

# The power of quantum systems on a line

Dorit Aharonov\*  
School of Computer Science and Engineering,  
Hebrew University, Jerusalem, Israel

Daniel Gottesman†  
Perimeter Institute  
Waterloo, Canada

Sandy Irani‡  
Computer Science Department  
University of California, Irvine, USA

Julia Kempe§  
School of Computer Science,  
Tel Aviv University, Israel

February 19, 2009

## Abstract

We study the computational strength of quantum particles (each of finite dimensionality) arranged on a line. First, we prove that it is possible to perform universal adiabatic quantum computation using a one-dimensional quantum system (with 9 states per particle). This might have practical implications for experimentalists interested in constructing an adiabatic quantum computer. Building on the same construction, but with some additional technical effort and 12 states per particle, we show that the problem of approximating the ground state energy of a system composed of a line of quantum particles is QMA-complete; QMA is a quantum analogue of NP. This is in striking contrast to the fact that the analogous classical problem, namely, one-dimensional MAX-2-SAT with nearest neighbor constraints, is in P. The proof of the QMA-completeness result requires an additional idea beyond the usual techniques in the area: Not all illegal configurations can be ruled out by local checks, so instead we rule out such illegal configurations because they would, in the future, evolve into a state which can be seen locally to be illegal. Our construction implies (assuming the quantum Church-Turing thesis and that quantum computers cannot efficiently solve QMA-complete problems) that there are one-dimensional systems which take an exponential time to relax to their ground states at any temperature, making them candidates for being one-dimensional spin glasses.

## 1 Introduction

The behavior of classical or quantum spin systems frequently depends very heavily on the number of spatial dimensions available. In particular, there is often a striking difference between the behavior of one-dimensional systems and of otherwise similar two- or higher-dimensional systems. For instance, one-dimensional systems generally do not experience phase transitions except at zero temperature, whereas

---

\*E-mail: doria@cs.huji.ac.il. Supported by Israel Science Foundation grant number 039-7549, Binational Science Foundation grant number 037-8404, and US Army Research Office grant number 030-7790.

†E-mail: dgottesman@perimeterinstitute.ca. Supported by CIFAR, by the Government of Canada through NSERC, and by the Province of Ontario through MRI.

‡E-mail: irani@ics.uci.edu. Partially supported by NSF Grant CCR-0514082.

§E-mail: kempe@cs.tau.ac.il. This work was mainly done while the author was at CNRS and LRI, University of Paris-Sud, Orsay, France. Partially supported by the European Commission under the Integrated Project Qubit Applications (QAP) funded by the IST directorate as Contract Number 015848, by an ANR AlgoQP grant of the French Research Ministry, by an Alon Fellowship of the Israeli Higher Council of Academic Research by an Individual Research grant of the ISF, and by a European Research Council (ERC) Starting Grant.

phase transitions are common in other dimensions. As another example, satisfiability with nearest neighbor constraints between constant-size variables set on a two- or larger-dimensional grid is a hard problem (NP-complete, in fact), whereas for a one-dimensional line, it can be solved in polynomial time.

We thus ask: what is the computational power associated with a line of quantum particles? There are a number of ways to interpret the question. For instance, we can ask whether the evolution of such systems can be efficiently simulated by a classical computer. We can ask whether the system is powerful enough to create a universal quantum computer under various scenarios for how we may control the system — in which case we cannot hope to simulate this behavior efficiently on a classical computer, unless of course  $BQP = BPP$ . (BPP and BQP are the classes of problems efficiently solvable on a probabilistic classical computer and on a quantum computer, respectively.) We can also ask how difficult it is to calculate or approximate interesting properties of the system, such as its ground state energy (that is, the lowest eigenvalue of the Hamiltonian operator for the system).

For many one-dimensional quantum systems, it *is* indeed possible to classically compute many properties of the system, including in some cases the dynamical behavior of the system. The method of the density matrix renormalization group (DMRG) [Whi92, Whi93, Sch05] has been particularly successful here, although there is no complete characterization of which one-dimensional systems it will work on. Indeed, DMRG provides a good example of the difference between one- and two-dimensional systems; there are only a few results applying DMRG techniques to simulate special two-dimensional systems.

However, it has long been known that one-dimensional quantum systems can also, under the right circumstances, perform universal quantum computation. It is straightforward to create a quantum computer from a line of qubits with fully controllable nearest-neighbor couplings by using SWAP gates to move the qubits around. Even a one-dimensional quantum cellular automaton can perform universal quantum computation [Wat95, SFW06]; the smallest known construction has 12 states per cell. While many one-dimensional systems are relatively simple and can be simulated classically, the general one-dimensional quantum system must thus have complexities that are inaccessible classically.

In two interesting (and closely related) subfields of the area of quantum computation, studied extensively over the past few years, it has been conjectured that one-dimensional systems are not as powerful as systems of higher dimensionality. These cases are adiabatic evolution, and the QMA-completeness of the local Hamiltonian problem.

## 1.1 Results related to adiabatic computation

In an adiabatic quantum computer, the Hamiltonian of the system is slowly shifted from a simple Hamiltonian, whose ground state is easy to calculate and to create, to a more complicated Hamiltonian, whose ground state encodes the solution to some computational problem. The quantum adiabatic theorem guarantees that if the Hamiltonian is changed slowly enough, the system stays close to its ground state; the time required to safely move from the initial Hamiltonian to the final Hamiltonian is polynomial in the minimal spectral gap between the ground state and the first excited state over the course of the computation.

The first adiabatic quantum algorithm was introduced by Farhi *et al.* [FGGS00, FGG<sup>+</sup>01], though the idea of encoding the solution to an optimization problem in the ground state of a Hamiltonian appeared as early as 1988 [ACdF90, ACdF88]. The computational power of the model was clarified in [AvDK<sup>+</sup>04], where it was shown that adiabatic quantum computers could in fact perform any quantum computation. Adiabatic computation can thus be viewed as an alternative model to quantum circuits or quantum Turing machines.

Adiabatic quantum computers might be more robust against certain kinds of noise [CFP02], and the slow change of parameters used in an adiabatic quantum computer might be more amenable to some kinds of experimental implementation. To this end, it is important to understand what sorts of physical systems can be used to build an adiabatic quantum computer, and in particular, we would like the interactions of

the physical Hamiltonian to be as simple as possible. Several research groups have devoted effort to this question recently. To describe their work, let us make the following definition:

**Definition 1.1 (Local Hamiltonian)** *Let  $H$  be a Hermitian operator (interpreted as a Hamiltonian, giving the energy of some system). We say that  $H$  is an  $r$ -state Hamiltonian if it acts on  $r$ -state particles. When  $r = 2$ , namely, when the particles are qubits, we often omit mention of the number of states per particle. We say that  $H$  is  $k$ -local if it can be written as  $H = \sum_i H_i$ , where each  $H_i$  acts non-trivially on at most  $k$  particles. Note that this term does not assume anything about the physical location of the particles. We say that  $H$  is a  $d$ -dim Hamiltonian if the particles are arranged on a  $d$ -dimensional grid and the terms  $H_i$  interact only pairs of nearest neighbor particles. Note that the term  $d$ -dim implies that the Hamiltonian is 2-local.*

We will in this paper assume for simplicity that each local term  $H_i$  in the Hamiltonian is positive definite and of norm at most 1. It is always possible to ensure this by shifting and rescaling the Hamiltonian  $H \mapsto \lambda(H + E_0)$ . Provided we also similarly shift and rescale all other energies in the problem (such as the spectral gap), this transformation will not at all alter the properties of the Hamiltonian. We will further assume, to avoid pathologies, that there are at most polynomially many terms  $H_i$  and that the coefficients in each  $H_i$  are specified to polynomially many bits of precision.

[AvDK<sup>+</sup>04] made a first step towards a practical Hamiltonian, by showing that a 2-dim 6-state Hamiltonian suffices for universal adiabatic quantum computation. Kempe, Kitaev and Regev [KKR06], using perturbation-theory gadgets, improved the results and showed that qubits can be used instead of six-dimensional particles, namely, adiabatic evolution with 2-local Hamiltonians is quantum universal. Their interactions, however, were not restricted to a two-dimensional grid. Terhal and Oliveira [OT05] combined the two results, using other gadgets, to show that a similar result holds with 2-dim 2-state Hamiltonians.

Naturally, the next question was whether a one-dimensional adiabatic quantum computer could be universal. Since one-dimensional systems are fairly simple, it was conjectured that the answer was “no.” However, in this paper, we prove

**Theorem 1.2** *Adiabatic computation with 1-dim 9-state Hamiltonians is universal for quantum computation.*

As mentioned, this could be important experimentally, as one-dimensional systems are easier to build than two-dimensional ones, and adiabatic systems might be more robust than standard ones. Any simplification to the daunting task of building a quantum computer is potentially useful. However, a more systematic way of dealing with errors will be needed before it is possible to build a large adiabatic quantum computer. The results of [CFP02] only imply adiabatic quantum computation is robust against certain types of control errors, and it remains an interesting open question to show that adiabatic quantum computation can be made fault-tolerant in the face of general small errors. An important step in this direction is [JFS06], which introduces some quantum error-correcting codes in the adiabatic model.

It is also worth noting that the adiabatic construction given in the proof of Theorem 1.2 actually has a degenerate ground state, which makes it less robust against noise than a non-degenerate adiabatic computer, although still resistant to timing errors, for instance. Using a 1-dim 13-state Hamiltonian similar to the QMA-complete 1-dim 12-state Hamiltonian family described below, it is possible to break the degeneracy. This construction is outlined in Section 4.1.

It is natural to ask whether our result holds even when we restrict ourselves to translationally invariant Hamiltonians. We sketch one way to do this in Section 2. However, it does not appear to be possible to make the translational invariant adiabatic construction nondegenerate.

## 1.2 Results related to QMA-completeness

Theorem 1.2 means that efficient simulation of general one-dimensional adiabatic quantum systems is probably impossible. One might expect that calculating, or at least approximating, some specific property of a system, such as its ground state energy, would be more straightforward, as this does not require complete knowledge of the system. This intuition is misleading, and in fact it can be *harder* to find the ground state energy than to simulate the dynamics of a system. More concretely, it has long been known that it is NP-hard to find the ground state energy of some classical spin systems, such as a three-dimensional Ising model. Kitaev [KSV02] extended these results to the quantum realm, by defining the quantum analogue of NP, called QMA (for “quantum Merlin-Arthur”). QMA is thus, roughly speaking, the class of problems that can be efficiently checked on a quantum computer, provided we are given a “witness” quantum state related to the answer to the problem. QMA is believed to be strictly larger than both BQP and NP. That is, a problem which is QMA-complete can most likely not be solved efficiently with a quantum computer, and the answer cannot even be checked efficiently on a classical computer.

Kitaev proved that the problem of approximating the ground state energy of a quantum system with a 5-local Hamiltonian is complete for this class. (The ability to solve a QMA-complete problem implies the ability to solve any problem in QMA.) The exact definition of the local Hamiltonian problem is this:

**Definition 1.3 (Local Hamiltonian problem)** *Let  $H = \sum_i H_i$  be an  $r$ -state  $k$ -local Hamiltonian. Then  $(H, E, \Delta)$  is in  $r$ -STATE  $k$ -LOCAL HAMILTONIAN if the lowest eigenvalue  $E_0$  of  $H$  is less than or equal to  $E$ . The system must satisfy the promises that  $\Delta = \Omega(1/\text{poly}(n))$  and that either  $E_0 \leq E$  or  $E_0 \geq E + \Delta$ . We can define a language  $d$ -DIM  $r$ -STATE HAMILTONIAN similarly.*

Kitaev’s proof uses what is called the circuit-to-Hamiltonian construction of Feynman [Fey85], which we sketch in Subsection 1.5. The main result of [AvDK<sup>+</sup>04], namely the universality of adiabatic computation, was based on the observation that the circuit-to-Hamiltonian construction is useful also in the adiabatic context, where it is used to design the final Hamiltonian of the adiabatic evolution. It turns out that the same connection between adiabatic universality proofs and QMA-completeness proofs can also be made for many of the follow-up papers on the two topics. It is often the case that whenever adiabatic universality can be proven for some class of Hamiltonians, then the local Hamiltonian problem with the same class (roughly) can be shown to be QMA-complete — and vice versa.

Note, however, that there is no formal implication from either of those problems to the other. On one hand, proving QMA-completeness is in general substantially harder than achieving universal adiabatic quantum computation: In an adiabatic quantum computer, we can choose the initial state of the adiabatic computation to be any easily-created state which will help us solve the problem, so we can choose to work on any convenient subspace which is invariant under the Hamiltonian. For QMA, the states we work with are chosen adversarially from the full Hilbert space, and we must be able to check ourselves, using only local Hamiltonian terms, that they are of the correct form. On the other hand, proving adiabatic universality involves analyzing the spectral gap of all of the Hamiltonians  $H(t)$  over the duration of the computation, whereas QMA-completeness proofs are only concerned with one Hamiltonian.

Both [KKR06] and [OT05] in fact prove QMA-completeness; universal adiabatic quantum computation follows with a little extra work. Thus, the 2-DIM HAMILTONIAN problem is QMA-complete.

However, the 1-DIM  $r$ -STATE HAMILTONIAN problem remained open for all (constant)  $r$ . It was suspected that the problem was not QMA-complete, and in fact might be in BQP or even BPP. For one thing, ground state properties of one-dimensional quantum systems are generally considered particularly easy. For instance, Osborne has recently proven [Os07] that there are efficient classical descriptions for a class of one-dimensional quantum systems. DMRG techniques have been employed extensively to calculate ground state energies and other properties of a variety of one-dimensional quantum systems. Furthermore, the classical analogue of 1-DIM  $r$ -STATE HAMILTONIAN is easy: Take  $r$ -state variables arranged on a

line with constraints restricting the values of neighboring pairs of variables. If we assign a constant energy penalty for violating a constraint, the lowest energy state satisfies as many constraints as possible. This problem, a one-dimensional restriction of MAX-2-SAT with  $r$ -state variables, is in fact in P; it can be solved with a recursive divide-and-conquer algorithm or by dynamic programming.

For instance, we can divide the line in half, and run through all possible assignments of the two variables  $x_i$  and  $x_{i+1}$  adjacent to the division. For each pair of values for  $x_i$  and  $x_{i+1}$ , we can calculate the maximal number of satisfiable constraints by solving the sub-problems for the right and left halves of the line. Then we can compare across all  $r^2$  assignments for the one that gives the largest number of satisfied constraints. Thus, to solve the problem for  $n$  variables, it is sufficient to solve  $2r^2$  similar problems for  $n/2$  variables. By repeatedly dividing in half, we can thus solve the problem in  $O((2r^2)^{\log n}) = O(n^{\log(2r^2)})$  steps. MAX- $k$ -SAT in one dimension is also in P, and can in fact be reduced to MAX-2-SAT with large enough variables.

Despite the intuition that one-dimensional systems should not be too difficult, we prove:

**Theorem 1.4** *1-DIM 12-STATE HAMILTONIAN is QMA-complete.*

The theorem implies a striking qualitative difference between the quantum and the classical one-dimensional versions of the same problem — one is easy, whereas the other is complete for a class which seems to be strictly larger than NP. This might seem surprising, but in retrospect, we can provide an intuitive explanation. The reason is that the  $k$ -local Hamiltonian essentially allows us to encode an extra dimension, namely, time, by making the ground state a superposition of states corresponding to different times. In other words, the result implies that the correct analogue of one-dimensional local Hamiltonian is *two*-dimensional MAX- $k$ -SAT, which is of course NP-complete. Indeed, there are many cases in physics where one-dimensional quantum systems turn out to be most closely analogous to two-dimensional classical systems. One such example is given by Suzuki who showed that the  $d$ -dimensional Ising model with a transverse field is equivalent to the  $(d + 1)$ -dimensional Ising model at finite temperatures [Suz76], although the extra dimension there is quite different from the dimension of time addressed in this paper.

### 1.3 Implications of our results

One consequence of Theorem 1.4 is that there exist one-dimensional systems which take exponentially long to relax to the ground state, with any reasonable environment or cooling strategy. To see this, we invoke a quantum version of the modern Church-Turing thesis,<sup>1</sup> which would state that any reasonable physical system can be efficiently simulated on a quantum computer. As it invokes a notion of a “reasonable” physical system, this is of course not a provable statement, but one can be convinced of it by looking at examples. The details of doing so for any particular system may be complicated, but no counter-example is known, and barring undiscovered laws of physics modifying quantum mechanics, the quantum Church-Turing thesis is believed to be true.

In this particular case, we could thus use a quantum computer to simulate the system plus the environment. If the system reliably relaxes to the ground state in a time  $T$ , then by causality, we need only simulate an environment of size  $O(T^3)$  (for a three-dimensional environment); the quantum simulation can be performed using standard techniques and runs in polynomial time. We can then use the simulation to solve 1-DIM 12-STATE HAMILTONIAN problems. Since the latter problem is difficult, that implies the simulation must also take a long time to reach the ground state. This observation assumes, of course, that QMA is exponentially hard to solve by a quantum computer.

Similarly, we can argue that the class of systems presented in this paper will take an exponentially long time to relax to the thermal equilibrium state at very low temperatures (again, assuming QMA is hard for a quantum computer). To see this, suppose that there is a state of the system with energy less than  $E$  (the

---

<sup>1</sup>The modern Church-Turing thesis is discussed, for instance, in [vEB90].

bound on the ground state energy from Definition 1.3), and note that at a temperature less than  $O(\Delta/n)$  (with  $\Delta$  polynomially small in the system size  $n$ ), the Gibbs distribution gives a constant probability of the system occupying the low-lying state, even if there are exponentially many eigenstates with energy just above  $E + \Delta$ . Of course, if there is no state with energy less than  $E + \Delta$ , then the system will never have energy less than  $E + \Delta$ , no matter what the temperature. Therefore, if the system were able to reach its thermal equilibrium state, by observing it, we would be sampling from the Gibbs distribution, and would, with good probability, be able to answer the QMA-complete 1-DIM 12-STATE HAMILTONIAN problem.

Exponentially long relaxation times are a characteristic of a *spin glass* [BY86], making this system a candidate to be a one-dimensional spin glass. Spin glasses are not completely understood, nor is there a precise definition of the term, but it has been unclear whether a true one-dimensional spin glass can exist. There are a number of properties relevant to being a spin glass, long relaxation times among them, and classical one-dimensional systems generally lack that property and most other typical spin glass properties. On the other hand, some other one-dimensional quantum systems are known [Fis95] that exhibit long relaxation times and some, but not all, other spin-glass-like properties. Our result may help shed further light on the situation. Unfortunately, we are unable to present a specific Hamiltonian with long relaxation times. As usual in complexity, it is difficult to identify specific hard instances of computationally hard problems. Similarly, we cannot say very much about other physically interesting properties of the system, such as the nature or existence of a liquid/glass phase transition.

This paper is a merger of [Ira07] and version 1 of [AGK07], two similar papers originally written independently. The extended abstract of the merged version appeared in [AGIK07]. Following the initial version of our paper, various groups have improved and extended our result. For instance, [NW08] gives a 20-state translation-invariant modification of our construction (improving on a 56-state construction by [JWZ07]) that can be used for universal 1-dimensional adiabatic computation. Kay [Kay08] gives a QMA-complete 1-dim  $r$ -state Hamiltonian with all two-particle terms identical. Nagaj [Nag08] has improved our QMA construction to use 11-state particles.

## 1.4 Notation

A *language* is computer science terminology for a computational problem; it can be defined as a set of bit strings. Frequently we consider the language to be a proper subset of all possible bit strings, but here we will be more concerned with languages which satisfy a *promise*, meaning we consider the language as a subset of a smaller set of bit strings which satisfy some constraint. We will refer to the set of strings which satisfy the constraint as the *instances* of the problem. The instances then can be partitioned into two sets  $L_{\text{yes}}$ , the set of instances which are in the language, and  $L_{\text{no}}$ , the set of instances which are not in the language. The *decision problem* associated with the language is to decide if a given bit string is in the language or not. The instances represent different possible inputs about which we are attempting to reach a yes/no decision. For example, satisfiability (or SAT) is the decision problem for the language composed of Boolean formulas which have satisfying assignments of the variables; each instance encodes a possible formula, some of which can be satisfied and some of which cannot. A *complexity class* is a set of languages.

The precise definition of the complexity class QMA is as follows:

**Definition 1.5** *A language  $L$  is in QMA iff for each instance  $x$  there exists a uniform polynomial-size quantum circuit<sup>2</sup>  $C_x$  such that (a) if  $x \in L_{\text{yes}}$ ,  $\exists|\psi\rangle$  (the “witness”, a polynomial-size quantum state) such that  $C_x|\psi\rangle$  accepts with probability at least  $2/3$ , and (b) if  $x \in L_{\text{no}}$ , then  $\forall|\psi\rangle$ ,  $C_x|\psi\rangle$  accepts with probability at most  $1/3$ . We only consider strings which are instances of  $L$ .*

---

<sup>2</sup>A *uniform circuit* is one whose description can be generated in polynomial time by a Turing machine.

We say we *reduce* language  $L$  to language  $M$  if we have a function  $f$  converting bit strings  $x$  to bit strings  $f(x)$  such that if  $x \in L_{\text{yes}}$  then  $f(x) \in M_{\text{yes}}$  and if  $x \in L_{\text{no}}$ , then  $f(x) \in M_{\text{no}}$ . We do not put any requirement on the behavior of  $f$  when the input  $x$  does not satisfy the promise.  $f$  must be computable classically in polynomial time. The point is that if we can solve the decision problem for  $M$  then we can automatically also solve the decision problem for  $L$  by first converting  $x$  to  $f(x)$  and then checking if  $f(x)$  is in  $M$ . There are other notions of reduction, but this will be sufficient for our purposes. A language  $L$  is *complete* for a complexity class  $C$  if any language in  $C$  can be reduced to  $L$ . Thus, solving a  $C$ -complete language  $L$  implies the ability to solve any problem in  $C$  with polynomial time overhead;  $L$  is thus one of the most computationally difficult problems in  $C$ .

We will be interested in reducing an arbitrary QMA language  $L$  to variants of the language  $r$ -STATE  $k$ -LOCAL HAMILTONIAN. We will use  $x$  to refer to an instance of the original problem  $L$ , and  $C_x$  to refer to the checking circuit associated with the instance  $x$ . That is, we assume that information about the instance is encoded in the structure of the checking circuit. The circuit  $C_x$  acts on  $n$  qubits. Some of the qubits will be ancilla qubits which must be initialized to the state  $|0\rangle$ , whereas others are used for the potential witness for  $x \in L$ .

In the context of adiabatic quantum computation, we will also use  $C_x$  to refer to a more general quantum circuit, unrelated to a QMA problem. In this case, we consider  $C_x$  to be in a form where all input qubits are  $|0\rangle$ . When the adiabatic quantum computation is intended to compute some classical function (such as factoring),  $x$  is the classical input to the function, which we hardwire into the structure of the circuit itself.

## 1.5 Outline of the approach

To explain the main idea behind the proofs of our two main theorems, we recall Kitaev's proof that 5-LOCAL HAMILTONIAN is QMA-complete.

The fact that  $r$ -STATE  $k$ -LOCAL HAMILTONIAN is in QMA is not difficult. The witness is the ground state of the Hamiltonian (or indeed any state with energy less than  $E$ ), and a standard phase estimation technique can be used to measure its energy. The accuracy of the measurement is not perfect, of course, which is why we need the promise that the energy is either below  $E$  or above  $E + \Delta$ .

To prove completeness, we need to reduce an arbitrary QMA language  $L$  to 5-LOCAL HAMILTONIAN. The translation to local Hamiltonians is done by creating a Hamiltonian whose ground state is of the form  $\sum_t |\phi_t\rangle|t\rangle$  (ignoring normalization), where the first register  $|\phi_t\rangle$  is the state of the checking circuit  $C_x$  at time  $t$  and the second register acts as a clock. We call this state the *history state* of the circuit.<sup>3</sup> The main term in Kitaev's Hamiltonian sets up constraints ensuring that  $|\phi_{t+1}\rangle = U_t|\phi_t\rangle$ , where  $U_t$  is the  $(t + 1)$ -st gate in  $C_x$ . This term, denoted  $H_{\text{prop}}$ , ensures that the ground state is indeed a history state, reflecting a correct propagation in time according to  $C_x$ . The clock is used to associate the correct constraint with each branch of the superposition; any state which does not have the correct time evolution for the circuit will violate one or more constraints and will thus have a higher energy.

Kitaev's Hamiltonian includes more terms, which guarantee that the input to  $C_x$  is correct, that the final state of the circuit is accepted, and that the state of the clock is a valid clock state and not some arbitrary state. In the context of adiabatic evolution, these additional terms are not needed, since we have control over the initial state, but they are needed to prove QMA-completeness, since we must be able to check (i) that the witness being tested for the QMA-complete problem has the correct structure, corresponding to a valid history state for a possible input to  $C_x$ , and (ii) that  $C_x$  accepts on that input.

In order to prove the universality of adiabatic quantum computation, [AvDK<sup>+</sup>04] let Kitaev's Hamiltonian  $H_{\text{prop}}$  be the final Hamiltonian in the adiabatic evolution, and set the initial Hamiltonian to be diagonal in the standard basis, forcing the initial ground state to be the correct input state  $|\phi_0\rangle|0\rangle$  of the circuit. At

---

<sup>3</sup>We call a state a history state if it has the above form for *any* input  $|\phi_0\rangle$  to  $C_x$ , not just the correct input.

any point during the adiabatic quantum computation, the state of the system is then in a subspace  $\mathcal{I}$  spanned by the states  $|\phi_t\rangle|t\rangle$ , and the spectral gap of any convex combination of the initial and final Hamiltonians restricted to this subspace is at most polynomially small in the size of the original circuit. Therefore, the adiabatic quantum computation can be performed with at most polynomial overhead over the original circuit, to reach a state which is very close to the history state, from which the result of the original circuit can be measured with reasonable probability.

Our idea for proving universality of adiabatic evolution and QMA-completeness in one dimension is similar. One would consider a quantum circuit  $C_x$ , and design a 1-dim  $r$ -state Hamiltonian which will verify correct propagation according to this circuit, and then use this Hamiltonian as the final Hamiltonian for the adiabatic evolution or as the instance of the QMA-complete problem corresponding to the instance  $x$  of the original problem.

This idea, however, is not easy to realize when our Hamiltonian is restricted to work on a one- or two-dimensional lattice, since we cannot directly access a separate clock register, as only the subsystems nearest to the clock would be able to take advantage of it in order to check correct propagation in time. Instead, following the strategy of [AvDK<sup>+</sup>04], we must first modify the original circuit  $C_x$  into a new circuit  $\tilde{C}_x$ . This allows us to distribute the clock, making the time implicit in the global structure of the state. When time is encoded somehow in the configuration, then it can be ensured that the transition rules allow only propagation to the correct next step. The construction of [AvDK<sup>+</sup>04] used the following arrangement: the qubits of the original circuit were put initially in the left-most column in the two-dimensional grid, and one set of gates was performed on them. To advance time and perform another set of gates, the qubits were moved over one column, leaving used-up “dead” states behind. Doing all this required going up to 6 states per particle instead of 2, with the qubits of the original circuit encoded in various two-dimensional subspaces. The time could therefore be read off by looking at the arrangement and location of the two-dimensional subspaces containing the data.

This construction relies heavily on the ability to copy qubits to the next column in order to move to the next block of gates in the computation, so a new strategy is needed in one dimension. For the modified circuit  $\tilde{C}_x$ , we instead place the qubits in a block of  $n$  adjacent particles. We do one set of gates, and then move all of the qubits over  $n$  places to advance time in the original circuit  $C_x$ . It is considerably more complicated to move qubits  $n$  places than to move them over one column in a two-dimensional arrangement, and we thus need extra states. The adiabatic construction can be done using 9 states per particle, which we increase to 12 for QMA-completeness.

A straightforward application of existing techniques can then complete the proof of one-dimensional universal adiabatic quantum computation, but there is an additional wrinkle for proving QMA-completeness. In particular, previous results used local constraints to ensure that the state of the system had a valid structure; for instance, in [AvDK<sup>+</sup>04], terms in the Hamiltonian check that there are not two qubit states in adjacent columns. However, using only local constraints, there is no way to check that there are exactly  $n$  qubit data states in an unknown location in a one-dimensional system — there are only a constant number of local rules available, which are therefore unable to count to an arbitrarily large  $n$ . We instead resort to another approach, which might be useful elsewhere. While our modified circuit has invalid configurations (containing, for instance, too many qubit states) which cannot be locally checked, we ensure that, under the transition rules of the system, any invalid configurations will evolve in polynomial time into a configuration which can be detected as illegal by local rules. Thus, for every state which is not a valid history state, either the propagation is wrong, which implies an energy penalty due to the propagation Hamiltonian, or the state evolves to an illegal configuration which is locally detectable, which implies an energy penalty due to the local check of illegal configurations. We call this result the *clairvoyance lemma*.

The structure of the paper is as follows: in Section 2 we describe how to map circuits to one-dimensional arrangements. Section 3 shows our result on the universality of adiabatic computation on the line, and



Section 4 gives the result on QMA-completeness. We conclude in Section 5.

## 2 The basic construction

In this section, we will describe our construction which maps a quantum circuit  $C_x$  to a modified circuit  $\tilde{C}_x$ . We first present the larger 12-state QMA construction, and explain how to modify it to get the 9-state adiabatic construction at the end of this section. There are a number of properties which  $\tilde{C}_x$  must satisfy. It should perform the same computation as  $C_x$ , of course. In addition, each gate in  $\tilde{C}_x$  must interact only nearest-neighbor particles in a line; on the other hand, we allow those particles to be 12-dimensional. The gate performed at any given time cannot depend explicitly on the time (the number of gates already performed), but can depend on location. We will define the gates in terms of a variety of possible transition rules, and to remove any ambiguity, we will ensure that for a legal state of the system, only one transition rule will apply at any given time. For QMA-completeness, we need some additional properties ensuring that enough of the constraints are locally checkable.

The problem of moving to one dimension is somewhat similar to a quantum cellular automaton in that the transition rules need to depend only on the local environment and not on some external time coordinate, but differs from a cellular automaton in a number of ways. A cellular automaton acts on all locations simultaneously and in the same way, whereas our transition rules are required to only cause one pair of particles to change per time step (provided the system is in a state of the correct form), and the transition rules differ slightly from location to location.

A better analogy is to a single-tape Turing machine. We will have a single “active site” in our computation, analogous to the head of a Turing machine, which moves around manipulating the computational qubits, changing states as it does so in order to perform different kinds of actions. The transition rules of a Turing machine are independent of the location. In the construction below, we have a few different kinds of locations, for instance for the different sorts of quantum gates used in  $C_x$ . Each kind of location has a different set of transition rules that apply (although many of the rules are the same for all types of location), and the position of a site determines which set of rules applies to that site. For the adiabatic construction, we could instead proceed with translational-invariant rules as with a Turing machine by initializing each site with an additional marker indicating what type of location it is supposed to be. This would require a substantial increase in the number of states per site (perhaps to 100 or so), and we do not know how to use translational invariant nearest-neighbor rules for the QMA result. Instead we use position-dependent rules. This enables us to use essentially the same construction for both the adiabatic result and the QMA-completeness result, and reduces the required number of states in the adiabatic case.

In the beginning we put the quantum circuit we wish to simulate into a canonical form. Let  $n$  be the number of qubits in the circuit, labeled  $1 \dots n$  from left to right in a line. We assume that the circuit is initialized to the all  $|0\rangle$  state and consists of  $R$  rounds. Each round is composed of  $n - 1$  nearest-neighbor gates (some or all of which may be the identity gate); the first gate in a round acts on qubits 1 and 2, the second on qubits 2 and 3, and so on. Any quantum circuit can be put into this form with the number of rounds at most proportional to the number of gates. We will assume that the first round and the last round contain only identity gates, as we will need those two rounds for additional checking in the QMA-completeness result.

In the 1D arrangement, there will be a total of  $nR$  12-state particles, arranged on a line in  $R$  blocks of  $n$  qubits, each block corresponding to one round in the circuit. Roughly speaking, we imagine that in the beginning the qubits are all in the first block. The first round of quantum gates is performed one by one on these qubits, after which they are moved to the next block, where the second round of gates is performed, and so on. Once the qubits reach the last block and undergo the last round of gates, their state will be the final state of the circuit.

The main difficulty we will have to overcome is our inability to count. Since there are only a constant number of states per site, there is no way to directly keep track of how far we have moved. Our solution is to move the full set of qubits over only one space at a time. We keep moving over until the qubits reach the next block. We can tell that we have reached a new block, and are ready to perform a new set of gates, by making the transition rules different at the boundary between two blocks. When the qubits are correctly aligned with a block, a qubit state and a non-qubit state will be adjacent across a block boundary, whereas while the qubits are moving, adjacent pairs of particles which cross a block boundary will either have two qubit states or two non-qubit states.

We denote a block boundary by  $\cdot|$ . The 12 states in each site consist of 2-state subsystems (different versions of a qubit holding data), represented by elongated shapes (e.g.,  $\textcircled{R}$ ), and 1-state subspaces, represented by round shapes (e.g.,  $\textcircled{\circ}$ ). Two of the 2-state systems and two of the 1-state site types will be “flags” or “active” sites, which will be represented by dark shapes and can be thought of as pointers on the line that carry out the computation. Light-colored shapes represent a site that is inactive, waiting for the active site to come nearby. There will only be one active site in any (valid) configuration. We have the following types of states:

Inactive states	Flags (active states)
$\textcircled{R}$ : Qubits to right of active site	$\textcircled{\$}$ : Gate marker (moves right)
$\textcircled{L}$ : Qubits to left of active site	$\textcircled{\rightarrow}$ : Right-moving flag
$\textcircled{\circ}$ : Unborn (to right of all qubits)	$\textcircled{\leftarrow}$ : Left-moving flag, moves qubits right one space
$\textcircled{\times}$ : Dead (to left of all qubits)	$\textcircled{\odot}$ : Turning flag

The  $\textcircled{\$}$  and  $\textcircled{\rightarrow}$  active sites are qubit states. In the analogy to a Turing machine, they can be thought of as the head sitting on top of a qubit on the tape. The  $\textcircled{\leftarrow}$  and  $\textcircled{\odot}$  flags are one-dimensional subspaces, which sit between or next to the particles which are in qubit states.

**Definition 2.1** We use the term configuration to refer to an arrangement of the above types of states without regard to the value of the data stored in the qubit subsystems. Valid (or legal) configurations of the chain have the following structure:

$$\textcircled{\times} \cdots \textcircled{\times}(\mathbf{qubits})\textcircled{\circ} \cdots \textcircled{\circ},$$

where the  $\textcircled{\times} \cdots \textcircled{\times}$  string or  $\textcircled{\circ} \cdots \textcircled{\circ}$  string might not appear if the qubits are at one end of the computer. The  $(\mathbf{qubits})$  string consists of either  $n$  sites (for the first two possibilities) or  $n + 1$  sites (for the last three choices) and is of one of the following forms:

$$\begin{array}{c} \textcircled{L} \cdots \textcircled{L} \textcircled{\$} \textcircled{R} \cdots \textcircled{R} \\ \textcircled{L} \cdots \textcircled{L} \textcircled{\rightarrow} \textcircled{R} \cdots \textcircled{R} \\ \textcircled{L} \cdots \textcircled{L} \textcircled{\leftarrow} \textcircled{R} \cdots \textcircled{R} \textcircled{R} \\ \textcircled{\odot} \textcircled{R} \quad \cdots \quad \textcircled{R} \textcircled{R} \\ \textcircled{L} \quad \cdots \quad \textcircled{L} \textcircled{\odot} \end{array}$$

In the first three cases, either the  $\textcircled{L}$  string or the  $\textcircled{R}$  string might be absent when the active site is at one end of the qubits. (The  $\textcircled{\odot}$  flag is only needed at the left or right end of a string of qubits.) When the flag is  $\textcircled{\$}$ , the  $n$  qubit sites are lined up inside a single block; for all other values of the flag, the  $(\mathbf{qubits})$  string crosses a block boundary.

The initial configuration is  $|\textcircled{\$} \textcircled{R} \cdots \textcircled{R} | \textcircled{\circ} \cdots \textcircled{\circ} | \cdots | \textcircled{\circ} \cdots \textcircled{\circ} |$  where the  $\textcircled{\$}$  and the  $\textcircled{R}$  qubits are in the state corresponding to the input of the original circuit  $C_x$ . Following the transition rules below, the  $\textcircled{\$}$  sweeps to the right, performing gates as it goes. When it reaches the end of the qubit states at the border of the next block, it becomes a  $\textcircled{\odot}$ , which in turn creates a  $\textcircled{\leftarrow}$  flag which sweeps left, moving each qubit one space to the right in the process. When the  $\textcircled{\leftarrow}$  flag reaches the left end, it stops by turning into a  $\textcircled{\odot}$  which

then creates a  $\hat{\ominus}$  flag, which moves right through the qubits without disturbing them. The  $\hat{\ominus}$  flag hits the right end of the set of qubits and becomes a  $\hat{\odot}$ , which begins a  $\ominus$  flag moving left again. The  $\ominus$ - $\hat{\odot}$ - $\hat{\ominus}$ - $\hat{\odot}$  cycle continues until the qubits have all been moved past the next block boundary. Then we get a  $\hat{\odot}$ | $\hat{\text{R}}$  arrangement, which spawns a new  $\hat{\text{L}}$ , beginning the gate cycle again. The evolution stops when the qubits reach the last block boundary and the gate flag reaches the end of the line, i.e., the final configuration is  $|\otimes \cdots \otimes| \cdots |\otimes \cdots \otimes| \text{L} \cdots \text{L} \hat{\text{L}}$ .

We have the following transition rules. Something of the form  $XY$  is never across a block boundary, whereas  $X|Y$  is.  $|Y$  and  $X|$  represent the left and right end of the chain, respectively.

1. (Gate rule)  $\hat{\text{L}}\hat{\text{R}} \rightarrow \hat{\text{L}}\hat{\text{L}}$  (performing the appropriate gate between the two encoded qubits)
2. (Turning rules right side)  $\hat{\text{L}}|\hat{\odot} \rightarrow \hat{\text{L}}|\hat{\odot}$ ,  $\hat{\odot}\hat{\odot} \rightarrow \ominus\hat{\odot}$ ,  $\hat{\odot}|\hat{\odot} \rightarrow \ominus|\hat{\odot}$ ,  $\hat{\odot} \rightarrow \ominus|$
3. (Sweeping left rules)  $\text{L}\hat{\ominus} \rightarrow \ominus\hat{\text{R}}$ ,  $\text{L}|\hat{\ominus} \rightarrow \ominus|\hat{\text{R}}$
4. (Turning rules left side)  $\otimes\hat{\ominus} \rightarrow \otimes\hat{\odot}$ ,  $\otimes|\hat{\ominus} \rightarrow \otimes|\hat{\odot}$ ,  $|\hat{\ominus} \rightarrow |\hat{\odot}$ ,  $\hat{\odot}\hat{\text{R}} \rightarrow \otimes\hat{\ominus}$
5. (Sweeping right rules)  $\hat{\ominus}\hat{\text{R}} \rightarrow \hat{\text{L}}\hat{\ominus}$ ,  $\hat{\ominus}|\hat{\text{R}} \rightarrow \hat{\text{L}}|\hat{\ominus}$ ,  $\hat{\ominus}\hat{\odot} \rightarrow \hat{\text{L}}\hat{\odot}$
6. (Starting new round rule)  $\hat{\odot}|\hat{\text{R}} \rightarrow \otimes|\hat{\text{L}}$




Figure 1: A full cycle, with  $n = 3$ . Time goes down the first and then the second column.

We can see by inspection that these rules create the cycle described above; Figure 1 follows an example with  $n = 3$  step-by-step through a full cycle. A single cycle takes  $n(2n + 3)$  moves, so the full computation from the initial configuration to the final one takes a total of  $K = n(2n + 3)(R - 1) + n - 1$  steps. (The additional  $n - 1$  steps are for the gate flag to move through the final block.)

The only tricky part is to note the following fact:

**Fact 2.2** *For any configuration containing only one active site, there is at most one transition rule that applies to move forward one time step and at most one transition rule to move backwards in time by one step. For valid configurations, there is always exactly one forwards transition and exactly one backwards transition, except for the final and initial states, which have no future and no past, respectively.*

**Proof:** We can see this by noting that the transition rules only refer to the type of active site present and to the state of the site to the right (for  $\textcircled{R}$ ,  $\textcircled{L}$ , and  $\textcircled{O}$ ) or left (for  $\textcircled{L}$ ) of the active site (or left and right respectively for backwards transitions). For a configuration with only one active site, there is therefore at most one transition rule that applies. We can also see that for every legal configuration, there is indeed a transition rule that applies; the few cases that may appear to be missing cannot arise because of the constraints on alignment of the string of qubits with the blocks. ■

The illegal cases without forward transitions are  $\textcircled{R}|\textcircled{R}$ ,  $\textcircled{R}|\textcircled{O}$ ,  $\textcircled{L}|\textcircled{O}$ , and  $\textcircled{L}|$ , and the illegal cases without backwards transitions are  $\textcircled{L}|\textcircled{R}$ ,  $\textcircled{L}|\textcircled{L}$ ,  $\textcircled{L}|\textcircled{L}$ , and  $|$ . These arrangements will play an important role in Section 4.

*Remark:* For adiabatic computation, we can simplify this construction a bit. For the proof of QMA-completeness, we will need to distinguish the particles to the left of the active site from the particles to its right, but this is not necessary for universal adiabatic computation. Therefore, we may combine the  $\textcircled{L}$  and  $\textcircled{R}$  states into a single qubit subspace  $\textcircled{R}$ , and we may combine  $\textcircled{L}$  and  $\textcircled{O}$  into a single state  $\textcircled{O}$ . That leaves us with essentially the same construction, but with only 9 states per particle instead of 12.

### 3 Universality of adiabatic evolution in one dimension

We first recall the definition of adiabatic computation [FGGS00, FGG<sup>+</sup>01]. A quantum system slowly evolves from an easy to prepare ground state of some initial Hamiltonian  $H_0$  to the ground state of a tailored final Hamiltonian  $H_1$ , which encodes the problem. The state evolves according to Schrödinger's equation with a time-varying Hamiltonian  $H(s) = (1-s)H_0 + sH_1$ , with  $s = t/T$ , where  $t$  is the time and  $T$  is the total duration of the adiabatic computation. Measuring the final ground state in the standard basis gives the solution of the problem. The adiabatic theorem tells us that if  $T$  is chosen to be large enough, the final state will be close to the ground state of  $H_1$ . More specifically, we can define instantaneous Hamiltonians  $H(t/T)$ , and let  $g_{\min}$  be the minimum of the spectral gaps over all the Hamiltonians  $H(t/T)$ .  $T$  is required to be polynomial in  $1/g_{\min}$ .

To prove Theorem 1.2, given a general quantum circuit, we would like to design an efficient simulation of this circuit by adiabatic evolution on a one-dimensional line. The construction in Section 2 associates to any circuit  $C_x$  with  $n$  qubits a modified circuit  $\tilde{C}_x$  using  $nR$  9-state particles in one dimension (with some states merged according to the remark in the end of Sec. 2), together with transition rules which evolve the system for  $K = n(2n+3)R + n - 1$  steps through  $K+1$  configurations from initial to final. We denote the quantum state of the  $t$ -th configuration by  $|\gamma(t)\rangle$  for  $t = 0, \dots, K$ . Observe that the  $|\gamma(t)\rangle$  are orthogonal to each other, since each has a different configuration. We note that the last state  $|\gamma(K)\rangle$  contains the final state of the original circuit, encoded in the last block of particles.

Next we construct a one-dimensional adiabatic evolution whose final Hamiltonian has as its ground state the history state of this evolution, namely,

$$\frac{1}{\sqrt{K+1}} \sum_{t=0}^K |\gamma(t)\rangle.$$

To define  $H_1 \equiv H_{\text{prop}}$ , we simply translate the transition rules of the previous section to 9-state 2-local Hamiltonians. In other words, any transition rule which takes a state  $|\alpha\rangle$  to  $|\beta\rangle$  is translated to a Hamiltonian of the form  $\frac{1}{2}(|\alpha\rangle\langle\alpha| + |\beta\rangle\langle\beta| - |\alpha\rangle\langle\beta| - |\beta\rangle\langle\alpha|)$ . E.g., the first rule becomes  $\frac{1}{2}(|\textcircled{R}\textcircled{R}\rangle\langle\textcircled{R}\textcircled{R}| + U|\textcircled{R}\textcircled{L}\rangle\langle\textcircled{R}\textcircled{L}|U^\dagger - U|\textcircled{R}\textcircled{L}\rangle\langle\textcircled{L}\textcircled{R}| - |\textcircled{L}\textcircled{R}\rangle\langle\textcircled{R}\textcircled{L}|U^\dagger)$ , summed over a basis of states for the encoded qubits and over all nearest neighbor pairs of particles (recall we have combined  $\textcircled{R}$  and  $\textcircled{L}$ ). Here  $U$  is the gate corresponding to this location and acts on the two qubits encoded as  $|\textcircled{R}\textcircled{L}\rangle$ . Restricted to the eight-dimensional subspace spanned by  $|\textcircled{R}\textcircled{R}\rangle$  and  $|\textcircled{L}\textcircled{L}\rangle$  (each for values of  $|00\rangle$ ,  $|01\rangle$ ,  $|10\rangle$ , and  $|11\rangle$  for the

two qubits) for a single pair of particles, that would give the matrix

$$\frac{1}{2} \begin{pmatrix} I & -U \\ -U^\dagger & I \end{pmatrix},$$

with  $I$  the  $4 \times 4$  identity matrix.

