



The Complexity of Translationally Invariant Spin Chains with Low Local Dimension

Johannes Bausch, Toby Cubitt and Maris Ozols

Abstract. We prove that estimating the ground state energy of a translationally invariant, nearest-neighbour Hamiltonian on a 1D spin chain is QMA_{EXP} -complete, even for systems of low local dimension (≈ 40). This is an improvement over the best previously known result by several orders of magnitude, and it shows that spin-glass-like frustration can occur in translationally invariant quantum systems with a local dimension comparable to the smallest-known non-translationally invariant systems with similar behaviour. While previous constructions of such systems rely on standard models of quantum computation, we construct a new model that is particularly well-suited for encoding quantum computation into the ground state of a translationally invariant system. This allows us to shift the proof burden from optimizing the Hamiltonian encoding a standard computational model, to proving universality of a simple model. Previous techniques for encoding quantum computation into the ground state of a local Hamiltonian allow only a linear sequence of gates, hence only a linear (or nearly linear) path in the graph of all computational states. We extend these techniques by allowing significantly more general paths, including branching and cycles, thus enabling a highly efficient encoding of our computational model. However, this requires more sophisticated techniques for analysing the spectrum of the resulting Hamiltonian. To address this, we introduce a framework of graphs with unitary edge labels. After relating our Hamiltonian to the Laplacian of such a unitary labelled graph, we analyse its spectrum by combining matrix analysis and spectral graph theory techniques.

Contents

1. Background and Motivation	3451
2. Extended Introduction and Overview of Results	3452
2.1. Historical Context	3452

2.2.	Main Result	3455
2.3.	Proof Ideas and Techniques	3457
2.4.	Structure of the Paper	3464
3.	Preliminaries	3464
3.1.	Reversible Turing Machines	3464
3.2.	Quantum Complexity Classes	3466
3.3.	Geometrically k -Local Hamiltonians	3468
3.4.	QMA versus QMA _{EXP}	3469
3.5.	Laplacian Matrix and Algebraic Connectivity of Graphs	3471
3.6.	Kitaev's Geometrical Lemma for Graphs	3472
4.	Quantum Ring Machine	3474
4.1.	Definition	3474
4.2.	Universality	3476
5.	Unitary Labelled Graphs	3480
5.1.	Definitions	3480
5.2.	Semi-classical Unitary Labelled Graphs	3482
5.3.	Kitaev's Geometrical Lemma for Unitary Labelled Graphs	3485
6.	Quantum Thue Systems	3486
6.1.	Thue Systems	3486
6.2.	Quantum Thue Systems and Their Hamiltonian	3488
6.3.	Quantum Thue Systems as a Computational Model	3489
7.	Hardness Results	3491
7.1.	A Special Kind of Quantum Thue System	3491
7.2.	Hardness Result	3491
8.	Turing's Wheelbarrow	3494
8.1.	Introduction	3494
8.2.	Notation	3497
8.3.	Transition Rules and History State	3498
8.3.1.	Initialization	3498
8.3.2.	Ghost	3499
8.3.3.	Base-6 Counter	3499
8.3.4.	Unary Counter	3501
8.3.5.	Computation	3501
8.4.	Branching in the History State	3503
8.5.	Simplicity of Turing's Wheelbarrow	3507
8.6.	Special Properties	3508
8.7.	Final Dimension Reduction	3508
9.	Conclusion	3509
	Acknowledgements	3510
	Open Access	3510
	References	3511

1. Background and Motivation

Complex physical behaviour can emerge from even very simple rules. Yet if the system is *too* simple, one can often rule out the possibility of any exotic behaviour. Just how simple can a system be to nonetheless feature complex properties? Much of the progress in Hamiltonian complexity and related areas over the last decade can be viewed as improving our understanding of where this boundary between simple and complex lies.

For example, consider 1D spin chains with translationally invariant nearest neighbour interactions. