi \leq \lambda_{\min}(G_N) - 2^{-t} + \mathrm{O}(2^{-3t/2}) \\ = \mathrm{O}(2^{-t/2}) & t < |N|. \end{cases}$$

Furthermore, $\tilde{\eta}(N,\varphi,t)$ *is monotonically increasing in the interval*

$$\varphi \in \left[\lambda_{\min}(G_N) - \mathrm{O}(2^{-3t/2}), \lambda_{\min}(G_N) - 2^{-t} + \mathrm{O}(2^{-3t/2})\right].$$

Proof. We take $\mathcal{M}''$ from lemma 7.9, and leave the input for the $|\nu\rangle$ section unconstrained. The rest follows by corollary 7.1 and lemma 7.9. The fact $\tilde{\eta}$ is monotonically increasing in the given region is proven in lemma 7.10. □

One fact that we have glossed over is that we can assume that the QTM in theorem 7.7 is well-formed as defined in [BV97, Def. 3.3]—as its evolution is trivially unitary—and well-behaved as in [BV97, Def. 3.12]; the latter condition simply means that the QTM halts in a final state such that the halting head state is in the same cell.

Moreover, we can assume further "good" properties we wish: unidirectionality (meaning each state is only ever entered from one direction) which Bernstein and Vazirani show can be simulated if not originally present ([BV97, Lem. 5.5]).

7.5.3 A Phase Comparator History State Hamiltonian

In this section we translate the QTM designed in section 7.5.1 into a history state Hamiltonian. This technique is by now standard [KSV02; GI09] and used ubiquitously throughout literature (see e.g. [Bau20, Sec. 4.1], or [BC18a, Sec. 1], for an overview). We start with the following refinement regarding standard form Hamiltonians (those with an initial and final penalty at the start and end of the computation).

Theorem 7.8 (Adaptation of Theorem 3.4 from [Wat19]). *Let $H(\mu)$ be standard-from Hamiltonian with minimum output penalty μ on the final time step, such that the encoded QTM has runtime T. Then if $\mu = \frac{k}{256T}$ for $0 \leq k \leq 1$, the following bound holds:*

$$\frac{0.99k}{256T^2} \leq \lambda_{\min}\left(H\left(\frac{k}{256T}\right)\right) \leq \frac{1.05k}{256T^2}$$

Following from this, we take the arguably shortest rigorous route, and directly formulate the following theorem.

Theorem 7.9 (Phase Comparator Hamiltonian). *Let $N \in \mathbb{N}$, and $\varphi \in [0,1]$. For any Gottesman-Irani Hamiltonian G_N there exists a constant $d > 0$, and Hermitian operators $h^{(1)} \in \mathcal{B}(\mathbb{C}^d)$, $h^{(2)} \in \mathcal{B}(\mathbb{C}^d \times \mathbb{C}^d)$, such that*

1. *$h^{(1)}, h^{(2)} \geq 0$, with matrix entries in $\mathbb{Z}$.*
2. *$h^{(2)} = A + e^{i\pi\varphi}B + e^{-i\pi\varphi}B^\dagger + e^{i\pi 0.\,\mathrm{enc}(N)}C + e^{-i\pi 0.\,\mathrm{enc}(N)}C^\dagger$, where*
   - *$B, C \in \mathcal{B}(\mathbb{C}^d)$ with coefficients in $\mathbb{Z}$, and*
   - *$A \in \mathcal{B}(\mathbb{C}^d)$ is Hermitian and with coefficients in $\mathbb{Z} + \mathbb{Z}/\sqrt{2} + e^{i\pi/4}\mathbb{Z}$.*

Define a translationally-invariant nearest-neighbour Hamiltonian on a spin chain of length L via

$$H_{\mathrm{QTM}}(L) := \sum_{i=1}^{L} h_i^{(1)} + \sum_{i=1}^{L-1} h_{i,i+1}^{(2)}.$$

Denote with $|\blacksquare\rangle$ and $|\textcircled{\scriptsize II}\rangle$ two special basis states of $\mathbb{C}^d$, and for $m \in \mathbb{N}$, denote the bracketed *subspace*

$$\mathcal{S}_{\mathrm{br}}(m) := |\textcircled{\scriptsize II}\rangle \otimes (\mathbb{C}^d)^{\otimes m} \otimes |\blacksquare\rangle \otimes (\mathbb{C}^d)^{\otimes (L-m)} |\textcircled{\scriptsize II}\rangle\,. \tag{7.18}$$

Then $H_{\mathrm{QTM}}(L)$ has the following properties.

3. *$H_{\mathrm{QTM}}(L) = \bigoplus_{m=1}^{L-1} H(L,m) \oplus R$, where $H(L,m) := H_{\mathrm{QTM}}(L)|_{\mathcal{S}_{\mathrm{br}}(m)}$; i.e. $H_{\mathrm{QTM}}(L)$ is block-diagonal with respect to the subspaces spanned by $\mathcal{S}_{\mathrm{br}}(m)$, and R captures the remaining block.*
4. *$R \geq 1$.*

5. $\lambda_{\min}(H(L,m)) \geq 1$ *if* $m = 0, 1$.

6. *There exist* $L_N = \operatorname{poly} N$ *and* $m_N = \operatorname{poly}\log_2 N$ *and an integer constant* b*, such that the ground state energy* $\lambda_{\min}(H(L,m))$ *of the other blocks satisfies*

$$\lambda_{\min}(H(L,m)) \begin{cases} \leq \frac{1.05}{256L^b}\frac{\pi^2}{24} & (m,L) = (m_N, L_N) \wedge \varphi \geq \lambda_{\min}(G_N) - \mathrm{O}(N^{-6C}) \\ \geq \frac{0.99}{256L^b}\left(1 - \frac{\pi^2}{24}\right) & (m,L) = (m_N, L_N) \wedge \varphi \leq \lambda_{\min}(G_N) - N^{-4C} + \mathrm{O}(N^{-6C}) \\ \geq \frac{0.99}{256L^b}\left(1 - \frac{\pi^2}{24}\right) & m < |N| \vee (m,L) \neq (m_N, L_N), \end{cases}$$

where $T(L) = L^{b/2}$ *is the runtime of the encoded computation.*

7. *If* $(m,L) = (m_N, L_N)$*, then* $\lambda_{\min}(H(L,m))$ *is monotonically decreasing with* φ *for*

$$\varphi \in \left[\lambda_{\min}(G_N) - N^{-4C} + \mathrm{O}(N^{-6C}), \lambda_{\min}(G_N) - \mathrm{O}(N^{-6C})\right].$$

Proof. A QTM can be translated into a 2-local quantum Thue system [BCO17]; its associated Hamiltonian is then a 2-local nearest-neighbour translationally-invariant Hamiltonian. Classical QTS transition rules yield integer matrix entries; the only nontrivial matrix entries stem from transition rules involving quantum letters (i.e. those that label sites where the quantum state is encoded, and transitions between those are unitary)—which in turn are simply the quantum gates available to the computation we will encode; as the phase gradient gates U_N, U_φ, and a universal gate set comprising CNOT, Hadamard and T will suffice for our purposes, the first two claims follow.

It is furthermore clear that one can statically penalise all but the bracketed configurations, such that $R \geq 1$, and all bracketed states in $\mathcal{S}_{\mathrm{br}}(0)$ and $\mathcal{S}_{\mathrm{br}}(1)$ (see [GI09, Sec. 5] for how this can be done). Moreover, the QTM we will construct will treat $|\blacksquare\rangle$ as a passive state, i.e. it will never move the block; as such, on the bracketed states themselves, the Hamiltonian is block-diagonal with respect to the position of the $|\blacksquare\rangle$ marker state, which is assumed to be at distance m along the spin chain. The next three claims follow.

Let N^{-C} be the promise gap of the local Hamiltonian problem $\lambda_{\min}(G_N)$ for

some constant $C \in \mathbb{N}$.[8] In order to resolve this promise gap with the given precision of t bits, we would require

$$2^{-t} \leq N^{-C} \quad \Longleftarrow \quad t \geq 2 \times 4C \lceil \log_2 N \rceil =: m_N \tag{7.19}$$

where the extra factor of 4 was added such that at least four times as many bits than necessary are resolved. The computation we encode then performs the following steps.

1. Translate the segment length L into binary onto a track, and do the same with m.

2. Perform $\tilde{\mathcal{M}}$ from theorem 7.7, using $t = m$ as precision input.

3. Verify that

   (a) $L = L_N := 2 + 4t + t + N$, and

   (b) $t = m = m_N$ as defined in eq. (7.19), and

   (c) t is large enough such that $\tilde{\eta}(N, \varphi, t) \leq \pi^2/24$ even in case $t < |N|$ (i.e. $\mathrm{O}(2^{-t/2}) \leq \pi^2/24$ in theorem 7.7).

   If any of the above conditions do not hold, set a penalty flag.

It is a standard exercise to ensure that in all computational branches the size of the history state (i.e. the length of the computation) has the same length; this is usually done by introducing a global clock and idling steps (we point the reader to [CPW15, Sec. 4]). We can assume that the runtime of all of the above computation is precisely $T(L) = L^{b/2}$ for some even integer constant $b > 0$.

We assume our history state Hamiltonian features a single in- and output penalty, as in section 3.6; it is straightforward to show that $H(L, m)$ is also standard form as in [Wat19, Sec. 5]. We choose the output penalty to penalise the complement of the accepting subspace defined by the projector $\Pi = (\mathbb{1} - |11\rangle\langle 11|_f)$ on the final time step of the computation, as well as the case where $L \neq L_N$ (this can both be done locally

[8] We can always shrink the promise gap to obtain this scaling for some integer C.

for standard form Hamiltonians); the penalty will have strength $1/256T^2 = 1/256L^b$. Since we do not want our Hamiltonian to have another explicit dependence on L, we remark that this can be done by rotating an ancilla

$$|0\rangle_a \longmapsto \delta\,|0\rangle_a + \sqrt{1-\delta^2}\,|1\rangle_a \quad \text{for} \quad \delta = \frac{1}{256L^b} \tag{7.20}$$

as L and b are both known; then the penalty can be conditioned onto $|0\rangle\langle 0|_a \otimes \Pi =: \Pi'$. Then for a valid history state $|\chi\rangle$ and by using theorem 7.8, we have that

$$\mathrm{Tr}[|\chi\rangle\langle\chi|\,\Pi'] = \frac{1}{256T^2}(1 - E(L,N,\varphi)),$$

where $E(L,N,\varphi)$ is the weight on the accepting subspace of the computation, which is the product of $\tilde{\eta}$ (i.e. $\tilde{\mathcal{M}}$'s output) and the test that $L = L_N$ and $m = |N|$.[9]

The case where $m = t$ is too short to expand N in full is captured by the output of the QTM $\tilde{\mathcal{M}}$, i.e. by theorem 7.7 we have that $\tilde{\eta}(N,\varphi,t) = \mathrm{O}(2^{-t/2})$ in this case, and hence also $E(L,N,\varphi) = \mathrm{O}(2^{-t/2})$.

Let us thus focus on the case when t is large enough (i.e. $t \geq |N|$). If $L \neq L_N$ or $m \neq m_N$, $E(L,N,\varphi) = 0$ by construction, so we only need to analyse the remaining case of $L = L_N$ and $m = m_N$. By theorem 7.7 the accepting state overlap is then

$$\tilde{\eta}(N,\varphi,t) \begin{cases} \geq 1 - \frac{\pi^2}{24} & \varphi \geq \lambda_{\min}(G_N) - \mathrm{O}(N^{-6C}) \\ \leq \frac{\pi^2}{24} & \varphi \leq \lambda_{\min}(G_N) - N^{-4C} + \mathrm{O}(N^{-6C}). \end{cases}$$

where we made use of the fact that $2^{-t} = 2^{-m_N} = N^{-4C}$ and $2^{-3t/2} = 2^{-3m_N/2} = N^{-6C}$, and thus overall

$$E(L,N,\varphi) \begin{cases} \geq 1 - \frac{\pi^2}{24} & (L,m) = (L_N,m_N) \wedge \varphi \geq \lambda_{\min}(G_N) - \mathrm{O}(N^{-6C}) \\ \leq \frac{\pi^2}{24} & (L,m) = (L_N,m_N) \wedge \varphi \leq \lambda_{\min}(G_N) - N^{-4C} + \mathrm{O}(N^{-6C}) \\ \leq \frac{\pi^2}{24} & t = m < |N| \vee (L,m) \neq (L_N,m_N). \end{cases} \tag{7.21}$$

[9] We note that the last test can fail to produce the right result if $t = m$ was too small to begin with to expand enough bits of N; but in this case, the output of the QTM $\tilde{\mathcal{M}}$ already asserts a small acceptance probability, by theorem 7.7.

The bounds in the statement then follow from combining eq. (7.21) with theorem 7.8.

Finally, to prove $\lambda_{\min}(H(L,m))$ is monotonically decreasing with φ, we note that $1-\tilde{\eta}(N,\varphi,t)$ is monotonically decreasing (since theorem 7.7 shows $\tilde{\eta}(N,\varphi,t)$ is monotonically increasing). Since the Hamiltonian is standard form, this acts as a penalty of the form $\frac{1}{256T}(1-\tilde{\eta}(N,\varphi,t))$, which is monotonically decreasing, and which can be shown by standard techniques to be equivalent to adding on a positive semi-definite projector [Wat19]. As adding a positive semi-definite matrix to another matrix can never lead to a *decrease* in the combined eigenvalues, the claim follows. □

7.5.4 Combining the Comparator Hamiltonian with a 2D Marker Tiling

We import the 2D Marker Hamiltonian from section 3.8, which describes a checkerboard pattern for which each checkerboard square of size $L\times L$, with a special marker offset at position $m<L$ on one of the edges, has a net negative energy contribution $\propto 1/4^{f(L,m)}$.

To do this, we introduce a special marker state $|\star\rangle$[10] which interacts with the Marker Hamiltonian. The function f is then defined by the placement of $|\star\rangle$, and so we can use the placement of $|\star\rangle$ (controlled by a classical tiling pattern within each of the squares) to define f to have the appropriate properties for our proof.

We paraphrase the following result, tightening the bounds on the Marker Hamiltonian's ground state energy as we go.

Theorem 7.10 (Adjustment from Theorem 3.5)**.** *Let $\mathcal{H}=(\mathbb{C}^d)^{\otimes\Lambda}$ be a square spin lattice Λ with spins of dimension d, and let $c>0$. Further let $f(L)$ be a function such that $f(L)\leq L$ is an integer and computable in time and space $\leq kL$, for some constant $k\in\mathbb{N}$. Then there exists a translationally-invariant nearest-neighbour Hamiltonian $H^{(\boxplus,f)}=\sum_{\langle i,j\rangle}h_{i,j}^{(\boxplus,f)}$ with the following properties:*

1. $H^{(\boxplus,f)}=\bigoplus_{L>0}H^{(\boxplus,f)}(L)\oplus R'$.

[10] Not to be confused with the bracketing state $|\blacksquare\rangle$ in eq. (7.18); and note we also re-use the letter m here to indicate the offset of $|\star\rangle$. This is not necessarily the same offset as the offset for $|\blacksquare\rangle$.

2. $R' \geq 0$.

3. $H^{(\boxplus,f)}(L)$ *has a unique ground state corresponding to a checkerboard tiling. Let* $H^{(\boxplus,f)}(L)|_S$ *be the restriction of the Hamiltonian to a single checkerboard square, then for all* $L \geq 2$,

$$-\frac{9/4}{4^{f(L)}} \leq \lambda_{\min}(H^{(\boxplus,f)}(L)|_S) \leq -\frac{9/4 - 9/4^{f(L)}}{4^{f(L)}}. \tag{7.22}$$

Proof. We construct an augmented checkerboard tiling as in section 3.7 which places the a special marker $|\star\rangle$ offset at $f(L)$, using a classical tiling within the checkerboard square; as $f(L)$ was computable within time and space $\leq kL$ for some constant $k \in \mathbb{N}$ the existence of such a tiling follows by lemma 3.13.

Following theorem 3.4 and theorem 3.5 and the notation therein, the bounds for a Marker Hamiltonian of length L to be found are denoted

$$-\frac{1}{2} - \mathrm{lwr}(w) \leq \lambda_{\min}(\Delta'_w) \leq -\frac{1}{2} - \mathrm{upr}(w)$$

where Δ'_w denotes precisely one segment of the Marker Hamiltonian, encoding a computation of length w—which is L here, but to follow the notation of the lemmas we will amend we stick to w: this computation runtime can then be augmented to $f(w)$.

Lower Bound. Note that in the proof of lemma 3.16, the lower bound was obtained by realising

$$\frac{a-1}{a+1} \leq 1 \ \forall w \quad \Longleftarrow \quad \mathrm{lwr}(w) = \frac{3}{4^w}.$$

Analysing eq. (3.28) of section 3.8 more carefully, we note the same bound also holds for $\mathrm{lwr}(w) = 9/4 \times 4^{-w}$.

Upper Bound. In [Bau+18b, Lem. 8], it is easy to check that the inequality

$$\frac{a-1}{a+1} \leq 4^{-w}$$

also holds when starting with $p_w(-1/2-(9/4-\delta)\times 4^{-w})$, for any

$$\delta \geq \frac{9(5\times 4^2-2)}{4(2+2^{1+4w}-4\times 4^w)} \geq \frac{9}{4^w}.$$

Thus $\mathrm{upr}(w) = (9/4+9/4^w)\times 4^{-w}$ suffices.

Finally, the $-1/2$ offset are removed as in [Bau+18b, Th. 10]. □

7.5.5 From Phase Comparison to Phase Transition

Following section 3.9, we will now combine the QTM Hamiltonian H_{QTM} with the 2D Marker Hamiltonian $H^{(\boxplus,f)}$, to translate the outcome of the comparison $\varphi \lessgtr \lambda_{\min}(G_N)$, where G_N is the Gottesman-Irani Hamiltonian simulated by H_{QTM}, into the question of existence of a negative eigenstate within one square of the 2D Marker Hamiltonian. We will assume $H^{(\boxplus,f)}$ is such that all checkerboard squares have square sizes $L \in 4\mathbb{N}$; this can always be achieved by adding a fixed-dimensional tiling, which we leave implicit in the following.

The aim is to produce an overall Hamiltonian which has *negative energy density* when the encoded computation is accepting, but a *positive energy density* when the encoded computation rejects. The checkerboard structure allows us to create a repeated structure across the lattice; as each square will contribute a finite amount of either positive or negative energy, the density will follow suit.

Lemma 7.11. *Let $H \coloneqq H_{\mathrm{QTM}}\otimes \mathbb{1} + \mathbb{1}\otimes H^{(\boxplus,f)}$ on a spin lattice. Then its ground state is a product state $|\psi\rangle\otimes|T\rangle_c$, where $|T\rangle_c$ is the checkerboard tiling from $H^{(\boxplus,f)}$, and $|\psi\rangle$ the ground state of H_{QTM}. Consider an $L\times L$ square denoted $S(L)$ within the tiling and let $H|_{S(L)}$ be the Hamiltonian restricted to such a square. Then, adopting the notation from eq. (7.21),*

$$\lambda_{\min}(H|_{S(L)}) \begin{cases} < 0 & \text{if } (L,m) = (L_N,m_N) \wedge \varphi \geq \lambda_{\min}(G_N) - \mathrm{O}(N^{-6C}) \\ \geq 0 & \text{if } (L,m) = (L_N,m_N) \wedge \varphi \leq \lambda_{\min}(G_N) - N^{-4C} + \mathrm{O}(N^{-6C}) \\ \geq 0 & \text{if } (L,m) \neq (L_N,m_N). \end{cases}$$

Furthermore, if $L = L_N$ and $m = m_N$, then there is exactly one point in φ where

$\lambda_{\min}(H|_{S(L)})$ changes from < 0 to $= 0$ which occurs in the interval

$$\varphi \in \left[\lambda_{\min}(G_N) - N^{-4C} + \mathrm{O}(N^{-6C}), \lambda_{\min}(G_N) - \mathrm{O}(N^{-6C})\right].$$

Proof. We choose the Marker falloff $f(L)$ such that

$$\frac{9}{4}4^{-f(L)} = \frac{9}{16}\frac{1}{256L^b} \quad \Longrightarrow \quad f(L) = 5 + \log_4(L^b).$$

We now compare the energy (given by theorem 7.10) with the energy of the Hamiltonian encoding the QTM (given in theorem 7.9) and see the following bounds hold for sufficiently large L:

$$\frac{0.99}{256L^b}\left(1 - \frac{\pi^2}{24}\right) \geq \frac{9}{4}4^{-f(L)} \quad \text{and} \quad 4^{-f(L)}\left(\frac{9}{4} - \frac{90}{4\times 2^L}\right) \geq \frac{1.05}{256L^b}\frac{\pi^2}{24}$$

where b is the runtime exponent of $T = T(L) = L^{b/2}$, as given in theorem 7.9. Note $f(L)$ is trivially computable in time and space kL for some constant k.[11] This yields a Marker Hamiltonian with a ground state energy as in theorem 7.10, such that the ground state energy is "sandwiched" with ample margins between the upper and lower bounds of the ground state energy of the Hamiltonian encoding the computation.

As the spectrum is product by construction, the joint spectrum is then $\mathrm{spec}(H) = \mathrm{spec}(H_{\mathrm{QTM}}) + \mathrm{spec}(H^{(\boxplus,f)})$, and the rest follows from lemma 3.18.

Finally the fact there is exactly one point where $\lambda_{\min}(H|_{S(L)})$ changes from < 0 to $= 0$ is due to the fact that (as per point 7 of theorem 7.9) $\lambda_{\min}(H_{\mathrm{QTM}}(L))$ is strictly decreasing for $\varphi \in [\lambda_{\min}(G_N) - N^{-4C} + \mathrm{O}(N^{-6C})), \lambda_{\min}(G_N) - \mathrm{O}(N^{-6C})]$. As per the above analysis, the point at which $|\lambda_{\min}(H_{\mathrm{QTM}}(L))| = |\lambda_{\min}(H^{(\boxplus,f)}(L)|_S)|$ occurs for energy values corresponding to φ in this interval, hence this point at which $|\lambda_{\min}(H_{\mathrm{QTM}}(L))| < |\lambda_{\min}(H^{(\boxplus,f)}(L)|_S)|$ changes to $|\lambda_{\min}(H_{\mathrm{QTM}}(L))| = |\lambda_{\min}(H^{(\boxplus,f)}(L)|_S)|$ can only happen at a single point. □

Since we want the trivial ground state in the gapped phase to have eigenvalue

[11] Indeed: define a tiling pattern that counts in base 4, and does so b times; then counts another 5 steps. Penalise tile configurations indicating that the base-4 expansion of L is not of the form `100...`, corresponding to a number $L = 4^x$ for some integer x. This can all be done with $k = 1$.

zero, we want to shift the Hamiltonian $H \coloneqq H_{\mathrm{QTM}} \otimes \mathbb{1} + \mathbb{1} \otimes H^{(\boxplus, f)}$ constructed above by 1; this is a standard trick, summarised in the following lemma.

Lemma 7.12. *There exists a Hamiltonian H', with the same properties as $H' \coloneqq H + \sum_i P_i$, where P_i is a projector, such that on a lattice $\Lambda(L)$*

$$\lambda_{\min}(H'(\varphi)) \begin{cases} = 1 + \left\lfloor \frac{L}{L_N} \right\rfloor^2 \lambda_{\min}(H(\varphi)|_{S(L_N)}) & \text{if } \varphi \geq \lambda_{\min}(G_N) - \mathrm{O}(N^{-6C}) \\ \geq 1 & \text{if } \varphi \leq \lambda_{\min}(G_N) - N^{-4C} + \mathrm{O}(N^{-6C}), \end{cases}$$

where $H \coloneqq H_{\mathrm{QTM}} \otimes \mathbb{1} + \mathbb{1} \otimes H^{(\boxplus, f)}$.

Proof. From lemma 7.11 we know the energy of a single square $S(L_N)$. If the ground state $\lambda_{\min}(H(\varphi)|_{S(L_N)}) < 0$, then the overall ground state of the lattice becomes a checkerboard of these squares. As per corollary 3.7 in chapter 3, it can be shown that incomplete squares contribute zero energy, giving a total energy of $\left\lfloor \frac{L}{L_N} \right\rfloor^2 \lambda_{\min}(H(\varphi)|_{S(L_N)})$. If $\lambda_{\min}(H(\varphi)|_{S(L_N)}) \geq 0$, then the lattice has ≥ 0.

Finally, by using the energy shift trick of [Bau20, Lem. 23], we can add on an energy shift of 1 to the Hamiltonian, giving the bounds stated in the lemma. □

The final step is then to combine H' from lemma 7.12 with a trivial, a dense, and a guard Hamiltonian, to lift the ground state energy to a ground state energy density statement — this is exactly what was done in theorem 7.5. This modifies the Hamiltonian so that phase transitions can occur between the ground state of the checkerboard Hamiltonian and the ground state of a trivial zero energy state.

Theorem 7.11 (Existence of Two Phases)**.** *Let G_N be a Gottesman-Irani Hamiltonian with promise gap $\sim N^{-C}$ for some constant C. Then there exists a Hamiltonian $H^\Lambda(N, \varphi) = \sum_{\langle i,j \rangle} h^N_{i,j}(\varphi) + \sum_{i \in \Lambda} h^N_i$, and an order parameter $O_{A/B}$ acting on a subset $F \subset \Lambda$ of lattice sites, $|F|$ constant, such that, as $\Lambda \to \infty$ the following holds.*

- *if $\varphi \leq \lambda_{\min}(G_N) - N^{-4C} + \mathrm{O}(N^{-6C})$, then*

  i H^Λ is gapped with spectral gap $\geq 1/2$.

  ii product ground state.

iii has order parameter expectation value $\langle O_{A/B} \rangle = 1$.

- *if* $\varphi \geq \lambda_{\min}(G_N) - \mathrm{O}(N^{-6C})$, *then*

  i H^Λ *is gapless.*

  ii has a ground state with algebraically decaying correlations.

  iii has order parameter expectation value $\langle O_{A/B} \rangle = 0$.

Proof. Take H_{dense} to be a Hamiltonian that has an asymptotically dense spectrum in $[0, \infty)$ on $\mathcal{H}_2$. For convenience we choose H_{dense} to be the 1D critical XY-model [LSM61]. H_{trivial} to be diagonal in the computational basis, with a single product ground state $|0\rangle^{\otimes \Lambda}$, minimum eigenvalue 0 and spectral gap 1 acting on $\mathcal{H}_3$, and H_{guard} acting on $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2 \oplus \mathcal{H}_3$ via

$$H_{\text{guard}} \coloneqq \sum_{i \sim j} \left(\mathbb{1}_{1,2}^{(i)} \otimes \mathbb{1}_3^{(j)} + \mathbb{1}_3^{(i)} \otimes \mathbb{1}_{1,2}^{(j)} \right).$$

Take H' from lemma 7.12, and set

$$H^\Lambda(N, \varphi) \coloneqq H' \otimes \mathbb{1}_2 \oplus 0_3 + \mathbb{1}_1 \otimes H_{\text{dense}} \oplus 0_3 + 0_{1,2} \oplus H_{\text{trivial}} + H_{\text{guard}}.$$

Then

$$\operatorname{spec}(H^\Lambda) = \{0\} \cup (\operatorname{spec}(H') + \operatorname{spec}(H_{\text{dense}})) \cup G$$

for some $G \subset [1, \infty)$, as in the proof of theorem 7.5. From lemma 7.12 we can assume that for lattice sizes going to infinity,

$$\lambda_{\min}(H') \begin{cases} \geq 1 & \varphi \leq \lambda_{\min}(G_N) - N^{-4C} + \mathrm{O}(N^{-6C}) \\ \longrightarrow -\infty & \varphi \geq \lambda_{\min}(G_N) - \mathrm{O}(N^{-6C}), \end{cases}$$

as eventually there will exist a checkerboard square size such that $L = L_N$ and $t = m_N \geq |N|$ is satisfiable; the only differentiating condition left in lemma 7.11 is then $\varphi \leq \lambda_{\min}(G_N) - N^{-4C} + \mathrm{O}(N^{-6C})$. Then if $\lambda_{\min}(H) \geq 0$, we have that $\operatorname{spec}(H) + \operatorname{spec}(H_{\text{dense}}) \subseteq [1, \infty)$. The ground state of H^Λ is the trivial ground state

with spectral gap 1. Otherwise, if $\lambda_{\min}(H) \longrightarrow -\infty$, H^{Λ} becomes asymptotically gapless and dense via H_{dense}.

The order parameter is then defined as

$$O_{A/B} = \frac{1}{|F|} \sum_{i \in F} \left(0_{1,2} \oplus |0\rangle\langle 0|_3\right)^{(i)}$$

which makes it clear that in case the ground state is determined by H_{trivial}, the expectation value $\langle O_{A/B} \rangle = 1$; otherwise zero.

The claim of the algebraically decaying correlation functions follows from the fact that the critical XY-model has algebraically decaying correlations functions [LSM61]. □

It is clear that for our construction we could choose F to only contain a single spin, but we leave the statement in its generic form.

7.5.6 Existence of Exactly One Critical Point

The following lemma shows that there is exactly one critical point between the two phases of the Hamiltonian. We will give the statement of the lemma here, but defer its proof to the two-parameter case (which is the more generic setting).

Lemma 7.13 (Existence of Exactly One Critical Point)**.** *Consider the Hamiltonian $H^{\Lambda}(N, \varphi)$ from theorem 7.11. This has exactly one critical point in the interval*

$$\varphi^* \in \left[\lambda_{\min}(G_N) - N^{-4C} + \mathrm{O}(N^{-6C}), \lambda_{\min}(G_N) - \mathrm{O}(N^{-6C})\right].$$

Proof. As per lemma 7.11 there is exactly one φ within the given interval where $\lambda_{\min}(H|_{S(L)})$ goes from < 0 to ≥ 0, which (as per the proof of theorem 7.11) corresponds to the phase transition from gapped to gapless. □

7.5.7 Reduction of Translationally Invariant Local Hamiltonian to 1-CRT-PRM

We first remark that the parameter range for φ where the critical point can possibly be found is shrinking polynomially, due to the shrinking promise gap of the Gottesman-

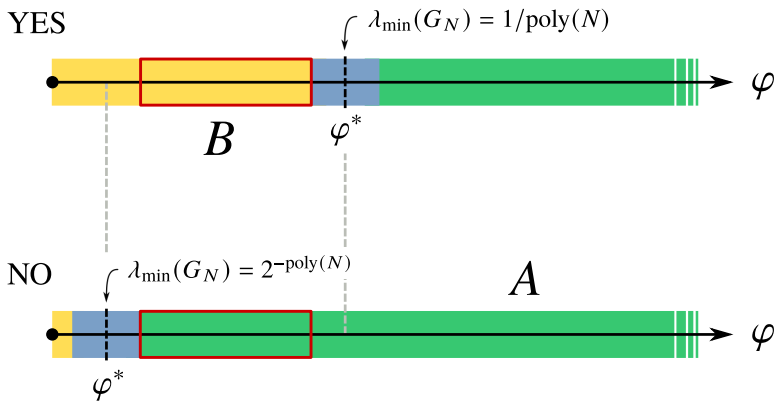

Figure 7.4: The two YES and NO cases (left resp. right) of phase diagrams, for the Hamiltonian in theorem 7.11. By lemma 7.14, the difference between the two dashed vertical lines can be scaled up to $\Omega(1)$. The light blue area indicates a $1/\mathrm{poly}(N)$-sized interval of uncertainty; yet in either case, by lemma 7.13, there exists precisely *one* critical point φ^* therein. The red region indicates an interval (of up to size $\Omega(1)$, by lemma 7.14) which is either entirely in phase A or B; determining which of the two cases holds is $\mathsf{QMA}_{\mathsf{EXP}}$-hard.

Irani Hamiltonian that we encode, and the associated comparison of its ground state energy $\lambda \lessgtr \varphi$. There is now two approaches to scaling up the φ parameter, so that we can get an $\mathrm{O}(1)$-area within the phase diagram where we cannot locate the critical point. This is summarised in the following remark.

Remark 7.1. *Let $0 \le x < t$. There exists a modification to the phase comparator QTM in section 7.5.1 that allows one to perform the rescaled phase comparison $a \lessgtr 2^{-x} b$ for the two unitaries U_a and U_b, where we assumed $1/10 \le b \le 1$, with an error (in the amplitudes of the resulting QPE output) upper-bounded by 2^{-2t}, and with overhead $\mathrm{poly}(2^t)$.*

Proof. By [SMM09], we know that for an unknown "black-box" unitary U, we can implement any power U^y for $y > 0$ of it to precision ϵ (in trace norm) in time $\mathrm{O}(\lfloor y \rfloor + \log(1/\epsilon)/\epsilon)$ (i.e. with that many calls to U), as long as the phase φ we want to estimate in U is not too close to 0, and no other phase lies in $(0, \varphi)$ (a gappedness constraint).

For us, we want to implement U_b^y for $y = 2^{-x} \times 2^m$ for $m = 0, \ldots, t$. By assumption, b satisfies the gappedness condition (as 0 and φ are the only eigenphases of U, and

$\epsilon \leq \varphi$). Since U_b acts on a single qubit only and $0 \leq x < t$, we know that a precision of $\epsilon = 2^{-4t}$ suffices to implement U_b^y such that each controlled rotation gate is off (in operator norm) by at most $\epsilon \times 2^m \leq 2^{-3t}\ \forall m$. As we have t controlled rotation gates, the overall deviation is at most $t \times 2^{-3t} \leq 2^{-2t}$. All amplitudes within the rescaled phase comparator thus at most deviate by 2^{-2t}, and the overhead is poly 2^t, as claimed. □

Lemma 7.14 (Existence of Two Phases with O(1) YES/NO Threshold)**.** *Let p, q be the polynomials defined in theorem 7.3 such that $1/p(N) - 1/q(N) = \Omega(N^{-C})$. There exists a variant of the Hamiltonian $H^{\Lambda}(N, \varphi)$ such that the two cases for φ in theorem 7.11 read*

1. *if $\varphi^* \leq A(N) = N^C(1/q(N) - \mathrm{O}(N^{-6C}))$, and*
2. *if $\varphi^* \geq B(N) = N^C(1/p(N) - N^{4C} + \mathrm{O}(N^{-6C}))$.*

The two bounds satisfy $B(N) - A(N) = \Omega(1)$.

Proof. Follows immediately from remark 7.1, by scaling φ down to be within $\Theta(1)$ of G_N's promise gap (which is a factor $1/\mathrm{poly}\, N \geq 2^{-t}$ for a polynomial we can compute efficiently, proportional to $p(N)$). □

Corollary 7.2. *1-CRT-PRM is $\mathsf{QMA}_{\mathsf{EXP}}$-hard for an $\Omega(1)$ gap.*

Proof. Immediate from lemma 7.14. □

The fact that we can rescale the range of φ to lie within an $\Omega(1)$ region is, in a sense, unsurprising: the same could be said to hold for the local Hamiltonian problem, where one can scale the overall Hamiltonian by a factor to have a $\Omega(1)$ promise gap as well. Note, however, that this is a meaningless transformation: the *precision* to which one wants to obtain the ground state energy is relative to the norm of the Hamiltonian (cf. "relative promise gap" or "relative UNSAT penalty", [BC18a]). There are thus two scale choices for the local Hamiltonian problem: i. the norm of the local terms, or ii. the norm of the overall (finite-sized) Hamiltonian. Naturally, in the first case, one could obtain a stronger local interaction without increasing the

individual coupling's norm by increasing the interaction degree (see e.g. [CN15]). The safer definition is thus the second one—or by limiting the interaction degree of the Hamiltonian to some constant.

In our case, there is no natural "finite size" Hamiltonian relative to which one can define a meaningful precision; the arguably right scale with respect to which one thus has to define φ's order of magnitude is either the local coupling strength (which is constant in our case), or relate it to the parameter N itself. In either case, and after the scaling has been applied, it makes sense to speak of φ to be hard to approximate to $\Omega(1)$ precision, even if that means that φ is now indeterminate in a range $[0, \operatorname{poly} N]$, as stated in theorem 7.1. It is also clear that if we know the polynomial $p(N)$ in lemma 7.14, and we know that it is tight for NO instances of the embedded Hamiltonian G_N, then it would suffice to scan φ within a constant region.[12]

7.5.8 Verifying the Local-Global Promise

Finally, we need to check that the Hamiltonian used to prove hardness—defined above—satisfies the local-global promises as per definition 7.3 and definition 7.4.

Lemma 7.15. *Consider an instance of the Hamiltonian* $H^{\Lambda(L)}(N,\varphi) = \sum_{\langle i,j\rangle} h^N_{i,j}(\varphi) + \sum_{i\in\Lambda} h^N_i$*, as defined in theorem 7.11, with local terms describable in* $|N|$ *bits. Then the Hamiltonian satisfies the global-local phase assumption definition 7.3 for the order parameter*

$$O_{A/B} = \frac{1}{|F|}\sum_{i\in F}\left(0_{1,2}\oplus|0\rangle\langle 0|_3\right)^{(i)}$$

defined in theorem 7.11, and for $L_0 = N^{2+a+b}$*. It also satisfies the global-local gap promise in definition 7.4 for the same* L_0*.*

Proof. Consider an $L > L_N = 2+4t+t+N$. Then, by lemma 7.12, we see that

$$\lambda_{\min}(H'(\varphi))\begin{cases} = 1 + \lfloor \frac{L}{L_N} \rfloor^2 \lambda_{\min}(H(\varphi)|_{S(L_N)}) & \varphi \le \lambda_{\min}(G_N) - N^{-4C} + \mathrm{O}(N^{-6C}) \\ \ge 1 & \varphi \ge \lambda_{\min}(G_N) - \mathrm{O}(N^{-6C}). \end{cases}$$

[12] We remark that this does not work for 2-CRT-PRM in section 7.6, as there the ground state energy is not determined by a single $\mathsf{QMA}_{\mathsf{EXP}}$ query.

Thus, when $\lambda_{\min}(H(\varphi)) \geq 0$ the ground state of $H'(\varphi)$ that of H_{trivial}, i.e. $|0\rangle^{\Lambda(L)}$, and $\langle O_{A/B}\rangle = 1$. On the other hand, when $\lambda_{\min}(H(\varphi)|_{S(L_N)}) < 0$, then eventually the ground state is a highly complex quantum plus classical state with $\langle O_{A/B}\rangle = 0$.

Thus, when $\lambda_{\min}(H(\varphi)|_{S(L_N)}) < 0$, for the highly quantum ground state to appear the lattice size L must meet the following condition:

$$\left\lfloor \frac{L}{L_N} \right\rfloor^2 \lambda_{\min}(H(\varphi)|_{S(L_N)}) < 1. \tag{7.23}$$

Otherwise the ground state is the zero energy state $|0\rangle^{\Lambda(L)}$.

From theorem 7.10, when $\lambda_{\min}(H(\varphi)|_{S(L_N)}) < 0$ and $\varphi \leq \lambda_{\min}(G_N) - N^{-4C} + \mathrm{O}(N^{-6C})$, then the ground state energy is

$$\left\lfloor \frac{L}{L_N} \right\rfloor^2 \left(\lambda_{\min}(H_{\text{QTM}}(L_N)) + \lambda_{\min}(H^{(\boxplus,f)}(L_N)|_S\right) \leq \left\lfloor \frac{L}{L_N} \right\rfloor^2 \left(-\frac{-c_1}{L_N^b}\right),$$

where we have used that $\lambda_{\min}(H_{\text{QTM}}(L_N)) + \lambda_{\min}(H^{(\boxplus,f)}(L_N)|_S = -\Omega(T^{-2}) = -c_1 L^{-b}$ for some constant c_1 (this can be seen by combining theorem 7.8 and lemma 7.11). Thus by for $L > L_0$, where

$$L_0 \geq c_1^{1/2} L_N^{1+b/2},$$

eq. (7.23) will be satisfied. Since $L_N = \mathrm{O}(N)$, we choose $L_0 = \mathrm{O}(N^{2+b})$.

For all $L \geq L_0$ the expectation value of $O_{A/B}$ is then constant, regardless of whether $\varphi \geq \lambda_{\min}(G_N) - \mathrm{O}(N^{-6C})$ or $\varphi \leq \lambda_{\min}(G_N) - N^{-4C} + \mathrm{O}(N^{-6C})$. Thus the Hamiltonian satisfies the global-local phase promise in definition 7.3.

Global-Local Gap Promise: The proof for the global-local gap promise is almost the same. For the L_0 above, we see that if $\varphi \geq \lambda_{\min}(G_N) - \mathrm{O}(N^{-6C})$, then the system has a very negative energy and a spectral gap $\Delta(L) = O(1/L^2)$. If $\varphi \leq \lambda_{\min}(G_N) - N^{-4C} + \mathrm{O}(N^{-6C})$, then the ground state is $|0\rangle^{\Lambda(L)}$ with zero energy and has the same gap as H_{trivial}: $\Delta \geq 1$. □

7.6 $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$ Hardness of 2-CRT-PRM

In this section sketch an outline of the proof of $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$-completeness of 2-CRT-PRM. We only sketch the proof for the purposes of brevity, noting that conceptually the proof is conceptually similar to the 1-CRT-PRM case. The full proof can be found in [WB21].

To prove this we will make a reduction from ∀-TI-APX-SIM, as defined definition 7.8, which was proved to be $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$-complete in chapter 6.

Proof Outline. The proof method here will be similar to the 1-parameter case, but instead of a reduction to the Local Hamiltonian problem, we perform a reduction to ∀-TI-APX-SIM, which is the question of approximating the expectation value of all low-energy states of a Hamiltonian with respect to a local observable. This means we construct—just as described in section 7.5—a Hamiltonian $H_N(\varphi)$ which, in its ground state, encodes the following computation.

1. Perform QPE to extract N from local terms.
2. Perform a phase comparison QPE on the unitary encoding φ and $\exp(itK_N)$, where K_N is a translationally-invariant local spin Hamiltonian with a $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$-complete ∀-TI-APX-SIM problem (on a spin chain of length N). The joint witness stems from an unconstrained input state. If this input state was an eigenstate of K_N with eigenvalue λ, the phase comparator QPE extracts the difference $\lambda - \varphi$ to bit precision $\sim |N|$.
3. If $\varphi < \lambda$, an output flag is set to $|0\rangle$; otherwise it is set to $|1\rangle$.
4. Another flag qubit captures the output bit of the $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$ computation performed within the history state of $H_N(\varphi)$.

An energy penalty is then given to the joint energy eigenvalue comparison *and* output bit of the $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$ computation, in the sense that

1. All eigenstates of K_N that are not considered "low energy" are penalised.
2. Those eigenstates of K_N that fall below the "low energy" cutoff are not inflicted with a penalty; but they are subject to a penalty from the *observable*

operator $(B - \theta\mathbb{1})$ for some scalar offset $\theta > 0$. Here B is the operator for which determining the expectation of on low energy states of K_N is $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$-complete.

The result is that the low-energy eigenspace of $H_N(\varphi)$ plus penalties is greater or smaller than some polynomial falloff we can calculate to high precision, and which will depend on the scalar expectation value offset θ that serves as the second parameter.

As in the one-parameter case, we can combine this energy penalty with a bonus of a Marker Hamiltonian to obtain a joint 1D spin Hamiltonian with the property that it has a *negative* ground state energy if we have a YES-instance—i.e. expectation values of low-energy states lie below some threshold—and the scalar offset is below a cutoff; and a *positive* ground state energy for a No instance, *or* for a too-small scalar offset θ. With standard techniques this dichotomy is then amplified to a gapless resp. gapped phase in the thermodynamic limit.

In order to understand why this two-parameter family of Hamiltonians has a $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$-hard-to-compute phase diagram, note that to figure out the relevant φ region within which a phase transition can occur takes multiple queries to a QMA oracle, as we need to identify $\lambda_{\min}(K_N)$ to sufficient precision; and then we don't yet know whether the output is a YES or No case, so there is *two* possible θ-regions around which to explore the phase diagram.

7.6.1 Additional Preliminaries

It was shown in corollary 6.1 that ∀-TI-APX-SIM is $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$-complete for a Hamiltonian which we label $K_N \in \mathcal{B}(\mathbb{C}^d)^{\otimes N}$. For ease of reading, we restate the lemma (and modify it slightly):

Lemma 7.16 (From corollary 6.1)**.** *There exists a fixed one-local observable A and interaction terms $k_{i,i+1} \in \mathcal{B}(\mathbb{C}^d \otimes \mathbb{C}^d)$ acting between pairs of nearest neighbour qudits, which define a Hamiltonian on a 1D chain of length N, $K_N = \sum_{i=1}^{N-1} k_{i,i+1}$ such that for all states $|\psi\rangle$ that satisfy $\langle\psi| K_N |\psi\rangle \leq \lambda_0(K_N) + \delta$ for $\delta = \Omega(1/\mathrm{poly}(N))$ either of the following holds:*

YES*:* $1 - 1/\mathrm{poly}(N) \leq \langle\psi| A |\psi\rangle \leq 1$*, or*

No: $0 \leq \langle\psi| A |\psi\rangle \leq 1/\mathrm{poly}(N)$.

Determining which case is true is $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$*-complete.*

This is not quite what was proven in corollary 6.1, where e.g. the overlap with the observable in the first case was $1/T_{K_N} - \mathrm{O}(2^{-\mathrm{poly}(N)}) \leq \langle\psi| A |\psi\rangle \leq 1/T_{K_N}$, for T_{K_N} the length of the encoded computation. However, we can adjust the construction using an idling technique from [CLN18] which increases the weight on the output bit of the computation (such that the encoded computation has its runtime increased to $P_1(N)T_{K_N}$, for some polynomial $P_1(N)$ we are free to choose). Furthermore, we ask that there be some marker flag, placed next at or next to the output qubit, which indicates when the first part of the computation—before the idling—has finished (this allows us to keep A as a 1-local operator). These techniques are by now standard, and we will not go into details.

We denote the set of eigenstates of K_N for which the energy expectation value is below the cutoff as

$$S_\delta \coloneqq \left\{\psi : H|\psi\rangle = \lambda|\psi\rangle \text{ where } \lambda \leq \lambda_{\min}(K_N) + \delta\right\}, \tag{7.24}$$

which means $\langle\psi| H |\psi\rangle \leq \lambda_{\min}(K_N) + \delta$ for all $|\psi\rangle \in \mathrm{Span}(S_\delta)$. We also define the shifted observable

$$B \coloneqq A + \mathbb{1}, \tag{7.25}$$

which if A is a projector has eigenvalues in the set $\{1, 2\}$, which we label as $\lambda_0(B) = 1$ and $\lambda_1(B) = 2$. This offset merely simplifies some of the maths in due course. Together with lemma 7.16, this choice of B in eq. (7.25) immediately yields the following corollary.

Corollary 7.3. *We use the notation of lemma 7.16, for an observable A that is a one-local projector, and B as defined in eq. (7.25). Any state $|\psi\rangle \in S_\delta$ for the Hamiltonian K_N then has expectation value either $\leq \lambda_0(B) + \mathrm{O}(1/P_1(N))$, or $\geq \lambda_1(B) - \mathrm{O}(1/P_1(N))$ for a polynomial $P_1(N)$ we are free to choose.*

7.6.2 A Modified Phase Comparator QTM

The following lemma follows the same setup as lemma 7.4.

Lemma 7.17 (Multi-QPE QTM)**.** *Let K_z be the translationally invariant Hamiltonian on chain of length z described in lemma 7.16. Take the same setup as in lemma 7.4, but where the Hamiltonian G_z is replaced by K_z. The output of this QTM $\mathcal{M}(N,\varphi,t,|\nu\rangle)$ will then be*

$$
\begin{aligned}
|\chi\rangle = &\sum_{z\in V_t}\sum_{x\leq 0}\sum_{g}\alpha_x(z',g)\gamma_z\kappa_g(z')\,|11\rangle_f\,|z\rangle\,|x\rangle\sum_j\sigma_j(g,z')\,|j\rangle\left|g_{z',j}\right\rangle\left|\xi_{z',g}\right\rangle + \\
&\sum_{z\in V_t}\sum_{x>0}\sum_{g}\alpha_x(z',g)\gamma_z\kappa_g(z')\,|10\rangle_f\,|z\rangle\,|x\rangle\sum_j\sigma_j(g,z')\,|j\rangle\left|g_{z',j}\right\rangle\left|\xi_{z',g}\right\rangle + \\
&\sum_{z\notin V_t}\sum_{x\leq 0}\sum_{g}\alpha_x(z',g)\gamma_z\kappa_g(z')\,|01\rangle_f\,|z\rangle\,|x\rangle\sum_j\sigma_j(g,z')\,|j\rangle\left|g_{z',j}\right\rangle\left|\xi_{z',g}\right\rangle + \\
&\sum_{z\notin V_t}\sum_{x>0}\sum_{g}\alpha_x(z',g)\gamma_z\kappa_g(z')\,|00\rangle_f\,|z\rangle\,|x\rangle\sum_j\sigma_j(g,z')\,|j\rangle\left|g_{z',j}\right\rangle\left|\xi_{z',g}\right\rangle,
\end{aligned}
\tag{7.26}
$$

where we have expanded $|g_{z'}\rangle = \sum_j \sigma_j(g,z')\,|j\rangle\left|g_{z',j}\right\rangle$ such that the $|j\rangle$ denote the eigenvectors of B defined in eq. (7.25).

Proof. Follows from the output state given in lemma 7.4 and the form of the unconstrained state taken as "input". □

We now need an equivalent expression to eq. (7.4) which captures the expected output penalty that we wish to inflict later on. Here, we will modify the flag projector slightly; instead of using $|11\rangle\langle 11|_f$ that just singles out those eigenstates of K_z that have low energy, we also add in the observable B as defined in eq. (7.25), which acts on the *output* of the computation.[13] As K_z encodes a $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$-hard computation, this output bit is a single qubit; and can be assumed to satisfy the bounds given in corollary 7.3.

Taking the output state $|\chi\rangle$ of $\mathcal{M}(N,\varphi,t,|\nu\rangle)$ from lemma 7.17, and letting B

[13] This penalty can be made 1-local using standard methods. We omit this here.

be the local observable from eq. (7.25) & corollary 7.3, we set

$$\begin{aligned}\eta(N,\varphi,\theta,t,|\nu\rangle) &:= \mathrm{Tr}\left(\left[|11\rangle\langle 11|_f \otimes (B-\theta\mathbb{1})\otimes\mathbb{1}\right]|\chi\rangle\langle\chi|\right)\\ &= \sum_{z\in V_t}|\gamma_z|^2\sum_{x\le 0}\sum_g |\alpha_x(z',g)|^2|\kappa_g(z')|^2\left(\sum_j \lambda_j(B)|\sigma_j(z',g)|^2-\theta\right).\end{aligned} \tag{7.27}$$

Here, as before, γ_z represents the amplitudes of QPE on U_N, while $\alpha_x(z',g)$ represent the amplitudes of QPE over φ and K_z on eigenstate $|g\rangle$, and $\sigma_j(z',g)$ are the coefficients of the eigenstates of B. $\kappa_g(z')$ are coefficients of basis expansions of $|\nu\rangle$ in the energy eigenbasis of K_z.

We further define

$$\eta_{\max}(N,\varphi,\theta,t) := \max_{|\nu\rangle}\eta(N,\varphi,\theta,t,|\nu\rangle). \tag{7.28}$$

This is the maximum acceptance probability that the computation can output for a given N,φ,θ,t, maximised over all states $|\nu\rangle$.

Remark 7.2. *For $t \ge |N|$, $\eta(N,\varphi,\theta,t,|\nu\rangle)$ assumes its maximum for an eigenstate $|\nu\rangle = |g\rangle$ of K_N.*

Proof. For $t \ge |N|$, $\gamma_N = 1$ and all other $\gamma_z = 0$ for $z \ne N$. Hence

$$\begin{aligned}&\mathrm{argmax}_{|\nu\rangle}\,\eta(N,\varphi,\theta,t,|\nu\rangle)\\ =&\mathrm{argmax}_{|\nu\rangle}\sum_g |\kappa_g(N)|^2\left(\sum_{x\le 0}|\alpha_x(N,g)|^2\left(\sum_j\lambda_j(B)|\sigma_j(N,g)|^2-\theta\right)\right)\\ =&\mathrm{argmax}_{|\nu\rangle}\sum_g |\kappa_g(N)|^2\left(\sum_{x\le 0}|\alpha_x(N,g)|^2\right)\Gamma(g,\theta).\end{aligned}$$

where $|\nu\rangle = \sum_g \kappa_g(N)|g\rangle$ and $\Gamma(g,\theta) := \left(\sum_j \lambda_j(B)|\sigma_j(N,g)|^2-\theta\right)$. Thus $\eta(N,\varphi,\theta,t,|\nu\rangle)$ is a convex combination of the $\left(\sum_{x\le 0}|\alpha_x(N,g)|^2\right)\Gamma(g,\theta)$, and its maximum is assumed at an extremal point. The claim follows. □

Thus, for energy eigenstates $\{|g\rangle\}_g$ we can write:

$$\max_{|\nu\rangle} \eta(N,\varphi,\theta,t,|\nu\rangle) = \left(\sum_{x\leq 0} |\alpha_x(N,g)|^2 \left(\sum_j \lambda_j(B)|\sigma_j(N,g)|^2 - \theta\right)\right) \tag{7.29}$$

We see that $\eta_{\max}$ is dependent on $\Gamma(g,\theta) \coloneqq \left(\sum_j \lambda_j(B)|\sigma_j(N,g)|^2 - \theta\right)$ which is proportional to the expectation value of B on the input state. If this state is a low energy state, then $\langle\psi|\, B\,|\psi\rangle$ is promised to satisfy $\langle\psi|\, B\,|\psi\rangle \approx 1$ or ≈ 2.

We then go through roughly the same proof as the 1-CRT-PRM proof: we encode the QTM in a circuit-to-Hamiltonian mapping, then combine this with a tiling Hamiltonian and a negative energy Hamiltonian. Except now $\eta_{\max}$ depends on both φ (in the same way as it did in the 1-CRT-PRM proof) and θ, such that θ changes the value of $\eta_{\max}$ in a linear fashion. Since $\eta_{\max}$ determines the ground state energy of the circuit-to-Hamiltonian mapping, we see that the energy varies roughly linearly with θ provided φ and the input state are fixed.

Analogous to theorem 7.11, we obtain the following central result:

Theorem 7.12. *Let $K_N \in (\mathbb{C}^d)^{\otimes N}$ be a the Hamiltonian from corollary 6.1 and lemma 7.16 such that $\delta = \Omega(N^{-D})$ for some constant D. Define the order parameter $O_{A/B}$ acting on a const-sized subset of the lattice as in theorem 7.11. We can explicitly construct a Hamiltonian $H^\Lambda(N,\varphi,\theta) = \sum_{\langle i,j\rangle} h^N_{i,j}(\varphi,\theta) + \sum_{i\in\Lambda} h^N_i$ such that, in the infinite lattice size limit the following conditions hold. For any $\varphi \in [\lambda_{\min}(K_N) + \delta/3, \lambda_{\min}(K_N) + 2\delta/3]$ and supposing for the $\lambda_{j^*}(B) \in \{1,2\}$ which satisfies $|\langle\psi_0|\, B\,|\psi_0\rangle - \lambda_{j^*}(B)| = \mathrm{O}(1/P_1(N))$, where $|\psi_0\rangle$ is the ground state of K_N, then:*

- *if $\theta \geq \lambda_{j*}(B) - \frac{2}{5} + 1/P_2(N)$:*

  i H^Λ is gapped with spectral gap 1.

  ii product ground state.

  iii has order parameter expectation value $\langle O_{A/B}\rangle = 1$.

- *if $\theta \leq \lambda_{j*}(B) - \frac{1}{2} - 1/P_2(N)$:*

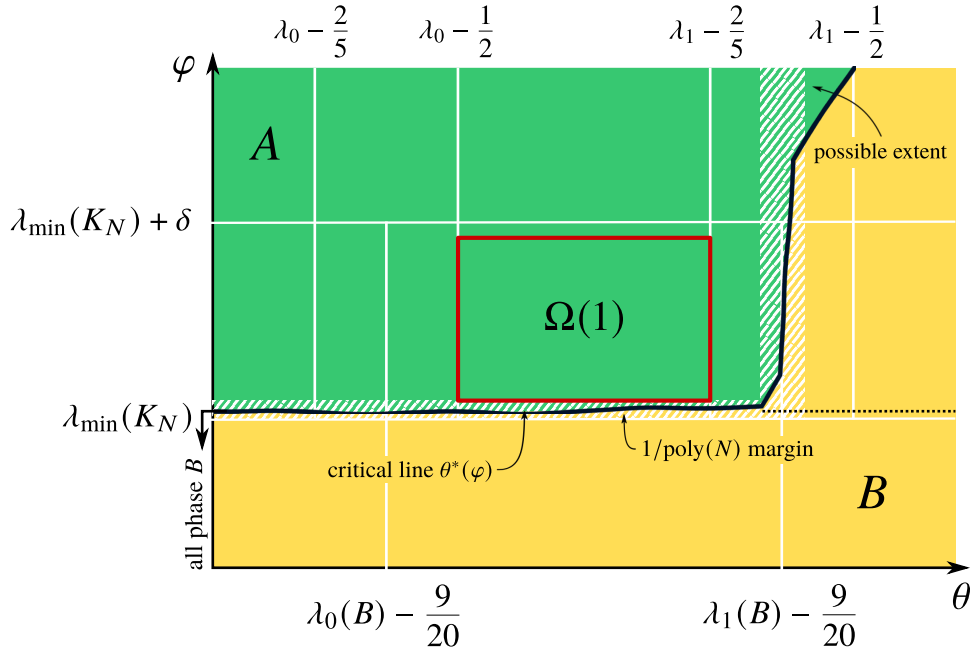

Figure 7.5: The YES case two-parameter phase diagram. The shaded white area shows an uncertainty region (of size $1/\mathrm{poly}\,N$); its *inner* extent indicates the minimal area circumscribed by the critical line $\theta^*(\varphi)$; it can be shown that the true critical line has precisely one critical point whenever φ is fixed and θ is varied, as well as vice versa; i.e., the critical line $\theta^*(\varphi)$ is a function, and monotonous, within an $\Omega(1)$ area of the phase space. It encompasses an $\Omega(1)$ area (given φ is scaled such that effecively $\delta = \Omega(1)$, as explained in corollary 7.4) of the phase space for which the system is guaranteed to be completely in phase A in this case. The location of the rectangle is efficiently computable relative to the point along the φ axis below which the system is completely in phase B, irrespective of θ. The NO case phase diagram is shown in fig. 7.6.

i H^Λ *is gapless.*

ii has a ground state with algebraically decaying correlations.

iii has order parameter expectation value $\langle O_{A/B} \rangle = 0$.

Furthermore, for any φ *in the given interval, this Hamiltonian has exactly one critical point in terms of* θ*, which we denote* θ^**. This occurs in the interval*

$$\theta^* \in \left[\lambda_{j*}(B) - \frac{1}{2} - 1/P_2(N), \lambda_{j*}(B) - \frac{2}{5} + 1/P_2(N)\right].$$

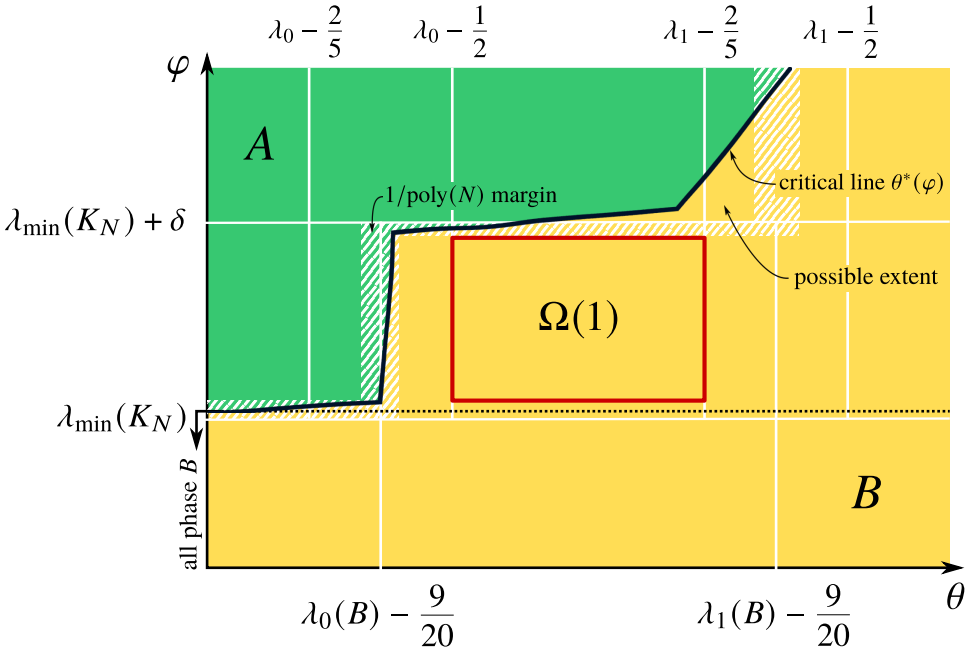

Figure 7.6: The NO case two-parameter phase diagram. The shaded white area shows an uncertainty region (of size $1/\mathrm{poly}\,N$); its *outer* extent indicates the maximal area circumscribed by the critical line $\theta^*(\varphi)$ It encompasses an $\Omega(1)$ area of the phase space for which the system is guaranteed to be completely in phase B in this case, and its location is efficiently computable relative to the point along the φ axis below which the system is completely in phase B, irrespective of θ. The YES case phase diagram is shown in fig. 7.5.

7.6.3 Reduction of ∀-TI-APX-SIM to 2-CRT-PRM

As per the 1-CRT-PRM case, we will find it useful to rescale the QPE process implemented φ as per remark 7.1. This has the effect of mapping $\varphi \longrightarrow N^{-D}\varphi$. This allow us to write the following corollary:

Corollary 7.4 (φ-Rescaled Hamiltonian)**.** *Theorem 7.12 holds for a modified Hamiltonian if* $\varphi \in [N^D(\lambda_{\min}(K_N)+\delta/3), N^D(\lambda_{\min}(K_N)+2\delta/3)]$ *such that* $N^D\delta = \Omega(1)$.

Having introduced this rescaling of φ, we now show that determining the phase transition point $\theta^*(\varphi)$ is $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$-hard by showing that determining $\theta^*(\varphi)$ for a specific $\mathrm{O}(1)$ interval of φ is gives the answer to a ∀-TI-APX-SIM instance.

Theorem 7.13 ($\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$-hardness)**.** *There is a polynomial time Turing reduction from ∀-TI-APX-SIM to 2-CRT-PRM, and hence 2-CRT-PRM is* $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$*-hard.*

Proof. We refer the reader to fig. 7.5 and fig. 7.6 to aid this proof.

From lemma 7.16 it is known that determining whether $\langle\psi|\, B\,|\psi\rangle > \beta$ or $\langle\psi|\, B\,|\psi\rangle < \alpha$ for states $|\psi\rangle \in S_\delta$ is $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$-hard. From theorem 7.12 we know that for all $\varphi \in [N^D(\lambda_{\min}(K_N) + \delta/3), N^D(\lambda_{\min}(K_N) + 2\delta/3)]$ it holds that the critical point $\theta^*(\varphi)$ is determined by whether $\langle\psi|\, B\,|\psi\rangle > \beta$ or $\langle\psi|\, B\,|\psi\rangle < \alpha$ for states $|\psi\rangle \in S_\delta$ and $\beta - \alpha = \Omega(1)$. □

Corollary 7.5. *2-CRT-PRM is* $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$*-complete.*

Proof. Hardness and containment follow from theorems 7.6 and 7.13, respectively. □

The two phase diagrams in the YES and NO cases are shown in figs. 7.5 and 7.6. It can then be verified that the Local-Global promise is satisfied by the Hamiltonian.

7.7 Discussion

Comparison to Undecidability and Uncomputability Results. We also take care to distinguish our results from the size driven quantum phase transitions [Bau+18a]. Here we are promised that in the thermodynamic we are always in a particular phase, but that the transition takes place at some uncomputably large lattice size. Our result differs significantly in that the spectral gap and phase are explicitly computable for some finite size lattice, and the system can be gapped or gapless in the thermodynamic limit.

We also emphasise the differences to the previous undecidability results [CPW15; Bau+18b; BCW19]. There are two key differences here: the promise of the global-local gap/phase means that the gap/phase are computable. The previous works prove results for systems with an infinite number of phase transitions[14], and although this is not totally unphysical, it is be no means a common property. The systems studied in this chapter contain either only a single phase transition, or a small finite number, which arguably better reflects the systems we see in nature.

[14] For various technical reasons it is not possible to define phases for the systems studied in [CPW15; Bau+18b], and we refer the reader to section 3.1 for a more extended discussion. Instead the authors remark that there are an uncountably infinite number of points where the system changes from being gapped to gapless.

Containment of the 1-CRT-PRM Case. In the 1-CRT-PRM result, we prove that determining the gap/phase 1-CRT-PRM for an $\Omega(1)$ region of the $\varphi \in [0,1]$ parameter space is $\mathsf{QMA}_{\mathsf{EXP}}$-hard. It is natural to ask whether we can prove containment in $\mathsf{QMA}_{\mathsf{EXP}}$ or at least $\mathrm{P}^{\mathsf{QMA}_{\mathsf{EXP}}[\mathrm{const}]}$? In particular, the fact that for the Hamiltonian constructed, the spectral gap is actually either $\geq 1/2$ or $\leq 1/\mathrm{poly}\, N$ suggests we might be able to distinguish the two cases by estimating the spectral gap to only constant precision in section 7.4 rather than the $1/\mathrm{poly}\, N$ precision we currently perform the algorithm to. However, both of the local-global algorithms (used for determining the spectral gap or the order parameter at a given point, respectively) require knowledge about the ground state to the relevant precision. In case we promise that the Hamiltonian's ground state energy can be resolved within a constant number of bits, containment in aforementioned stricter classes follows. Naturally, this leaves open the question whether an algorithm exists that can answer the spectral gap or order parameter problems to constant precision without knowing the ground state energy to the same precision.

Precise Variant. As mentioned below the definition of 1-CRT-PRM, definition 7.5, there is a natural "precise" variant, Precise-1-CRT-PRM, where we want to approximate the critical point to exponential precision. As explained in [Koh+20, Th. 4.1], for exponential precision one can allow the embedded computation to run for time $\exp\mathrm{poly}(L)$ in the size of the spin chain segment L; as such, one can extract exponentially many bits of the parameter N, and the distinction between translationally invariant and non-translationally-invariant models vanish. In this case, as detailed in [WBG20, Cor. 29&31], the APX-SIM variant is simply PSPACE-complete, by simulating a PSPACE computation within the history state. As such, it follows that the Precise-1-CRT-PRM problem—and by a similar argument also Precise-2-CRT-PRM—are PSPACE-hard. Containment in PSPACE, for a suitable definition of a local-global gap (that is now allowed to shrink exponentially in the system size), follows from a precise variant of Ambainis's algorithm to determine spectral gap; and because $\mathrm{P}^{\mathrm{PreciseQMA}} = \mathrm{P}^{\mathrm{PSPACE}} = \mathrm{PSPACE}$. [DGF20].

Open Questions. The following points are natural continuations of this line of work.

1. One major open question is to pin down the exact hardness of the 1-CRT-PRM problem, as we could show $\mathsf{QMA}_{\mathsf{EXP}}$ hardness, but only an upper bound of $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$, mostly because it is not obvious a priori how to reduce the algorithm that decides the Local-Global promise to constant precision with only constantly-many queries without negating them (as one would, otherwise, require co-$\mathsf{QMA}_{\mathsf{EXP}}$ queries). In all likelihood, a problem variant as formulated in 1-CRT-PRM with the promise parameters α and β provided as input is indeed contained in $\mathsf{QMA}_{\mathsf{EXP}}$, but that a more physically-motivated variant that just asks about the approximation to some precision (i.e., the non-decision variant; or a variant that just asks for the last bit of a poly-precision approximation to be 0 or 1) to be $\mathsf{P}^{\mathsf{QMA}_{\mathsf{EXP}}}$-hard.

2. The Knabe and Martingale methods for determining the spectral gap apply to frustration free Hamiltonians, but the Hamiltonians used to prove our results here are not. Can we prove a similar result to ours for frustration free Hamiltonians, for the classes $\mathrm{QMA}_1\mathrm{EXP}$ or similar? The class QMA_1 naturally characterises the Local Hamiltonian problem in where YES case correspond to frustration free Hamiltonians [Bra11; BT10].

Chapter 8

General Conclusions

For the past 20 years Hamiltonian complexity has been an enormously vibrant area of research, drawing on a wide range of techniques across physics, mathematics and computer science. We discuss some open problems related to the work in this thesis.

A key problem which has been noted for all of the results in thesis thesis, is the condition of naturalness and simplicity. The Hamiltonians studied here are generally highly unnatural, consisting of highly-tuned interactions on qudits with large local Hilbert space dimension. Is there any way of reducing the apparently unnaturalness of these Hamiltonians to find ones similar to which occur in the physical world? Some authors have argued that we should expect undecidability and uncomputability results to occur in almost all systems, with the exceptions of those which are extremely simple [Cue20], but how to prove this is not clear. The fact that (for example) proving gappedness appears to be difficult in all but the very constrained systems (e.g. frustration free) lends credence to this hypothesis. Indeed, how one should "measure" the complexity of inherent in a Hamiltonian in terms of its spatial and algebraic locality, local Hilbert space dimension, interaction terms, system size, interaction graph, etc is not clear, if it is possible at all.

Another major, related problem this author remains interested in is stability of complexity or uncomputability phenomena under perturbations. It has been shown that the spectral gap of a frustration-free Hamiltonian (satisfying some other conditions) is stable under local perturbations [MZ13]. If complexity and uncomputability results are to be relevant to real condensed matter systems — which

in reality have imperfections and external fields — we should expect them to still occur even after these local perturbations have been introduced. As far as this author is aware, it is an open problem whether any complexity/computability result retains its properties under local perturbations. We note, however, that such stability results are hard to come by and have only been proven for very limited systems [NYS18].

All the results in this paper have been proven for systems at zero temperature. Yet in reality we are always at some small but non-zero temperature. How hard is it to predict properties of such system, and in particular to any of the results in the thermodynamic limit survive? As mentioned in the introduction, classical hardness result are known for free energies/partition functions at non-zero temperature. Furthermore, for finite systems of size N, if one goes to temperatures $T = O(1/\text{poly}(N))$ then one will be close enough to the ground state such that predicting properties is QMA-hard. However, if we take a quantum system at non-zero temperature in the thermodynamic limit, its complexity is not clear. This problem is relevant for predicting phenomena such as phase transition, correlation functions etc. This is closely related to the quantum PCP conjecture, which if true, would imply that for a finite system predicting the properties of at non-zero temperature is QMA-hard [AAV13]. Some partial progress has been made on this question with the proof of the NLTS conjecture [ABN22].

Appendix A

Appendix

In this section we prove some properties of the Hamiltonian used in chapter 3. Importantly that it is a so-called "standard-form" Hamiltonian, and that this implies certain properties.

A.1 Standard Form Hamiltonians and the Clairvoyance Lemma

We begin with the following definition for a 1D chain of spins:

Definition A.1 (Standard Basis States, from Section 4.1 of [CPGW15a])**.** *Let the single site Hilbert space be $\mathcal{H} = \otimes_i \mathcal{H}_i$ and fix some orthonormal basis for the single site Hilbert space. Then a* Standard Basis State *for $\mathcal{H}^{\otimes L}$ are product states over the single site basis.*

We now define standard-form Hamiltonians – extending the definition from [CPGW15a]:

Definition A.2 (Standard-form Hamiltonian, from [Wat19], extended from [CPGW15a])**.** *We say that a Hamiltonian $H = H_{trans} + H_{pen} + H_{in} + H_{out}$ 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 $H_{trans,pen,in,out} = \sum_{i=1}^{L-1} h^{(i,i+1)}_{trans,pen,in,out}$, and $h_{trans,pen,in,out}$ satisfy the following conditions:*

1. *$h_{trans} \in \mathcal{B}\left((\mathbb{C}^C \otimes \mathbb{C}^Q)^{\otimes 2}\right)$ 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* $|ab\rangle |\psi^i_{abcd}\rangle \to |cd\rangle U_{abcd} |\psi^i_{abcd}\rangle$ *for all* i*. There is then a corresponding term in the Hamiltonian of the form* $(|cd\rangle \otimes U_{abcd} - |ab\rangle)(\langle cd| \otimes U^\dagger_{abcd} - \langle ab|)$.

2. $h_{pen} \in \mathcal{B}\left((\mathbb{C}^C \otimes \mathbb{C}^Q)^{\otimes 2}\right)$ *is a sum of penalty terms which act non-trivially only on* $(\mathbb{C}^C)^{\otimes 2}$ *and are diagonal in the standard basis, such that* $h_{pen} = \sum_{(ab)\ Illegal} |ab\rangle_C \langle ab| \otimes \mathbb{1}_Q$*, where* (ab) *are members of a disallowed/illegal subspace.*

3. $h_{in} = \sum_{ab} |ab\rangle \langle ab|_C \otimes \Pi_{ab}$*, where* $|ab\rangle \langle ab|_C \in (\mathbb{C}^C)^{\otimes 2}$ *is a projector onto* $(\mathbb{C}^C)^{\otimes 2}$ *basis states, and* $\Pi^{(in)}_{ab} \in (\mathbb{C}^Q)^{\otimes 2}$ *are orthogonal projectors onto* $(\mathbb{C}^Q)^{\otimes 2}$ *basis states.*

4. $h_{out} = |xy\rangle \langle xy|_C \otimes \Pi_{xy}$*, where* $|xy\rangle \langle xy|_C \in (\mathbb{C}^C)^{\otimes 2}$ *is a projector onto* $(\mathbb{C}^C)^{\otimes 2}$ *basis states, and* $\Pi^{(in)}_{xy} \in (\mathbb{C}^Q)^{\otimes 2}$ *are orthogonal projectors onto* $(\mathbb{C}^Q)^{\otimes 2}$ *basis states.*

We note that although h_{out} and h_{in} have essentially the same form, they will play a conceptually different role.

Lemma A.1. H_{QTM} *is a standard form Hamiltonian.*

Proof. Comparing with definition A.2, we see that all terms fall into one of the four classifications, and hence it is standard form. □

We now introduce the following definition.

Definition A.3 (Legal and Illegal Pairs and States, from [CPGW15a])**.** *The pair* ab *is an* illegal pair *if the penalty term* $|ab\rangle \langle ab|_C \otimes \mathbb{1}_Q$ *is in the support of the* H_{pen}

component of the Hamiltonian. If a pair is not illegal, it is legal. We call a standard basis state legal *if it does not contain any illegal pairs, and illegal otherwise.*

Then the following is a straightforward extension of Lemma 42 of [CPGW15a] with H_{in} and H_{out} terms included.

Lemma A.2 (Invariant subspaces, extended from Lemma 42 of [CPGW15a])**.** *Let H_{trans}, H_{pen}, H_{in} and H_{out} define a standard-form Hamiltonian as defined in definition A.2. 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 $|ab\rangle_{CD}|\psi\rangle \to |cd\rangle_{CD} U_{abcd}|\psi\rangle$ acts on $\mathcal{H}_C$ by restriction to $(\mathbb{C}^C)^{\otimes 2}$, i.e. it acts as $ab \to cd$). 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 $H = H_{pen} + H_{trans} + H_{in} + H_{out}$ where $\mathcal{K}_{S_i}$ is spanned by S_i.*

Lemma A.3 (Clairvoyance Lemma, extended from Lemma 43 of [CPGW15a])**.** *Let $H = H_{trans} + H_{pen} + H_{in} + H_{out}$ be a standard-form Hamiltonian, as defined in definition A.2, and let $\mathcal{K}_S$ be defined as in Lemma A.2. Let $\lambda_0(\mathcal{K}_S)$ denote the minimum eigenvalue of the restriction $H|_{\mathcal{K}_S \otimes \mathcal{H}_Q}$ of $H = H_{trans} + H_{pen} + H_{in} + H_{out}$ 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:

1. *All legal standard basis states for $\mathcal{H}_C$ are contained in $\mathcal{W}$.*
2. *$\mathcal{W}$ is closed with respect to the transition rules.*
3. *At most one transition rule applies in each direction to any state in $\mathcal{W}$. Furthermore, there exists an ordering on the states in each S such that the forwards transition (if it exists) is from $|t\rangle \to |t+1\rangle$ and the backwards transition (if it exists) is $|t\rangle \to |t-1\rangle$.*
4. *For any subset $S \subseteq \mathcal{W}$ that contains only legal states, there exists at least one state to which no backwards transition applies and one state to which no forwards transition applies. Furthermore, the unitaries associated with the transition $|t\rangle \to |t+1\rangle$ are $U_t = \mathbb{1}_Q$, for $0 \le t \le T_{init} - 1$ and $T_{init} < T$, and that*

the final state $|T\rangle$ is detectable by a 2-local projector acting only on nearest neighbour qudits.

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

1. *S contains only illegal states, and $H|_{\mathcal{K}_S\otimes\mathcal{H}_Q} \geq \mathbb{1}$.*

2. *S contains both legal and illegal states, and*

$$W^\dagger H|_{\mathcal{K}_S\otimes\mathcal{H}_Q} W \geq \bigoplus_i \left(\Delta^{(|S|)} + \sum_{|k\rangle\in K_i} |k\rangle\langle k|\right) \tag{A.1}$$

   where $\sum_{|k\rangle\in K_i} |k\rangle\langle k| := H_{pen}|_{\mathcal{K}_S\otimes\mathcal{H}_Q}$ and K_i is some non-empty set of basis states and W is some unitary.

3. *S contains only legal states, then there exists a unitary $R = W(\mathbb{1}_C \otimes (X \oplus Y)_Q)$ that puts $H|_{\mathcal{K}_S\otimes\mathcal{H}_Q}$ in the form*

$$R^\dagger H|_{\mathcal{K}_S\otimes\mathcal{H}_Q} R = \begin{pmatrix} H_{aa} & H_{ab} \\ H_{ab}^\dagger & H_{bb} \end{pmatrix}, \tag{A.2}$$

   where, defining $G := \mathrm{supp}\left(\sum_{t=0}^{T_{init}-1} \Pi_t^{(in)}\right)$ and $s := \dim G$,

   - *$X : G \to G$.*
   - *$Y : G^c \to G^c$.*
   - *H_{aa} is an $s \times s$ matrix.*
   - *$H_{aa}, H_{bb} \geq 0$ and are rank r_a, r_b respectively.*
   - *H_{aa} has the form*

$$H_{aa} = \bigoplus_i \left(\Delta^{(|S|)} + \alpha_i \,||S|-1\rangle\langle|S|-1|\right) + \sum_{t=0}^{T_{init}-1} |t\rangle\langle t| \otimes X^\dagger \Pi_t|_G X. \tag{A.3}$$

- H_{bb} *is a tridiagonal, stoquastic matrix of the form*

$$H_{bb} = \bigoplus_i (\Delta^{(|S|)} + \beta_i \, ||S|-1\rangle \langle |S|-1|). \tag{A.4}$$

- $H_{ab} = H_{ba}$ *is a real, negative diagonal matrix with rank* $\min\{r_a, r_b\}$.

$$H_{ab} = H_{ba} = \bigoplus_i \gamma_i \, ||S|-1\rangle \langle |S|-1| \, . \tag{A.5}$$

where either we get pairings between the blocks such that

$$\begin{pmatrix} \alpha_i & \gamma_i \\ \gamma_i & \beta_i \end{pmatrix} = \begin{pmatrix} 1-\mu_i & -\sqrt{\mu_i(1-\mu_i)} \\ -\sqrt{\mu_i(1-\mu_i)} & \mu_i \end{pmatrix} \quad or \quad \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \tag{A.6}$$

for $0 \leq \mu_i \leq 1$, *or we get unpaired values of* $\alpha_i = 0, 1$ *or* $\beta_i = 0, 1$ *for which we have no associated value of* γ_i.

Bibliography

[AR+20] Houssam Abdul-Rahman, Marius Lemm, Angelo Lucia, Bruno Nachtergaele, and Amanda Young. "A class of two-dimensional AKLT models with a gap". In: *Analytic Trends in Mathematical Physics Contemporary Mathematics* (2020), 1–21.

[Aff+87] Ian Affleck, Tom Kennedy, Elliott H. Lieb, and Hal Tasaki. "Rigorous results on valence-bond ground states in antiferromagnets". In: *Physical Review Letters* 59.7 (1987), pp. 799–802.

[AAV13] Dorit Aharonov, Itai Arad, and Thomas Vidick. "Guest column: "The Quantum PCP Conjecture"". In: *ACM SIGACT News* 44.2 (2013), 47–79.

[Aha+07] Dorit Aharonov, Daniel Gottesman, Sandy Irani, and Julia Kempe. "The Power of Quantum Systems on a Line". In: *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS07)* (2007).

[AI21] Dorit Aharonov and Sandy Irani. "Hamiltonian Complexity in the Thermodynamic Limit". In: *arXiv e-prints*, arXiv:2107.06201 (July 2021), arXiv:2107.06201. arXiv: 2107.06201 [quant-ph].

[AI22] Dorit Aharonov and Sandy Irani. "Hamiltonian Complexity in the Thermodynamic Limit". In: *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2022. Rome, Italy: Association for Computing Machinery, 2022, 750–763.

[Alt79] H. Alt. "Square rooting is as difficult as multiplication". In: *Computing* 21.3 (1979), pp. 221–232.

[Amb14] Andris Ambainis. “On Physical Problems that are Slightly More Difficult than QMA”. In: *2014 IEEE 29th Conference on Computational Complexity (CCC)* (2014).

[And08] P. W. Anderson. “More is Different: Broken Symmetry and the Nature of the Hierarchical Structure of Science”. In: *Emergence* (2008), 221–230.

[Ans20] Anurag Anshu. “Improved local spectral gap thresholds for lattices of finite size”. In: *Physical Review B* 101.16 (2020).

[ABN22] Anurag Anshu, Nikolas Breuckmann, and Chinmay Nirkhe. “NLTS Hamiltonians from good quantum codes”. In: *arXiv e-prints*, arXiv:2206.13228 (June 2022), arXiv:2206.13228. arXiv: 2206.13228 [quant-ph].

[AB10] Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. Cambridge Univ. Press, 2010.

[Bar82] Francisco Barahona. “On the computational complexity of Ising spin glass models”. In: *Journal of Physics A: Mathematical and General* 15.10 (1982), p. 3241.

[Bar18] Alexander Barvinok. *Combinatorics and Complexity of Partition Functions*. Springer International Publishing, 2018.

[Bau20] Johannes Bausch. “Perturbation Gadgets: Arbitrary Energy Scales from a Single Strong Interaction”. In: *Annales Henri Poincaré* 21.1 (2020), pp. 81–114. arXiv: 1810.00865.

[Bau19] Johannes Bausch. “Perturbation gadgets: Arbitrary Energy Scales from a single strong interaction”. In: *Annales Henri Poincaré* 21.1 (2019), 81–114.

[BC18a] Johannes Bausch and Elizabeth Crosson. “Analysis and limitations of modified circuit-to-Hamiltonian constructions”. In: *Quantum* 2 (2018), p. 94. arXiv: 1609.08571.

[BC18b] Johannes Bausch and Elizabeth Crosson. "Analysis and limitations of modified circuit-to-Hamiltonian constructions". In: *Quantum* 2 (2018), p. 94. arXiv: 1609.08571.

[Bau+18a] Johannes Bausch, Toby S. Cubitt, Angelo Lucia, David Perez-Garcia, and Michael M. Wolf. "Size-driven quantum phase transitions". In: *Proceedings of the National Academy of Sciences* 115.1 (2018), pp. 19–23. arXiv: 1512.05687.

[BCW19] Johannes Bausch, Toby S. Cubitt, and James D. Watson. "Uncomputability of Phase Diagrams". In: *arXiv e-prints*, arXiv:1910.01631 (2019), arXiv:1910.01631. arXiv: 1910.01631 [quant-ph].

[BCW21] Johannes Bausch, Toby S. Cubitt, and James D. Watson. "Uncomputability of phase diagrams". In: *Nature Communications* 12.1 (2021).

[Bau+18b] Johannes Bausch, Toby Cubitt, Angelo Lucia, and David Perez-Garcia. "Undecidability of the Spectral Gap in One Dimension". In: (Oct. 2018). arXiv: 1810.01858.

[BCO17] Johannes Bausch, Toby Cubitt, and Maris Ozols. "The Complexity of Translationally-Invariant Spin Chains with Low Local Dimension". In: *Annales Henri Poincaré* (2017), p. 52. arXiv: 1605.01718.

[BP17a] Johannes Bausch and Stephen Piddock. "The complexity of translationally invariant low-dimensional spin lattices in 3D". In: *Journal of Mathematical Physics* 58.11 (2017), p. 111901. arXiv: 1702.08830.

[BP17b] Johannes Bausch and Stephen Piddock. "The complexity of translationally invariant low-dimensional spin lattices in 3D". In: *Journal of Mathematical Physics* 58.11 (2017), p. 111901. arXiv: 1702.08830.

[Bei91] Richard Beigel. "Bounded queries to SAT and the Boolean hierarchy". In: *Theoretical Computer Science* 84.2 (1991), pp. 199 –223.

[Ben90] C H Bennett. "Undecidable Dynamics". In: *Nature* 346 (1990), 606–607.

[Ber66] Robert Berger. *The Undecidability of the Domino Problem*. American Mathematical Soc., 1966, p. 72.

[BV97] Ethan Bernstein and Umesh Vazirani. "Quantum Complexity Theory". In: *SIAM Journal on Computing* 26.5 (Oct. 1997), pp. 1411–1473.

[Ber+06] Dominic W. Berry, Graeme Ahokas, Richard Cleve, and Barry C. Sanders. "Efficient Quantum Algorithms for Simulating Sparse Hamiltonians". In: *Communications in Mathematical Physics* 270.2 (2006), 359–371.

[BCK15] Dominic W. Berry, Andrew M. Childs, and Robin Kothari. "Hamiltonian Simulation with Nearly Optimal Dependence on all Parameters". In: *2015 IEEE 56th Annual Symposium on Foundations of Computer Science* (2015).

[Bet31] H Bethe. "Zur Theorie der Metalle". In: *Zeitschrift für Physik* 71.3–4 (1931), pp. 205–226.

[BS99] Bibhas Bhattacharyya and Shreekantha Sil. "The Hubbard model with bond-charge interaction on a triangular lattice: a renormalization group study". In: *Journal of Physics: Condensed Matter* 11.17 (1999), pp. 3513–3523.

[BL08] Jacob D. Biamonte and Peter J. Love. "Realizable Hamiltonians for universal adiabatic quantum computers". In: *Phys. Rev. A* 78 (1 2008), p. 012352.

[Boh+19] Thomas C. Bohdanowicz, Elizabeth Crosson, Chinmay Nirkhe, and Henry Yuen. "Good Approximate Quantum LDPC Codes from Spacetime Circuit Hamiltonians". In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2019. Phoenix, AZ, USA: Association for Computing Machinery, 2019, 481–490.

[Bra08] Fernando G. S. L. Brandao. "Entanglement Theory and the Quantum Simulation of Many-Body Physics". In: (Oct. 2008).

[BK02] Sergey B. Bravyi and Alexei Yu. Kitaev. “Fermionic Quantum Computation”. In: *Annals of Physics* 298.1 (2002), 210–226.

[Bra11] Sergey Bravyi. “Efficient algorithm for a quantum analogue of 2-SAT”. In: *Contemporary Mathematics Cross Disciplinary Advances in Quantum Computing* (2011), 33–48.

[BBT06] Sergey Bravyi, Arvid J. Bessen, and Barbara M. Terhal. “Merlin-Arthur Games and Stoquastic Complexity”. In: *arXiv e-prints*, quant-ph/0611021 (Nov. 2006), quant–ph/0611021. arXiv: quant-ph/0611021 [quant-ph].

[Bra+21] Sergey Bravyi, Anirban Chowdhury, David Gosset, and Pawel Wocjan. “On the complexity of quantum partition functions”. In: *arXiv e-prints*, arXiv:2110.15466 (Oct. 2021), arXiv:2110.15466. arXiv: 2110.15466 [quant-ph].

[BG17] Sergey Bravyi and David Gosset. “Complexity of Quantum Impurity Problems”. In: *Communications in Mathematical Physics* 356.2 (2017), 451–500.

[BG15] Sergey Bravyi and David Gosset. “Gapped and gapless phases of frustration-free spin- /1 2 chains”. In: *Journal of Mathematical Physics* 56.6, 061902 (June 2015), p. 061902. arXiv: 1503.04035 [quant-ph].

[BT10] Sergey Bravyi and Barbara Terhal. “Complexity of Stoquastic Frustration-Free Hamiltonians”. In: *SIAM Journal on Computing* 39.4 (2010), 1462–1485.

[BP53] Ernst Carl Gerlach Stueckelberg de Breidenbach and Andreas Petermann. “La normalisation des constantes dans la théorie des quantaNormalization of constants in the quanta theory”. In: *Helv. Phys. Acta* 26 (1953), pp. 499–520.

[BT14] Nikolas P Breuckmann and Barbara M Terhal. “Space-time circuit-to-Hamiltonian construction and its applications”. In: *Journal of Physics A: Mathematical and Theoretical* 47.19 (2014), p. 195304.

[BG19] Anne Broadbent and Alex B. Grilo. “QMA-hardness of Consistency of Local Density Matrices with Applications to Quantum Zero-Knowledge”. In: *arXiv e-prints*, arXiv:1911.07782 (Nov. 2019), arXiv:1911.07782. arXiv: 1911.07782 [quant-ph].

[BFS11] Brielin Brown, Steven T. Flammia, and Norbert Schuch. “Computational Difficulty of Computing the Density of States”. In: *Physical Review Letters* 107.4 (2011).

[BZJ85] E. Brézin and J. Zinn-Justin. “Finite size effects in phase transitions”. In: *Nuclear Physics B* 257 (1985), pp. 867–893.

[Bus88] Jonathan F. Buss. “Relativized alternation and space-bounded computation”. In: *Journal of Computer and System Sciences* 36.3 (1988), 351–378.

[CLN18] Libor Caha, Zeph Landau, and Daniel Nagaj. “Clocks in Feynman’s computer and Kitaev’s local Hamiltonian: Bias, gaps, idling, and pulse tuning”. In: *Physical Review A* 97.6 (2018), p. 062306. arXiv: 1712.07395.

[CN15] Yudong Cao and Daniel Nagaj. “Perturbative gadgets without strong interactions”. In: *Quantum Information and Computation* 15.13 and 14 (2015), 1197–1222.

[Car96] John Cardy. *Scaling and renormalization in statistical physics*. Cambridge University Press, 1996.

[Cha75] Gregory J Chaitin. “A theory of program size formally identical to information theory”. In: *Journal of the ACM (JACM)* 22.3 (1975), pp. 329–340.

[CM17] Natalia Chepiga and Frédéric Mila. "Excitation spectrum and density matrix renormalization group iterations". In: *Phys. Rev. B* 96 (5 2017), p. 054425.

[CGW13] A. M. Childs, D. Gosset, and Z. Webb. "Universal Computation by Multiparticle Quantum Walk". In: *Science* 339.6121 (2013), 791–794.

[CGW16] Andrew M. Childs, David Gosset, and Zak Webb. "Complexity of the XY antiferromagnet at fixed magnetization". In: *Quantum Information and Computation* (2016), 1–18.

[CS19] Andrew M. Childs and Yuan Su. "Nearly Optimal Lattice Simulation by Product Formulas". In: *Physical Review Letters* 123.5 (2019), p. 050503.

[CPD21] Karthik Chinni, Pablo M. Poggi, and Ivan H. Deutsch. "Effect of chaos on the simulation of quantum critical phenomena in analog quantum simulators". In: *Phys. Rev. Research* 3 (3 2021), p. 033145.

[CBH19] Alex Churchill, Stella Biderman, and Austin Herrick. "Magic: The Gathering is Turing Complete". In: *arXiv e-prints*, arXiv:1904.09828 (Mar. 2019), arXiv:1904.09828. arXiv: 1904.09828 [cs.AI].

[CK92] László P. Csernai and Joseph I. Kapusta. "Dynamics of the QCD phase transition". In: *Physical Review Letters* 69.5 (1992), 737–740.

[CPGW15a] T. S. Cubitt, D. Perez-Garcia, and M. M. Wolf. "Undecidability of the spectral gap". In: (2015). arXiv: 1502.04573 [quant-ph].

[CMP18] Toby S. Cubitt, Ashley Montanaro, and Stephen Piddock. "Universal Quantum Hamiltonians". In: *Proceedings of the National Academy of Sciences* 115.38 (2018), 9497–9502.

[CPGW15b] Toby S. Cubitt, David Perez-Garcia, and Michael M. Wolf. "Undecidability of the spectral gap". In: *Nature* 528.7581 (Dec. 2015), pp. 207–211. arXiv: 1502.04573.

[CPW15] Toby S. Cubitt, David Perez-Garcia, and Michael M. Wolf. "Undecidability of the spectral gap". In: *Nature* 528.7581 (2015), pp. 207–211. arXiv: 1502.04573 [quant-ph].

[CM13] Toby Cubitt and Ashley Montanaro. "Complexity classification of local Hamiltonian problems". In: *arXiv e-prints*, arXiv:1311.3161 (Nov. 2013), arXiv:1311.3161. arXiv: 1311.3161 [quant-ph].

[Cue20] Gemma De las Cuevas. *Universality everywhere implies undecidability everywhere by Gemma de las Cuevas*. 2020.

[DB19] Alexander M. Dalzell and Fernando G. Brandão. "Locally accurate mps approximations for ground states of one-dimensional gapped local Hamiltonians". In: *Quantum* 3 (2019), p. 187.

[DT91] P. H. Damgaard and G. Thorleifsson. "Chaotic renormalization-group trajectories". In: *Phys. Rev. A* 44 (4 1991), pp. 2738–2741.

[DN05a] Christopher M. Dawson and Michael A. Nielsen. "The Solovay-Kitaev algorithm". In: (2005). arXiv: 0505030 [quant-ph].

[DN05b] Christopher M. Dawson and Michael A. Nielsen. "The Solovay-Kitaev algorithm". In: (2005). arXiv: 0505030 [quant-ph].

[Den+02] Eric Dennis, Alexei Kitaev, Andrew Landahl, and John Preskill. "Topological quantum memory". In: *Journal of Mathematical Physics* 43.9 (2002), 4452–4505.

[DEE99] Bernard Derrida, Jean-Pierre Eckmann, and Ayse Erzan. "Renormalisation groups with periodic and aperiodic orbits". In: *Journal of Physics A: Mathematical and General* 16 (Jan. 1999), p. 893.

[DGF20] Abhinav Deshpande, Alexey V. Gorshkov, and Bill Fefferman. "The importance of the spectral gap in estimating ground-state energies". In: *arXiv e-prints*, arXiv:2007.11582 (July 2020), arXiv:2007.11582. arXiv: 2007.11582 [quant-ph].

[DJ92] D. Deutsch and R. Jozsa. “Rapid solution of problems by quantum computation”. In: *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences* 439.1907 (1992), 553–558.

[Eba+21] Sepehr Ebadi et al. “Quantum phases of matter on a 256-atom programmable quantum simulator”. In: *Nature* 595.7866 (2021), 227–232.

[ER85] J.-P. Eckmann and D. Ruelle. “Ergodic theory of chaos and strange attractors”. In: *The Theory of Chaotic Attractors* (1985), 273–312.

[Elh+16] Serdar Elhatisari et al. “Nuclear Binding Near a Quantum Phase Transition”. In: *Phys. Rev. Lett.* 117 (13 2016), p. 132501.

[FS93] G Fath and J Solyom. “Solitonic excitations in the Haldane phase of an S=1 chain”. In: *Journal of Physics: Condensed Matter* 5.48 (1993), 8983–8998.

[FL16] Bill Fefferman and Cedric Yen-Yu Lin. “A Complete Characterization of Unitary Quantum Space”. In: *Leibniz International Proceedings in Informatics (LIPIcs)*. 9th Innovations in Theoretical Computer Science Conference, 2016. arXiv: 1604.01384.

[Fey84] Richard P. Feynman. “Quantum mechanical computers”. In: *Conference on Lasers and Electro-Optics* (1984).

[Fey82] Richard P. Feynman. “Simulating physics with computers”. In: *International Journal of Theoretical Physics* 21.6-7 (1982), 467–488.

[For94] Lance Fortnow. “The Role of Relativization in Complexity Theory”. In: *Bulletin of the European Association for Theoretical Computer Science* 52 (1994), pp. 52–229.

[Fra17] Fabio Franchini. *An Introduction to Integrable Techniques for One-Dimensional Quantum Systems*. Vol. 940. Lecture Notes in Physics. Cham: Springer International Publishing, 2017.

[GSMW13] Artur Garcia-Saez, Valentin Murg, and Tzu-Chieh Wei. "Spectral gaps of Affleck-Kennedy-Lieb-Tasaki Hamiltonians using tensor network methods". In: *Physical Review B* 88.24 (2013).

[GML54] M. Gell-Mann and F. E. Low. "Quantum Electrodynamics at Small Distances". In: *Phys. Rev.* 95 (5 1954), pp. 1300–1312.

[Gha13] Sevag Gharibian. "Approximation, proof systems, and correlations in a quantum world". Available at arXiv.org quant-ph/1301.2632. PhD thesis. University of Waterloo, 2013.

[GPY20] Sevag Gharibian, Stephen Piddock, and Justin Yirka. "Oracle Complexity Classes and Local Measurements on Physical Hamiltonians". In: *37th International Symposium on Theoretical Aspects of Computer Science (STACS 2020)*. Ed. by Christophe Paul and Markus Bläser. Vol. 154. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2020, 20:1–20:37.

[GS18] Sevag Gharibian and Jamie Sikora. "Ground State Connectivity of Local Hamiltonians". In: *ACM Trans. Comput. Theory* 10.2 (Apr. 2018).

[GY19] Sevag Gharibian and Justin Yirka. "The complexity of simulating local measurements on quantum systems". In: *Quantum* 3 (Sept. 2019), p. 189.

[GK82] V. L. Ginzburg and D. A. Kirzhnits. "The Problem of High-Temperature Superconductivity (General Review)". In: *High-Temperature Superconductivity* (1982), 1–40.

[Gol06] Oded Goldreich. "On Promise Problems: A Survey". In: *Lecture Notes in Computer Science Theoretical Computer Science* (2006), 254–290.

[GC18] Carlos E. González-Guillén and Toby S. Cubitt. "History-state Hamiltonians are critical". In: *arXiv e-prints*, arXiv:1810.06528 (Oct. 2018), arXiv:1810.06528. arXiv: 1810.06528 [quant-ph].

[GMV16] David Gosset, Jenish C. Mehta, and Thomas Vidick. "QCMA hardness of ground space connectivity for commuting Hamiltonians". In: *arXiv e-prints*, arXiv:1610.03582 (Oct. 2016), arXiv:1610.03582. arXiv: 1610.03582 [cs.CC].

[GM16] David Gosset and Evgeny Mozgunov. "Local gap threshold for frustration-free spin systems". In: *Journal of Mathematical Physics* 57.9 (2016), p. 091901. eprint: https://doi.org/10.1063/1.4962337.

[GTV15] David Gosset, Barbara M. Terhal, and Anna Vershynina. "Universal Adiabatic Quantum Computation via the Space-Time Circuit-to-Hamiltonian Construction". In: *Phys. Rev. Lett.* 114 (14 2015), p. 140501.

[GI09] Daniel Gottesman and Sandy Irani. "The quantum and classical complexity of translationally invariant tiling and Hamiltonian problems". In: *Foundations of Computer Science, 2009. FOCS'09. 50th Annual IEEE Symposium on*. IEEE. 2009, pp. 95–104.

[GP83] Peter Grassberger and Itamar Procaccia. "Measuring the strangeness of strange attractors". In: *Physica D: Nonlinear Phenomena* 9.1-2 (1983), 189–208.

[Had75] F. Hadlock. "Finding a Maximum Cut of a Planar Graph in Polynomial Time". In: *SIAM Journal on Computing* 4.3 (1975), 221–225.

[HB88] C. J. Hamer and Michael N. Barber. "Finite-Size Scaling in Hamiltonian Field Theory: The Transverse Ising Model". In: *Finite-Size Scaling Current Physics–Sources and Comments* (1988), 258–263.

[HHL09] Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd. "Quantum Algorithm for Linear Systems of Equations". In: *Phys. Rev. Lett.* 103 (15 2009), p. 150502.

[Har+93] Juris Hartmanis, Richard Chang, Jim Kadin, and Stephen G. Mitchell. "Some Observations about Relativization of Space Bounded Computations". In: *Current Trends in Theoretical Computer Science* (1993), 423–434.

[Has07a] M B Hastings. "An area law for one-dimensional quantum systems". In: *Journal of Statistical Mechanics: Theory and Experiment* 2007.08 (2007).

[Has07b] M. B. Hastings. "Entropy and entanglement in Quantum Ground States". In: *Physical Review B* 76.3 (2007).

[HWB19] Oscar Higgott, Daochen Wang, and Stephen Brierley. "Variational Quantum Computation of Excited States". In: *Quantum* 3 (2019), p. 156.

[HW94] Lloyd C. L. Hollenberg and N. S. Witte. "General nonperturbative estimate of the energy density of lattice Hamiltonians". In: *Physical Review D* 50.5 (1994), 3382–3386.

[Hub63] J Hubbard. "Electron correlations in narrow energy bands". In: *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* 276.1365 (1963), 238–257.

[Jan07] Dominik Janzing. "Spin-1/2 particles moving on a two-dimensional lattice with nearest-neighbor interactions can realize an autonomous quantum computer". In: *Phys. Rev. A* 75 (1 2007), p. 012307.

[Jon+19] Tyson Jones, Suguru Endo, Sam McArdle, Xiao Yuan, and Simon C. Benjamin. "Variational quantum algorithms for discovering Hamiltonian spectra". In: *Phys. Rev. A* 99 (6 2019), p. 062304.

[JGL10] Stephen P. Jordan, David Gosset, and Peter J. Love. "Quantum-Merlin-Arthur-complete problems for stoquastic Hamiltonians and Markov matrices". In: *Physical Review A* 81.3 (2010).

[JP79] R. Jullien and P. Pfeuty. "Zero-temperature renormalization-group method for quantum systems. II. Isotropic $X-Y$ model in a transverse field in one dimension". In: *Phys. Rev. B* 19 (9 1979), pp. 4646–4652.

[Jul+78] R. Jullien, P. Pfeuty, J. N. Fields, and S. Doniach. "Zero-temperature renormalization method for quantum systems. I. Ising model in a transverse field in one dimension". In: *Phys. Rev. B* 18 (7 1978), pp. 3568–3578.

[Kad66] Leo P. Kadanoff. "Scaling laws for ising models near T_c". In: *Physics Physique Fizika* 2 (6 1966), pp. 263–272.

[KL80] Richard M. Karp and Richard J. Lipton. "Some connections between nonuniform and uniform complexity classes". In: *Proceedings of the twelfth annual ACM symposium on Theory of computing - STOC 80* (1980).

[KL18] Michael J Kastoryano and Angelo Lucia. "Divide and conquer method for proving gaps of frustration free Hamiltonians". In: *Journal of Statistical Mechanics: Theory and Experiment* 2018.3 (2018), p. 033105.

[KR03a] J. Kempe and O. Regev. "3-Local Hamiltonian is QMA-complete". In: *Quantum Information and Computation* 3.3 (2003), 258–264.

[KKR06] Julia Kempe, Alexei Kitaev, and Oded Regev. "The Complexity of the Local Hamiltonian Problem". In: *SIAM Journal on Computing* 35.5 (2006), 1070–1097.

[KR03b] Julia Kempe and Oded Regev. "3-Local Hamiltonian is QMA-complete". In: *Quantum Computation and Information* 3.3 (2003), pp. 258–264. arXiv: 0302079 [quant-ph].

[KSV02] A. Kitaev, A. Shen, and M. Vyalyi. *Classical and Quantum Computation*. American Mathematical Society, 2002.

[Kna88] Stefan Knabe. "Energy gaps and elementary excitations for certain VBS-quantum antiferromagnets". In: *Journal of Statistical Physics* 52.3-4 (1988), 627–638.