Our initial Hamiltonian has the initial configuration  $|\gamma(0)\rangle$  as its ground state. To define  $H_0$ , we penalize all configurations that do not have  $\textcircled{3}$  in its  $|0\rangle$  state ( $|\textcircled{3}(0)\rangle$ ) in the first position:

$$H_0 = I - |\textcircled{3}(0)\rangle\langle\textcircled{3}(0)|_1$$

This Hamiltonian of course has a highly degenerate ground state, but the important point is that  $|\gamma(0)\rangle$  is the only state which satisfies  $H_0$  in the invariant subspace spanned by the  $|\gamma(t)\rangle$ . We assume we are able to select the correct ground state  $|\gamma(0)\rangle$  to be the actual initial state of the adiabatic computation. We could avoid the need for this by various methods, for instance by having an additional phase of adiabatic evolution from a simple non-degenerate Hamiltonian to  $H_0$ . Another option is to break the degeneracy of the ground state of  $H_1$  by using the full 12-state construction with the Hamiltonian described in Section 4. Even then, we either need to modify  $H_0$  to be non-degenerate (this approach requires a new analysis of the spectral gap, which we have not done), or to add an additional 13th start state. We outline the latter approach in Section 4.1.

To analyze the spectral gap of any Hamiltonian in the convex combination of  $H_0$  and  $H_1$ , we essentially follow the proof of [AvDK<sup>+</sup>04], using an improvement due to [DRS07]. We first observe that

**Claim 3.1** *The subspace  $\mathcal{K}_0$  spanned by the states  $|\gamma(t)\rangle$  is invariant under any of the Hamiltonians  $H(s)$ .*

This is easily seen to be true by the fact that it is invariant under the initial and final Hamiltonians. From this claim, it follows that throughout the evolution we are *inside*  $\mathcal{K}_0$ , so we only need to analyze the spectral gap in this subspace. We next move to the basis of  $\mathcal{K}_0$  spanned by the  $|\gamma(t)\rangle$  states. In this basis  $H_0$  restricted to  $\mathcal{K}_0$  looks like

$$H_0|_{\mathcal{K}_0} = \begin{pmatrix} 0 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \end{pmatrix},$$

and  $H_1$  restricted to  $\mathcal{K}_0$  becomes

$$H_1|_{\mathcal{K}_0} = \begin{pmatrix} \frac{1}{2} & -\frac{1}{2} & 0 & \dots & 0 \\ -\frac{1}{2} & 1 & -\frac{1}{2} & 0 & \ddots & \vdots \\ 0 & -\frac{1}{2} & 1 & -\frac{1}{2} & 0 & \ddots & \vdots \\ & \ddots & \ddots & \ddots & \ddots & \ddots & \\ \vdots & & 0 & -\frac{1}{2} & 1 & -\frac{1}{2} & 0 \\ & & & 0 & -\frac{1}{2} & 1 & -\frac{1}{2} \\ 0 & \dots & & 0 & -\frac{1}{2} & \frac{1}{2} \end{pmatrix}.$$

It can easily be seen that the history state (which in this basis is simply the all ones vector) is a zero eigenstate of this Hamiltonian. Finally we need to analyze the spectral gaps of all convex combinations of these two. This has been done in Sec. 3.1.2 of [AvDK<sup>+</sup>04] and simplified with improved constants in [DRS07]; we refer the reader there for the details. The result is that the spectral gap is at least  $1/[2(K+1)^2]$ , an inverse polynomial in  $n$  and  $R$ , which itself is an inverse polynomial in the number of gates in the original circuit. This proves Theorem 1.2.

## 4 1D QMA

The propagation Hamiltonian  $H_{\text{prop}}$  introduced in Section 3 is insufficient for QMA-completeness. Now we have a circuit  $C_x$  which checks the witness for a QMA problem, taking as input the witness and some ancilla qubits in the state  $|0\rangle$ . However, any correct history state for the circuit  $C_x$  will have zero eigenvalue for  $H_{\text{prop}}$ , even if the  $t = 0$  component of the history state is not correctly initialized, or if  $C_x$  does not accept the witness. Even worse,  $H_{\text{prop}}$  also has zero eigenvalue for any other state which is a uniform superposition of states connected by the transition rules, even if the superposition includes illegal configurations.

To solve these problems, we will introduce three new terms to the overall Hamiltonian  $H$ :

$$H = H_{\text{prop}} + H_{\text{init}} + H_{\text{final}} + H_{\text{penalty}}.$$

The initialization term  $H_{\text{init}}$  will constrain the initial state of the modified checking circuit  $\tilde{C}_x$  so that all ancilla qubits are initialized to  $|0\rangle$ , and the final Hamiltonian  $H_{\text{final}}$  verifies that the checking circuit does accept the witness as input.  $H_{\text{penalty}}$  will penalize illegal configurations using local constraints. As mentioned in the introduction, not all illegal configurations can be penalized directly. Instead, some of them will only be penalized because they evolve into a locally checkable illegal configuration.

To create the initialization term, we will assume without loss of generality that all the gates performed in the first block are the identity. We will use the gate flag  $\textcircled{\$}$  to check that qubits are correctly initialized instead of using it to do gates in the first block. Then we get the following Hamiltonian term:

$$H_{\text{init}} = \sum_i |\textcircled{\$}(1)\rangle\langle\textcircled{\$}(1)|_i.$$

The sum is taken over ancilla qubits  $i$  in their starting positions in the first block of  $n$  sites. (The remaining qubits in the first block of  $n$  sites are used to encode the potential witness.)  $H_{\text{init}}$  creates an energy penalty for any ancilla to be in the state  $|1\rangle$  when the gate flag passes over it. Since the gate flag sweeps through the whole block, this ensures that the ancilla qubits must be correctly initialized to  $|0\rangle$  or the state suffers an energy penalty.

Similarly, for the final Hamiltonian, we assume no gates are performed in the final block, and use  $\textcircled{\$}$  to check that the circuit accepts the output. That is, we get the following term in the Hamiltonian:

$$H_{\text{final}} = |\textcircled{\$}(0)\rangle\langle\textcircled{\$}(0)|_{\text{out}}$$

This causes an energy penalty if the output qubit “out” is in the state  $|0\rangle$  when the gate flag sweeps over it. In general, a correct history state for some potential input witness state will have, for the final block, a superposition of terms with the output qubit in the state  $|0\rangle$  and terms with the output qubit in the state  $|1\rangle$ , and the energy penalty is thus proportional to the probability that the circuit rejects the potential witness.

Now we move to  $H_{\text{penalty}}$ , which is a bit more involved. We will describe a set of local penalties that will enforce that each illegal configuration will be penalized either directly or because it will evolve into a configuration that will be penalized. We forbid the following arrangements:

1.  $\circ X, \circ|X$  ( $X$  is anything but  $\circ$ ),  $X\otimes, X|\otimes$  ( $X$  is anything but  $\otimes$ )
2. In the first block on the left:  $\circ, |\textcircled{R}$ . In the last block on the right:  $\otimes, \textcircled{L}$
3.  $\textcircled{R}X, \textcircled{R}|X$  ( $X$  is anything but  $\textcircled{R}$  or  $\circ$ ),  $X\textcircled{L}, X|\textcircled{L}$  ( $X$  is anything but  $\textcircled{L}$  or  $\otimes$ )
4.  $\otimes\textcircled{R}, \otimes|\textcircled{R}, \textcircled{L}\circ, \textcircled{L}|\circ, \otimes\circ, \otimes|\circ$
5.  $\textcircled{L}\textcircled{R}, \textcircled{L}|\textcircled{R}$

6. Any adjacent pair of active sites (e.g.,  $\hat{\otimes}\hat{\ominus}$  or  $\hat{\ominus}\hat{\otimes}$ ), with or without block boundaries
7.  $\hat{\otimes}\hat{\otimes}$ ,  $\hat{\otimes}\hat{\circ}$ ,  $\hat{\circ}\hat{\otimes}$ ,  $\hat{\otimes}\hat{\circ}$ ,  $\hat{\circ}\hat{\otimes}$ ,  $\hat{\otimes}\hat{\circ}$ ,  $\hat{\circ}\hat{\otimes}$  (but the first two are OK with a block boundary and the next four are OK without a block boundary)

Note that the arrangements forbidden in group 7 are precisely those missing a forward or backwards transition rule in the note following Fact 2.2.

We encode these rules into a penalty Hamiltonian in the straightforward way:

$$H_{\text{penalty}} = \sum_{XY} |XY\rangle\langle XY|_{i,i+1},$$

where the sum is taken over the forbidden arrangements XY listed above for all adjacent pairs of sites (tailored appropriately to the location of block boundaries and the first and last blocks).

**Claim 4.1** *A configuration that satisfies the rules in groups 1 through 6 is of one of the legal forms described in Section 2, or it is in one of three cases: (i) a (qubits) string of the form  $\hat{\circ}\cdots\hat{\circ}\hat{\otimes}\hat{\circ}\cdots\hat{\circ}$  (such a configuration is allowed in Section 2 only when  $\hat{\otimes}$  is all the way on one end of the (qubits) string), (ii) a (qubits) string that has incorrect length: different from  $n$  when the active site is  $\hat{\otimes}$  or  $\hat{\ominus}$  or different from  $n + 1$  when the active site is  $\hat{\ominus}$  or  $\hat{\otimes}$ , or (iii) a (qubits) string of length  $n$ , but with either an active site  $\hat{\ominus}$  and the (qubits) string aligned with a block boundary, or an active site  $\hat{\otimes}$  and the (qubits) string not aligned with a block boundary.*

Note that we do not need the rules in group 7 for this claim; they are used later to deal with the illegal states that cannot be checked locally.

**Proof:** The rules in group 1 enforce that the configuration is of the form  $\hat{\otimes}\cdots\hat{\otimes}(\text{qubits})\hat{\circ}\cdots\hat{\circ}$  by ensuring that all unborn ( $\hat{\circ}$ ) states are at the far right and all dead ( $\hat{\otimes}$ ) states are at the far left. Group 2 ensures that the particles are not all  $\hat{\circ}$  or all  $\hat{\otimes}$  (and in fact ensures that there is at least one block of  $n$  composed of something else). Group 3 ensures that within the (qubits) string, any  $\hat{\circ}$  qubits are on the right end and any  $\hat{\circ}$  qubits are on the left end. Groups 2 and 4 guarantee that the (qubits) string is not empty and does not consist of only  $\hat{\circ}$  or  $\hat{\circ}$  qubits. Group 5 then checks that the (qubits) string does not directly jump from  $\hat{\circ}$  to  $\hat{\circ}$ , so there is at least one active site in between, and the rules in group 6 ensure that there cannot be more than one adjacent active site. Since, by groups 1 and 3, an active site can only occur to the right of all  $\hat{\otimes}$  and  $\hat{\circ}$  sites and to the left of all  $\hat{\circ}$  and  $\hat{\circ}$  sites, all active sites are together in a group, so the rules in groups 5 and 6 ensure that there is exactly one active site. Comparing with Definition 2.1, we see that this leaves only the exceptions listed in the claim. ■

The three exceptions in the above claim cannot be ruled out directly via a local check, since counting cannot be done using local constraints.<sup>4</sup> We can only rule out the exceptions by considering both the penalty Hamiltonian  $H_{\text{penalty}}$  and the propagation Hamiltonian  $H_{\text{prop}}$ .

In order to do so, we first break down the full Hilbert space into subspaces which are invariant under both  $H_{\text{penalty}}$  and  $H_{\text{prop}}$ . Let us consider the minimal sets of configurations such that the sets are invariant under the action of the transition rules. This defines a partition of the set of configurations. Given such a minimal set  $S$ , we consider the subspace  $\mathcal{K}_S$  spanned by all the configurations in  $S$ . The penalty Hamiltonian is diagonal in the basis of configurations, so  $\mathcal{K}_S$  is invariant under  $H_{\text{penalty}}$ , and the set  $S$  is closed under the action of  $H_{\text{prop}}$ , so  $\mathcal{K}_S$  is invariant under  $H_{\text{prop}}$  as well. The space  $\mathcal{K}_S$  belongs to one of three types:

<sup>4</sup>If we are willing to add a 13th state, we could split the  $\hat{\otimes}$  state into two states, one to turn around on the left and one to turn around on the right. Then we could rule out  $\hat{\circ}\cdots\hat{\circ}\hat{\otimes}\hat{\circ}\cdots\hat{\circ}$  with local rules. However, this strategy will not work for strings with the wrong number of qubit states.

1. All states in  $S$  are legal. (We call this subspace  $\mathcal{K}_0$ .)
2. All states in  $S$  are illegal, but are all locally detectable, namely, none of them belong to the exceptions of Claim 4.1.
3. All states in  $S$  are illegal, but at least one of them is not locally detectable.

Note that the set  $S$  cannot contain both legal and illegal configurations, since legal configurations do not evolve to illegal ones, and vice-versa. We want to show that subspaces of type 2 or 3 have large energy. The main challenge will be to give a lower bound on the energy of a space  $\mathcal{K}_S$  of type 3, which we will do by lower bounding the fraction of configurations in  $S$  that violate one of the conditions 1-7 above. We want to prove:

**Lemma 4.2 (Clairvoyance lemma)** *The minimum eigenvalue of  $H_{\text{prop}} + H_{\text{penalty}}$ , restricted to any  $\mathcal{K}_S$  of type 2 or 3, is  $\Omega(1/K^3)$ , where  $K$  is the number of steps in  $\tilde{C}_x$ .*

**Proof:** To prove the theorem, we deal separately with each type of invariant subspace. Type 2 is straightforward: the full subspace  $\mathcal{K}_S$  has energy of at least 1, due to  $H_{\text{penalty}}$ .

We now focus on type 3. We observe several properties of type 3 subspaces. First, notice that the number of qubit states ( $\textcircled{L}$ ,  $\textcircled{R}$ ,  $\textcircled{\$}$ , and  $\textcircled{\ominus}$ ) and the number of active sites ( $\textcircled{\ominus}$ ,  $\textcircled{\$}$ ,  $\textcircled{\ominus}$ , and  $\textcircled{\ominus}$ ) are conserved under all transition rules. Since  $S$  is a minimal set preserved by the transition rules, all the configurations in  $S$  therefore contain the same number of qubit states. Since  $S$  is of type 3, it must contain an undetectable illegal configuration, such as a (**qubits**) string of the wrong length, and by Claim 4.1 has exactly one active site somewhere in the string. By Fact 2.2, there is at most one forwards transition rule and one backwards transition rule that applies to each configuration in  $S$ . Note also that configurations with one active site which is  $\textcircled{\ominus}$  or  $\textcircled{\ominus}$  always have both a forward and backwards transition rule available. Therefore, if  $S$  is of type 3 it must contain at least one configuration with a  $\textcircled{\ominus}$  or  $\textcircled{\$}$  active site.

Going further, we claim that a fraction of at least  $\Omega(1/n^2R)$  of the configurations in  $S$  violate one of the rules in groups 1 through 7 above. To see this, we divide into three cases:

- $S$  contains a configuration with a qubit string of the form  $\textcircled{L} \cdots \textcircled{L} \textcircled{\ominus} \textcircled{R} \cdots \textcircled{R}$ : In this case, the forward transition rule that applies to the illegal configuration is either  $\textcircled{\ominus} \textcircled{R} \rightarrow \textcircled{\times} \textcircled{\ominus}$  or  $\textcircled{\ominus} \textcircled{R} \rightarrow \textcircled{\times} \textcircled{\$}$ . This gives us a configuration which can be locally seen as illegal, such as  $\textcircled{L} \cdots \textcircled{L} \textcircled{\times} \textcircled{\ominus} \textcircled{R} \cdots \textcircled{R}$ . This violates group 1, as  $\textcircled{\times}$  sites should be to the left of all other kinds of sites. Therefore  $S$  contains at least  $1/2$  locally checkable illegal configurations. In fact, it is a much larger fraction, as all the backwards transition rules and further forward transition rules will also produce locally checkable illegal states.
- $S$  contains a configuration with a  $\textcircled{\$}$  active site: Via forward transitions, the  $\textcircled{\$}$  will move to the right until it hits either the right end of the (**qubits**) string or until it hits a block boundary. If it does not encounter both together, it has reached one of the illegal configurations forbidden in group 7. Via backwards transitions, the  $\textcircled{\$}$  moves to the left until it hits either the left end of the (**qubits**) string or until it hits a block boundary. Again, if it does not encounter both together, we must now be in one of the configurations forbidden in group 7. If the right and left ends of the (**qubits**) string are both lined up with block boundaries, and there are no other block boundaries in between, then the (**qubits**) string contains exactly  $n$  sites and is correctly aligned with the block boundary. Therefore it is actually a legal state, which cannot occur in  $S$  of type 3. Thus, we must have one of the locally checkable illegal states, and it must occur after at most  $n$   $\textcircled{\$}$  transitions (in either direction). The  $\textcircled{\$}$  configurations themselves, however, could be part of a larger cycle with transitions involving  $\textcircled{\ominus}$ ,  $\textcircled{\ominus}$ , and  $\textcircled{\ominus}$  active sites. However, the length of the (**qubits**) string is at most  $nR$ , and we can have at most



$n$   $\ominus$ - $\ominus$ - $\ominus$ - $\ominus$  cycles before getting a  $\otimes$ , so the locally checkable illegal states comprise a fraction at least  $\Omega(1/n^2R)$  of all the states in  $S$ .

- $S$  contains no configuration with a  $\otimes$  active site:  $S$  still must contain configurations with  $\ominus$ . These configurations can form part of a normal cycle, but the normal cycle eventually leads to a transition to a  $\otimes$  configuration. For forward transitions, this occurs when the left end of the (**qubits**) string is aligned with a block boundary, and for backwards transitions, it occurs when the right end of the (**qubits**) string is aligned with a block boundary. The only way to avoid this happening is for  $S$  to have a  $\ominus$  configuration with no forward transition and another with no backwards transition. Since the only such configurations contain one of the arrangements  $\otimes|\ominus$ ,  $\ominus|\ominus$ ,  $|\ominus$ , or  $\ominus|$ , all of which are forbidden by the rules in group 7,  $S$  must contain a locally checkable illegal configuration. Again, the length of the (**qubits**) string is at most  $nR$ , and we can have at most  $n$   $\ominus$ - $\ominus$ - $\ominus$ - $\ominus$  cycles in a row, so the locally checkable illegal states comprise a fraction at least  $\Omega(1/n^2R)$  of all the states in  $S$ .

We can now invoke Lemma 14.4 from [KSV02] to lower bound the energy of the overall Hamiltonian for a type 3 subspace:

**Lemma 4.3** *Let  $A_1, A_2$  be nonnegative operators, and  $L_1, L_2$  their null subspaces, where  $L_1 \cap L_2 = \{0\}$ . Suppose further that no nonzero eigenvalue of  $A_1$  or  $A_2$  is smaller than  $v$ . Then*

$$A_1 + A_2 \geq v \cdot 2 \sin^2 \theta / 2,$$

where  $\theta = \theta(L_1, L_2)$  is the angle between  $L_1$  and  $L_2$ .

In our case,  $A_1$  is the propagation Hamiltonian  $H_{\text{prop}}$ , and its null subspace, restricted to  $K_S$ , consists of equal superpositions over all configurations in the invariant subspace  $S$ . (There are multiple such states, with different values of the encoded qubit states.)  $A_2$  is the penalty Hamiltonian  $H_{\text{penalty}}$ , diagonal in the basis of configurations. Then  $\sin^2 \theta$  is the projection (squared) of the superposition of all shapes on the subspace of locally checkable illegal configurations; that is, it is the fraction of locally checkable illegal configurations in the invariant set. The minimum nonzero eigenvalue of  $H_{\text{penalty}}$  is 1, but (as in [KSV02]) the minimum nonzero eigenvalue of  $H_{\text{prop}}$  is  $\Omega(1/K^2)$ , where  $K = n(2n+3)(R-1) + n - 1$  is the number of steps in  $\tilde{C}_x$ . Thus, if  $S$  is a set containing a configuration which is illegal but cannot be locally checked, all states in  $K_S$  have an energy at least  $\Omega(1/K^3)$ . ■

We can now prove Theorem 1.4. We start by assuming the circuit  $C_x$  accepts correct witnesses and rejects incorrect witnesses with a probability exponentially close to 1. This can be achieved, for instance, by checking multiple copies of the witness. Then we will show that if there exists a witness which is accepted by  $C_x$  with probability at least  $1 - O(1/K^3)$ , then there is a state with energy at most  $O(1/K^4)$ , whereas if all possible witnesses are only accepted by  $C_x$  with probability at most  $1/K^3$ , then all states have energy at least  $\Omega(1/K^3)$ .

We already know that on subspaces of type 2 and 3, the minimum eigenvalue of  $H_{\text{prop}} + H_{\text{penalty}}$  is  $\Omega(1/K^3)$ . We therefore restrict attention to the subspace  $\mathcal{K}_0$  of type 1, built from only legal configurations, and can hence ignore  $H_{\text{penalty}}$ . From this point on, the proof follows [KSV02], but we include it for completeness.

The type 1 space  $\mathcal{K}_0$  contains only valid history states. However, some have incorrect ancilla inputs and are penalized by  $H_{\text{init}}$ , while others include an incorrect witness input, and are penalized by  $H_{\text{final}}$ . Only history states corresponding to a correct witness have low energy.

If there exists a witness which is accepted by  $C_x$  with probability at least  $1 - 1/K^3$ , then the history state  $|\Phi\rangle = \frac{1}{\sqrt{K+1}} \sum_t |\phi_t\rangle$  for that witness will have small enough energy:

$$\langle \Phi | H_{\text{prop}} + H_{\text{penalty}} + H_{\text{init}} | \Phi \rangle = 0,$$

so we only need to calculate  $\langle \Phi | H_{\text{final}} | \Phi \rangle$ . The only term in  $|\Phi\rangle$  that contributes is the one with  $\textcircled{3}$  on the output qubit in the last block, which has a coefficient  $1/\sqrt{K+1}$ . Therefore, the energy due to  $H_{\text{final}}$  is  $1/(K+1)$  times the probability that the witness is rejected by  $C_x$ . That is, it is  $O(1/K^4)$ .

Now let us consider what happens for a “no” instance. The following lemma will complete the proof of Theorem 1.4:

**Lemma 4.4** *If there is no witness that causes the circuit to accept with probability larger than  $1/K^3$ , then every state in  $\mathcal{K}_0$  has energy at least  $\Omega(1/K^3)$ .*

**Proof:** We will invoke Lemma 4.3 once more, with  $A_1 = H_{\text{prop}}$ , and  $A_2 = H_{\text{init}} + H_{\text{final}}$ . From now on, all the operators we discuss will be restricted to  $\mathcal{K}_0$ . The null subspace of the propagation Hamiltonian  $H_{\text{prop}}$  contains states which correspond to a history  $|\psi\rangle = \frac{1}{\sqrt{K+1}} \sum_t |\phi_t\rangle$  of the circuit  $\tilde{C}_x$  for some input state  $|\phi_0\rangle$ , not necessarily correct. Fix such a history state  $|\psi\rangle$ . We will upper bound its projection squared on the null space of  $A_2$ .