[Koh+20] Tamara Kohler, Stephen Piddock, Johannes Bausch, and Toby Cubitt. "Translationally-Invariant Universal Quantum Hamiltonians in 1D". In: *arXiv e-prints*, arXiv:2003.13753 (Mar. 2020), arXiv:2003.13753. arXiv: 2003.13753 [quant-ph].

[Koh+20] Tamara Kohler, Stephen Piddock, Johannes Bausch, and Toby Cubitt. "Translationally-Invariant Universal Quantum Hamiltonians in 1D". In: (2020). arXiv: 2003.13753.

[LL76] Richard E. Ladner and Nancy A. Lynch. "Relativization of questions about log space computability". In: *Mathematical Systems Theory* 10.1 (1976), 19–32.

[LVV15] Zeph Landau, Umesh Vazirani, and Thomas Vidick. "A polynomial time algorithm for the ground state of one-dimensional gapped local Hamiltonians". In: *Nature Physics* 11.7 (2015), 566–569.

[Lem20] Marius Lemm. "Finite-size criteria for spectral gaps in -dimensional quantum spin systems". In: *Analytic Trends in Mathematical Physics Contemporary Mathematics* (2020), 121–132.

[LM19] Marius Lemm and Evgeny Mozgunov. "Spectral gaps of frustration-free spin systems with boundary". In: *Journal of Mathematical Physics* 60.5 (2019), p. 051901.

[LN19] Marius Lemm and Bruno Nachtergaele. "Gapped PVBS models for all species numbers and dimensions". In: *Reviews in Mathematical Physics* 31.09 (2019), p. 1950028.

[LSW20] Marius Lemm, Anders W. Sandvik, and Ling Wang. "Existence of a Spectral Gap in the Affleck-Kennedy-Lieb-Tasaki Model on the Hexagonal Lattice". In: *Physical Review Letters* 124.17 (2020).

[LV93] Ming Li and Paul Vitányi. "An Introduction to Kolmogorov Complexity and Its Applications". In: (1993).

[LSM61] E.H. Lieb, T.H. Schultz, and D.C. Mattis. "Two soluble models of an antiferromagnetic chain". In: *Annals of Physics* 15.3 (1961), 472–473.

[Liu06] Yi-Kai Liu. "Consistency of Local Density Matrices Is QMA-Complete". In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*. Ed. by Josep Díaz, Klaus Jansen, José D. P. Rolim, and Uri Zwick. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 438–449.

[Liu07] Yi-Kai Liu. "The Local Consistency Problem for Stoquastic and 1-D Quantum Systems". In: *arXiv e-prints*, arXiv:0712.1388 (Dec. 2007), arXiv:0712.1388. arXiv: 0712.1388 [quant-ph].

[LCV07] Yi-Kai Liu, Matthias Christandl, and F. Verstraete. "Quantum Computational Complexity of the N-Representability Problem: QMA Complete". In: *Phys. Rev. Lett.* 98 (11 2007), p. 110503.

[Llo96] S. Lloyd. "Universal Quantum Simulators". In: *Science* 273.5278 (1996), 1073–1078.

[LS17] Mikhail Lukin and Sylvian Schwartz. "Exploring Many-Body Dynamics on a Programmable 50 Atom Quantum Simulator". In: *Frontiers in Optics 2017* (2017).

[MH21] Ryan L. Mann and Tyler Helmuth. "Efficient algorithms for approximating quantum partition functions". In: *Journal of Mathematical Physics* 62.2 (2021), p. 022201.

[MW05] C. Marriott and J. Watrous. "Quantum Arthur-Merlin games". In: *Computational Complexity* 14.2 (2005), pp. 122–152.

[MDS96] Miguel A. Martín-Delgado and Germán Sierra. "Real Space Renormalization Group Methods and Quantum Groups". In: *Phys. Rev. Lett.* 76 (7 1996), pp. 1146–1149.

[MBK82] Susan R. McKay, A. Nihat Berker, and Scott Kirkpatrick. "Spin-Glass Behavior in Frustrated Ising Models with Chaotic Renormalization-Group Trajectories". In: *Phys. Rev. Lett.* 48 (11 1982), pp. 767–770.

[MZ13] Spyridon Michalakis and Justyna P. Zwolak. “Stability of frustration-free Hamiltonians”. In: *Communications in Mathematical Physics* 322.2 (2013), pp. 277–302.

[Mie97] Jacek Miekisz. “Stable Quasicrystalline Ground States”. In: *Journal of Statistical Physics* 88.3/4 (1997), 691–711.

[MLM07] Ari Mizel, Daniel A. Lidar, and Morgan Mitchell. “Simple Proof of Equivalence between Adiabatic Quantum Computation and the Circuit Model”. In: *Phys. Rev. Lett.* 99 (7 2007), p. 070502.

[Moo91] Cristopher Moore. “Generalized shifts: unpredictability and undecidability in dynamical systems”. In: *Nonlinearity* 4.2 (1991), pp. 199–230.

[Moo90] Cristopher Moore. “Unpredictability and undecidability in dynamical systems”. In: *Phys. Rev. Lett.* 64 (20 1990), pp. 2354–2357.

[MN03] Alexei Morozov and Antti J. Niemi. “Can renormalization group flow end in a Big Mess?” In: *Nuclear Physics B* 666.3 (2003), 311–336.

[MBK15] Richard P Muller and Robin Blume-Kohout. “The Promise of Quantum Simulation”. In: *ACS nano* 9.8 (2015), 7738—7741.

[Nac96] Bruno Nachtergaele. “The spectral gap for some spin chains with discrete symmetry breaking”. In: *Communications in Mathematical Physics* 175.3 (1996), pp. 565–606.

[NYS18] Bruno Nachtergaele, Amanda Young, and Robert Sims. “Mathematical problems in Quantum physics”. In: *Contemporary Mathematics* (2018).

[Nag+21] Daniel Nagaj, Dominik Hangleiter, Jens Eisert, and Martin Schwarz. “Pinned quantum Merlin-Arthur: The power of fixing a few qubits in proofs”. In: *Phys. Rev. A* 103 (1 2021), p. 012604.

[NC10] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge: Cambridge University Press, 2010, p. 676.

[O'G+21] Bryan O'Gorman, Sandy Irani, James Whitfield, and Bill Fefferman. "Electronic Structure in a Fixed Basis is QMA-complete". In: *arXiv e-prints*, arXiv:2103.08215 (Mar. 2021), arXiv:2103.08215. arXiv: 2103.08215 [quant-ph].

[OT08] Roberto Oliveira and Barbara M. Terhal. "The Complexity of Quantum Spin Systems on a Two-Dimensional Square Lattice". In: *Quantum Info. Comput.* 8.10 (Nov. 2008), 900–924.

[OA01] D. Osadchy and J. E. Avron. "Hofstadter butterfly as quantum phase diagram". In: *Journal of Mathematical Physics* 42.12 (2001), pp. 5665–5671. arXiv: math-ph/0101019 [math-ph].

[PT81] John Palmer and Craig Tracy. "Two-dimensional Ising correlations: Convergence of the scaling limit". In: *Advances in Applied Mathematics* 2.3 (1981), pp. 329–388.

[Pap94] Christos H. Papadimitriou. *Computational complexity*. Addison-Wesley, 1994.

[Par85] Robert G. Parr. "Density Functional Theory in Chemistry". In: *Density Functional Methods In Physics* (1985), 141–158.

[Pat14] Matthew J. Patitz. "An introduction to tile-based self-assembly and a survey of recent results". In: *Natural Computing* 13.2 (2014), pp. 195–224.

[PKC15] Eva Pavarini, Erik Koch, and Piers Coleman. *Many-Body Physics: From Kondo to Hubbard*. 1st ed. Verlag des Forschungszentrum Jülich, 2015. Chap. 12.

[PJP82] K. A. Penson, R. Jullien, and P. Pfeuty. "Zero-temperature renormalization-group method for quantum systems. V. Frustration in two dimensions". In: *Phys. Rev. B* 25 (3 1982), pp. 1837–1847.

[Per+92] John P. Perdew et al. “Atoms, molecules, solids, and surfaces: Applications of the generalized gradient approximation for exchange and correlation”. In: *Physical Review B* 46.11 (1992), 6671–6687.

[PB20] Stephen Piddock and Johannes Bausch. “Universal Translationally-Invariant Hamiltonians”. In: (2020). arXiv: 2001.08050.

[PM18] Stephen Piddock and Ashley Montanaro. “Universal qudit Hamiltonians”. In: *arXiv e-prints*, arXiv:1802.07130 (Feb. 2018), arXiv:1802.07130. arXiv: 1802.07130 [quant-ph].

[PW09] David Poulin and Pawel Wocjan. “Sampling from the Thermal Quantum Gibbs State and Evaluating Partition Functions with a Quantum Computer”. In: *Phys. Rev. Lett.* 103 (22 2009), p. 220502.

[RT19] Ran Raz and Avishay Tal. “Oracle Separation of BQP and PH”. In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2019. Phoenix, AZ, USA: Association for Computing Machinery, 2019, 13–23.

[Ric53] H. G. Rice. “Classes of recursively enumerable sets and their decision problems”. In: *Transactions of the American Mathematical Society* 74.2 (1953), 358–358.

[Rob71] Raphael M. Robinson. “Undecidability and nonperiodicity for tilings of the plane”. In: *Inventiones mathematicae* 12.3 (Sept. 1971), pp. 177–209.

[SKO19] M. W. Wilde S. K. Oskouei S. Mancini. “Union bound for quantum information processing”. In: *Royal Society A* 475.2221 (2019).

[Sac11] Subir Sachdev. *Quantum Phase Transitions*. 2nd ed. Cambridge University Press, 2011.

[SV09] N. Schuch and F. Verstraete. “Computational complexity of interacting electrons and fundamental limitations of density functional theory”. In: *Nature Physics* 5.10 (2009), pp. 732–735. arXiv: 0712.0483 [quant-ph].

[SS14] Stephen H. Shenker and Douglas Stanford. “Black holes and the butterfly effect”. In: *Journal of High Energy Physics* 2014.3 (2014).

[SMM09] L. Sheridan, D. Maslov, and M. Mosca. “Approximating fractional time quantum evolution”. In: *Journal of Physics A: Mathematical and Theoretical* 42.18 (2009), p. 185302. arXiv: 0810.3843.

[SM21] Naoto Shiraishi and Keiji Matsumoto. “Undecidability in quantum thermalization”. In: *Nature Communications* 12.1 (2021).

[Sho97] Peter W. Shor. “Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer”. In: *SIAM Journal on Computing* 26.5 (1997), 1484–1509.

[SS03] Wolfgang L. Spitzer and Shannon Starr. “Improved Bounds on the Spectral Gap Above Frustration-Free Ground States of Quantum Spin Chains”. In: *Letters in Mathematical Physics* 63.2 (2003), 165–177.

[SMS13] Peter Staar, Thomas Maier, and Thomas C. Schulthess. “Dynamical cluster approximation with continuous lattice self-energy”. In: *Physical Review B* 88.11, 115101 (2013), p. 115101. arXiv: 1304.3624 [cond-mat.str-el].

[Sto83] Larry Stockmeyer. “The complexity of approximate counting”. In: *Proceedings of the fifteenth annual ACM symposium on Theory of computing - STOC ’83* (1983).

[SKS82] N M Svrakic, J Kertesz, and W Selke. “Hierarchical lattice with competing interactions: an example of a nonlinear map”. In: *Journal of Physics A: Mathematical and General* 15.8 (1982), pp. L427–L432.

[Tac22] Yuji Tachikawa. “Undecidable problems in quantum field theory”. In: *arXiv e-prints*, arXiv:2203.16689 (Mar. 2022), arXiv:2203.16689. arXiv: 2203.16689 [hep-th].

[Tot+95] K Totsuka, Y Nishiyama, N Hatano, and M Suzuki. “Isotropic spin-1 chains with bond alternation: analytic and numerical studies”. In: *Journal of Physics: Condensed Matter* 7.25 (1995), 4895–4920.

[Tur37] A. M. Turing. “On Computable Numbers, with an Application to the Entscheidungsproblem”. In: *Proceedings of the London Mathematical Society* s2-42.1 (1937), 230–265.

[VC05] F Verstraete and J I Cirac. “Mapping local Hamiltonians of fermions to local Hamiltonians of spins”. In: *Journal of Statistical Mechanics: Theory and Experiment* 2005.09 (2005).

[Voj00] Thomas Vojta. “Quantum phase transitions in electronic systems”. In: *Annalen der Physik* 9.6 (2000), pp. 403–440. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/1521-3889%28200006%299%3A6%3C403%3A%3AAID-ANDP403%3E3.0.CO%3B2-R.

[WKL02] J. Wang, Sabre Kais, and R. Levine. “Real-space renormalization group study of the Hubbard model on a non-bipartite lattice”. In: *International Journal of Molecular Sciences* 3.1 (2002), 4–16.

[Wat19] James D. Watson. “Detailed Analysis of Circuit-to-Hamiltonian Mappings”. In: *arXiv e-prints*, arXiv:1910.01481 (2019), arXiv:1910.01481. arXiv: 1910.01481 [quant-ph].

[WB21] James D. Watson and Johannes Bausch. “The Complexity of Approximating Critical Points of Quantum Phase Transitions”. In: *arXiv e-prints*, arXiv:2105.13350 (May 2021), arXiv:2105.13350. arXiv: 2105.13350 [quant-ph].

[WBG20] James D. Watson, Johannes Bausch, and Sevag Gharibian. “The Complexity of Translationally Invariant Problems beyond Ground State Energies”. In: *arXiv e-prints*, arXiv:2012.12717 (Dec. 2020), arXiv:2012.12717. arXiv: 2012.12717 [quant-ph].

[WC21] James D. Watson and Toby S. Cubitt. “Computational Complexity of the Ground State Energy Density Problem”. In: *arXiv e-prints*, arXiv:2107.05060 (July 2021), arXiv:2107.05060. arXiv: 2107.05060 [quant-ph].

[WC22] James D. Watson and Toby S. Cubitt. "Computational Complexity of the Ground State Energy Density Problem". In: *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2022. Rome, Italy: Association for Computing Machinery, 2022, 764–775.

[WOC21] James D. Watson, Emilio Onorati, and Toby S. Cubitt. "Uncomputably Complex Renormalisation Group Flows". In: *arXiv e-prints*, arXiv:2102.05145 (Feb. 2021), arXiv:2102.05145. arXiv: 2102.05145 [quant-ph].

[Whi92] Steven R. White. "Density matrix formulation for quantum renormalization groups". In: *Physical Review Letters* 69.19 (1992), 2863–2866.

[Wil71] Kenneth G. Wilson. "Renormalization Group and Critical Phenomena. I. Renormalization Group and the Kadanoff Scaling Picture". In: *Phys. Rev. B* 4 (9 1971), pp. 3174–3183.

[WK74] Kenneth G. Wilson and J. Kogut. "The renormalization group and the ε expansion". In: *Physics Reports* 12.2 (1974), pp. 75 –199.

[Yam97] Shoji Yamamoto. "Quantum fluctuations in the ground state of the S = 1 antiferromagnetic Heisenberg chain". In: *Physics Letters A* 225.1-3 (1997), 157–166.

[Zur90] Wojciech Hubert. Zurek. *Complexity, entropy, and the physics of information*. Addison-Wesley, 1990.