Let us first study the structure of this null space. The null space of  $H_{\text{init}}$  is spanned by states for which either the  $\textcircled{3}$  is not on an ancilla qubit in the first block or the  $\textcircled{3}$  is on an ancilla qubit, but the ancilla qubit is a  $|0\rangle$ . Let  $\Pi_{\text{init}}$  be the projection on the null space of  $H_{\text{init}}$ . The null space of  $H_{\text{final}}$  contains states for which either the  $\textcircled{3}$  is not on the output qubit in the last block or the  $\textcircled{3}$  is on the output qubit and the output qubit is a  $|1\rangle$ . Let  $\Pi_{\text{final}}$  be the projection on the null space of  $H_{\text{final}}$ . The two projectors commute, so the projector onto the null subspace of  $H_{\text{init}} + H_{\text{final}}$  is  $\Pi_{\text{init}}\Pi_{\text{final}}$ . We would thus like to upper bound

$$\langle \psi | \Pi_{\text{init}} \Pi_{\text{final}} | \psi \rangle.$$

To this end, we expand the input state  $|\phi_0\rangle$  for  $|\psi\rangle$  into  $|\phi_0\rangle = \alpha|\phi_0^v\rangle + \beta|\phi_0^i\rangle$ , where  $|\phi_0^v\rangle$  is a state with all valid ancilla input qubits and  $|\phi_0^i\rangle$  is an orthogonal state consisting of a superposition of states with one or more invalid ancilla input qubits. Note that  $|\psi\rangle = \alpha|\psi^v\rangle + \beta|\psi^i\rangle$ , where  $|\psi^v\rangle$  is the history state for input  $|\phi_0^v\rangle$  and  $|\psi^i\rangle$  is the history state for input  $|\phi_0^i\rangle$ . We have:

$$\langle \psi | \Pi_{\text{init}} \Pi_{\text{final}} | \psi \rangle = |\alpha|^2 \langle \psi^v | \Pi_{\text{init}} \Pi_{\text{final}} | \psi^v \rangle + |\beta|^2 \langle \psi^i | \Pi_{\text{init}} \Pi_{\text{final}} | \psi^i \rangle + 2\text{Re}(\alpha^* \beta \langle \psi^v | \Pi_{\text{init}} \Pi_{\text{final}} | \psi^i \rangle),$$

where we have used the fact that the two projectors commute, so  $\Pi_{\text{init}}\Pi_{\text{final}}$  is Hermitian.

We now bound each of the three terms. For the first term, note that  $\langle \psi^v | \Pi_{\text{init}} = \langle \psi^v |$ . We might as well assume the output qubit is the very last qubit, so the  $\textcircled{3}$  reaches it only at  $t = K$ , and we have that

$$\langle \psi^v | \Pi_{\text{final}} | \psi^v \rangle = \frac{1}{K+1} [K + \langle \phi_K^v | \Pi_{\text{final}} | \phi_K^v \rangle] = \frac{K+p}{K+1},$$

where  $p$  is the probability that the circuit  $C_x$  accepts  $|\phi_0^v\rangle$ . We thus find that the first term is at most  $|\alpha|^2(K + O(1/K^3))/(K+1)$ .

For the second term, it suffices to upper bound  $\langle \psi^i | \Pi_{\text{init}} | \psi^i \rangle$  since

$$\langle \psi^i | \Pi_{\text{init}} | \psi^i \rangle = \|\Pi_{\text{init}} |\psi^i\rangle\|^2 \geq \|\Pi_{\text{final}} \Pi_{\text{init}} |\psi^i\rangle\|^2 = \langle \psi^i | \Pi_{\text{init}} \Pi_{\text{final}} | \psi^i \rangle,$$

again using the fact that the two projectors commute. Since the computation  $\tilde{C}_x$  begins by sweeping a  $\textcircled{3}$  through the first  $n$  sites, and thereafter never has  $\textcircled{3}$  within the first  $n$  sites, only the first  $n$  states  $|\phi_t\rangle$  can possibly be outside the null space of  $H_{\text{init}}$ . Also, note that  $\langle \phi_t | \Pi_{\text{init}} | \phi_{t'} \rangle = 0$  unless  $t = t'$ . We can thus write

$$\langle \psi^i | \Pi_{\text{init}} | \psi^i \rangle = \frac{1}{K+1} \left[ (K+1-n) + \sum_{t=0}^{n-1} \langle \phi_t^i | \Pi_{\text{init}} | \phi_t^i \rangle \right].$$

We have  $\langle \phi_t | \Pi_{\text{init}} | \phi_t \rangle = 0$  if the  $t^{\text{th}}$  qubit of the input is an ancilla qubit and is in the state  $|1\rangle$ , and  $\langle \phi_t | \Pi_{\text{init}} | \phi_t \rangle = 1$  otherwise. Since the  $|\phi_t^i\rangle$  states are superpositions of terms with at least one incorrect ancilla input qubit, we have

$$\langle \psi^i | \Pi_{\text{init}} | \psi^i \rangle \leq K/(K+1).$$

For the third term, we have

$$\begin{aligned} |\text{Re}(\alpha^* \beta \langle \psi^v | \Pi_{\text{init}} \Pi_{\text{final}} | \psi^i \rangle)| &\leq |\langle \psi^v | \Pi_{\text{init}} \Pi_{\text{final}} | \psi^i \rangle| \\ &= |\langle \psi^v | \Pi_{\text{final}} | \psi^i \rangle|. \end{aligned}$$

Note, however, that  $\langle \phi_t^v | \phi_t^i \rangle = 0$ , since  $|\phi_t\rangle$  is related to  $|\phi_0\rangle$  by a unitary transformation and  $\langle \phi_0^v | \phi_0^i \rangle = 0$ . This gives us

$$\begin{aligned} |\langle \psi^v | \Pi_{\text{final}} | \psi^i \rangle| &= \frac{1}{K+1} |\langle \phi_K^v | \Pi_{\text{final}} | \phi_K^i \rangle| \\ &\leq \frac{1}{K+1} \|\Pi_{\text{final}} | \phi_K^v \rangle\| \\ &= O(K^{-5/2}), \end{aligned}$$

The above expression might not be 0 since both  $|\psi^v\rangle$  and  $|\psi^i\rangle$  might have a non-zero component which is accepted by the circuit, but since  $|\phi^v\rangle$  is only accepted rarely (with probability  $O(1/K^3)$ ), it cannot be too large.

Summing up all contributions, we have:

$$\langle \psi | \Pi_{\text{init}} \Pi_{\text{final}} | \psi \rangle \leq |\alpha|^2 \frac{K + O(1/K^3)}{K+1} + |\beta|^2 \frac{K}{K+1} + O(K^{-5/2}) = \frac{K}{K+1} + O(K^{-5/2}).$$

From this we can conclude that  $\sin^2(\theta)$  is at least  $\Omega(1/K)$ , where  $\theta$  is the angle between the null spaces of  $H_{\text{prop}}$  and  $H_{\text{init}} + H_{\text{final}}$ . Lemma 4.3 then tells us that every state has energy  $\Omega(1/K^3)$ . ■

## 4.1 Non-degenerate Universal Adiabatic Hamiltonian

If we wish to create a universal adiabatic quantum computer with a non-degenerate ground state, the Hamiltonians from Section 3 do not suffice. In fact, even if we used a 12-state construction with the final Hamiltonian  $H_1$  equal to  $H$  from the QMA-completeness construction above, then  $H_1$  would be non-degenerate, but the initial Hamiltonian  $H_0$  in Section 3 still is not. We sketch here a construction that adds an additional 13th state to make sure that the Hamiltonian remains non-degenerate throughout the adiabatic evolution.

The new state is a start state  $\textcircled{S}$ , which serves as the state for the input qubits in the initial configuration. We will add a set of terms to  $H_0$  and  $H_1$  (which also includes  $H_{\text{init}}$ ,  $H_{\text{final}}$ , and  $H_{\text{penalty}}$ , in addition to  $H_{\text{prop}}$ ). These terms will guarantee that the unique ground state of  $H_0$  is  $|\textcircled{S} \textcircled{S} \cdots \textcircled{S} | \textcircled{O} \cdots \textcircled{O} \cdots | \textcircled{O} \cdots \textcircled{O} \rangle$ . Furthermore,  $H_0$  will have the property that  $H_0 |\gamma(t)\rangle = 1$  for any  $t > 0$ , which allows us to use the calculation of the gap given in Section 3. To do this, we add a term that enforces the condition that if the first particle is in state  $\textcircled{S}$ , then the second particle is in state  $\textcircled{S}$ . This is accomplished by forbidding any configuration that has  $\textcircled{S}$  in the first location and any state except  $\textcircled{S}$  in the second. For example, to forbid states  $\textcircled{S} \textcircled{O}$  in particles 1 and 2, we add the term  $|\textcircled{S} \textcircled{O}\rangle \langle \textcircled{S} \textcircled{O}|_{12}$ . Similarly, for  $i = 2, \dots, n-1$ , we enforce that if particle  $i$  is in state  $\textcircled{S}$  then particle  $i+1$  must be in state  $\textcircled{S}$  as well. Then we add a term that says if particle  $n$  is in state  $\textcircled{S}$  then particle  $n+1$  is in state  $\textcircled{O}$ . For  $i > n$ , add terms that enforce that if particle  $i$  is in state  $\textcircled{O}$ , then so is particle  $i+1$ . Similar terms should also be added to  $H_1$  to assure that its ground state remains non-degenerate when the new start state is added. Finally, we need to slightly alter the transition rules used to build  $H_{\text{prop}}$ . For  $i = 1, \dots, n-1$ , replace the rule  $\textcircled{S} \textcircled{R} \rightarrow \textcircled{L} \textcircled{S}$  in locations  $i$  and  $i+1$  with  $\textcircled{S} \textcircled{S} \rightarrow \textcircled{L} \textcircled{S}(0)$ . With this change, Fact 2.2 still holds, with the small caveat that valid configurations

now must have any qubits in the first block set to  $|0\rangle$ . Note that none of the configurations in the subspace spanned by the  $|\gamma(t)\rangle$  violate any of the extra conditions. Furthermore, the additional terms guarantee the uniqueness of the ground state of  $H_0$  given above. In order to establish the non-degeneracy of any of the Hamiltonians  $H(s)$ , we need to prove that there are no low energy states outside of the subspace spanned by the  $|\gamma(t)\rangle$ , which follows from the Clairvoyance lemma and a simplified version of Lemma 4.4.

## 5 Discussion and Open Problems

We have shown that 1-dim 12-state Hamiltonians can be used for both universal adiabatic quantum computation and to produce very difficult, QMA-complete problems.

A similar result holds also for a related quantum complexity class, called QCMA, which is the subclass of QMA where the witnesses are restricted to be classical. To see this, observe first that our reduction from an arbitrary QMA language to the 1-DIM 12-STATE HAMILTONIAN problem is witness-preserving, at least once the acceptance probability has been amplified. Given a witness for the original QMA language, we can, in fact, efficiently construct on a quantum computer a witness for the corresponding instance of 1-DIM 12-STATE HAMILTONIAN using the adiabatic algorithm. This implies that if the witness for the original problem is efficiently constructible (which means we may as well assume it is a classical bit string describing the circuit used to construct the quantum witness), then the witness for 1-DIM 12-STATE HAMILTONIAN is also efficiently constructible. Thus, we have also shown that the sub-language of 1-DIM 12-STATE HAMILTONIAN which has the additional promise of an efficiently-constructible ground state is complete for QCMA.

There remain many interesting related open problems. Clearly, it is interesting to ask whether the size of the individual particles in the line can be further decreased, perhaps as far as qubits. We could, of course, interpret our 12-state particles as sets of 4 qubits, in which case our Hamiltonian becomes 8-local, with the sets of 8 interacting qubits arranged consecutively on a line. The perturbation-theory gadgets used to convert 8-local interactions to 2-local ones do not work in a one-dimensional system: The pairs of interacting qubits form a graph, which needs to have degree at least 3 (or 4 for some of the gadgets used by [OT05]). If we do apply the approach of [OT05], however, to the system described here, we get a 2-dim Hamiltonian on a strip of qubits of constant width. The constant is rather large, but this still constitutes an improvement over the result of [OT05].

Another approach to decreasing the dimensionality of the individual particles is to find a new protocol for removing the explicit reference to time from the circuit  $C_x$ . It seems likely that some further improvement is possible in this regard, but it is unlikely that an improved protocol by itself can take us all the way to qubits, as we need additional states to provide the control instructions. If it is possible to have universal adiabatic quantum computation and QMA-completeness with 1-dim 2-state Hamiltonians, we will probably need new techniques to prove it. It may also be that there is a transition at some intermediate number of states between 2 and 12 for which universality and QMA-completeness become possible. This would be analogous to the classical 2-dimensional Ising spin problem without magnetic field, which is in P for a single plane of bits, but is NP-complete when we have two layers of bits [Bar82].

Another interesting line of open questions is to investigate the energy gap. There are two energy gaps of relevance, both interesting. One is the “promise gap”  $\Delta$  in the definition of QMA. We have shown that we have QMA-completeness when  $\Delta$  is polynomially small relative to the energy per term. This can easily be improved to a constant value of  $\Delta$  by amplification:  $t$  copies of the ground state will either have energy less than  $tE$  or above  $t(E + \Delta)$ , amplifying the promise gap to  $t\Delta$ . A more interesting question is whether  $\Delta$  can be made a constant fraction of the total energy available in the problem, the *largest* eigenvalue of  $H$ ; if so, that would constitute a quantum version of the PCP theorem. Hastings and Terhal [HT] have argued that it is not possible to do this in any constant number of dimensions unless QMA = P: To approximate the

ground state energy, we can divide the system into blocks of a constant size, and diagonalize the Hamiltonian within each block. We can then consider the tensor product of the ground states for each separate block; since such a state ignores the energy due to Hamiltonian terms interacting different blocks, it will not be the true ground state, but its energy can differ from the ground state energy by at most the surface area of each block times the number of blocks (normalizing the maximum energy per term to be 1), whereas the maximum eigenvalue of  $H$  is roughly proportional to the *volume* of each block times the number of blocks. Thus, the tensor product state approximates the ground state energy up to a constant fraction of the total energy, and the fraction can be made arbitrarily small by increasing the size of each block. This gives a polynomial-time classical algorithm for approximating the ground state energy to this accuracy. When the interactions in  $H$  are not constrained by dimensionality, the argument breaks down, so this quantum version of PCP remains an interesting open problem for general  $r$ -state  $k$ -local Hamiltonians.

We can also look at the spectral gap, the gap between the ground state and the first excited state of  $H$ . For adiabatic computation, we are interested in the minimal spectral gap over the course of the computation. We have shown that universal adiabatic quantum computation is possible if the spectral gap is polynomially small relative to the energy per term. What happens if it is bounded below by a constant? Hastings has recently shown [Has] that it is possible to efficiently classically simulate the adiabatic evolution of a 1-dim  $r$ -state Hamiltonian system with constant spectral gap. We should therefore not expect to be able to perform universal adiabatic quantum computation with such Hamiltonians, as that would imply  $\text{BQP} = \text{P}$ . Hastings' argument builds on his earlier paper [Has07], which showed that the ground state of a gapped 1-dim  $r$ -state Hamiltonian has an efficient classical description as a matrix product state. By repeatedly updating the matrix product state description, one can then keep track of the adiabatic evolution with resources polynomial in the path length in parameter space and the inverse error of the approximation. For  $d$ -dim  $r$ -state Hamiltonians with constant spectral gap and  $d \geq 2$ , the question remains open in general, although Osborne has proven [Os06] that such systems can be efficiently classically simulated for logarithmic time.

For the QMA-completeness problem, we know very little about the possible values of the spectral gap. If the spectral gap is constant relative to the energy per term (or even  $\Omega(1/\ln \ln nR)$ ), the ground state has a matrix product state representation [Has07] and the problem is in NP. In our construction, there are enough low-energy states that the spectral gap cannot be much larger than the promise gap  $\Delta$ , but it might be much smaller. In the “no” instances, there are likely very many states which violate only a small number of transition rules, penalty terms, or initial conditions, and these have energy just above  $E + \Delta$ , so likely the spectral gap is exponentially small in the “no” instances. We know less about the size of spectral gap in the “yes” instances. States which do not correspond to valid histories have an energy at least  $\Delta$  larger than the ground state, but unfortunately, there may be different valid histories with energies less than  $E + \Delta$  but above the ground state energy. The difficulty is that the original problem might have many potential witnesses. Some may be good witnesses, accepted with high probability, whereas others may be mediocre witnesses, accepted with a probability near  $1/2$ . There could, in fact, be a full spectrum of witnesses with only exponentially small gaps between their acceptance probabilities. In order to show that the “yes” instances can be taken to have a spectral gap which is at least inverse polynomial in the system size, we would need a quantum version of the Valiant-Vazirani theorem [VV86], which would say that we can always modify a QMA problem to have a unique witness accepted with high probability.

## 6 Acknowledgements

This work was partly done while three of the authors (D. A., D. G., and J. K.) were visiting the Institute Henri Poincaré in Paris, and we want to thank the IHP for its hospitality. We wish to thank Daniel Fisher, Matt Hastings, Lev Ioffe, Tobias Osborne, Oded Regev, and Barbara Terhal for helpful discussions and comments and Oded for help with the design of the state icons.

## References

- [ACdF88] B. Apolloni, C. Carvalho, and D. de Falco. Quantum stochastic optimization. *Stochastic Processes and their Applications*, 33(5):233–244, 1988.
- [ACdF90] B. Apolloni, N. Cesa-Bianchi, and D. de Falco. A numerical implementation of ”quantum annealing”. In *Stochastic Processes, Physics and Geometry: Proceedings of the Ascona-Locarno Conference*, pages 97–111. World Scientific, River Edge, NJ, 1990.
- [AGIK07] Dorit Aharonov, Daniel Gottesman, Sandy Irani, and Julia Kempe. The power of quantum systems on a line. In *FOCS*, pages 373–383, 2007.
- [AGK07] D. Aharonov, D. Gottesman, and J. Kempe. The power of quantum systems on a line, 2007. arXiv:0705.4077 [quant-ph].
- [AvDK<sup>+</sup>04] D. Aharonov, W. van Dam, J. Kempe, Z. Landau, S. Lloyd, and O. Regev. Adiabatic quantum computation is equivalent to standard quantum computation. In *Proc. 45th FOCS*, pages 42–51. IEEE, 2004.
- [Bar82] F. Barahona. On the computational complexity of Ising spin glass models. *Jour. of Phys. A: Math. and Gen.*, 15:3241–3253, 1982.
- [BY86] K. Binder and A. P. Young. Spin glasses: Experimental facts, theoretical concepts, and open questions. *Rev. Mod. Phys.*, 58:801–976, 1986.
- [CFP02] A. Childs, E. Farhi, and J. Preskill. Robustness of adiabatic quantum computation. *Phys. Rev. A*, 65:012322, 2002.
- [DRS07] P. Deift, M. B. Ruskai, and W. Spitzer. Improved gap estimates for simulating quantum circuits by adiabatic evolution. *Quantum Information Processing*, 6(2):121–125, 2007.
- [Fey85] R. Feynman. Quantum mechanical computers. *Optics News*, 11:11–21, February 1985.
- [FGG<sup>+</sup>01] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda. A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem. *Science*, 292(5516):472–476, 2001.
- [FGGS00] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser. Quantum computation by adiabatic evolution, 2000. arXiv:quant-ph/0001106.
- [Fis95] D. S. Fisher. Critical behavior of random transverse-field Ising spin chains. *Phys. Rev. B*, 51(10):6411–6461, 1995.
- [Has] M. Hastings. Personal communication.
- [Has07] M. Hastings. An area law for one dimensional quantum systems, 2007. arXiv:0705.2024v2 [quant-ph].
- [HT] M. Hastings and B. Terhal. Personal communication.
- [Ira07] S. Irani. The complexity of quantum systems on a one-dimensional chain, 2007. arXiv:0705.4067v1 [quant-ph].

- [JFS06] S. P. Jordan, E. Farhi, and P. W. Shor. Error-correcting codes for adiabatic quantum computation. *Phys. Rev. A*, 74:052322, 2006.
- [JWZ07] Dominik Janzing, Pawel Wocjan, and Shengyu Zhang. A single-shot measurement of the energy of product states in a translation invariant spin chain can replace any quantum computation, 2007. arXiv:0710.1615v2 [quant-ph].
- [Kay08] Alastair Kay. The computational power of symmetric hamiltonians. *Phys. Rev. A*, 78:012346, 2008. arXiv:0801.3228v3 [quant-ph].
- [KKR06] J. Kempe, A. Kitaev, and O. Regev. The complexity of the Local Hamiltonian problem. *SIAM Journal of Computing*, 35(5):1070–1097, 2006.
- [KSV02] A.Y. Kitaev, A.H. Shen, and M.N. Vyalıy. *Classical and Quantum Computation*. AMS, Providence, RI, 2002.
- [Nag08] Daniel Nagaj. *Local Hamiltonians in Quantum Computation*. PhD thesis, MIT, 2008. arXiv:0808.2117v1 [quant-ph].
- [NW08] Daniel Nagaj and Pawel Wocjan. Hamiltonian quantum cellular automata in 1d, 2008. arXiv:0802.0886v4 [quant-ph].
- [Osb06] T. Osborne. Efficient approximation of the dynamics of one-dimensional quantum spin systems. *Phys. Rev. Lett.*, 97:157202, 2006.
- [Osb07] T. Osborne. Ground state of a class of noncritical one-dimensional quantum spin systems can be approximated efficiently. *Phys. Rev. A*, 75:042306, 2007.
- [OT05] R. Oliveira and B. Terhal. The complexity of quantum spin systems on a two-dimensional square lattice, 2005. quant-ph/0504050.
- [Sch05] U. Schollwöck. The density-matrix renormalization group. *Rev. Mod. Phys.*, 77:259–316, 2005.
- [SFW06] D. J. Shepherd, T. Franz, and R. F. Werner. Universally programmable quantum cellular automaton. *Phys. Rev. Lett.*, 97:020502, 2006.
- [Suz76] Masuo Suzuki. Relationship between d-dimensional quantal spin systems and (d+1)-dimensional ising systems. *Progress of Theoretical Physics*, 56(5):1454–1469, 1976.
- [vEB90] Peter van Emde Boas. *Handbook of Theoretical Computer Science*, volume A, chapter 1, pages 1–66. MIT Press, Cambridge, MA, USA, 1990.
- [VV86] L. G. Valiant and V. V. Vazirani. NP is as easy as detecting unique solutions. *Theor. Comput. Sci.*, 47(3):85–93, 1986.
- [Wat95] J. Watrous. On one-dimensional quantum cellular automata. In *Proc. 36th Annual IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 528–537, 1995.
- [Whi92] S. R. White. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.*, 69:2863, 1992.
- [Whi93] S. R. White. Density-matrix algorithms for quantum renormalization groups. *Phys. Rev. B*, 48:10345, 1993.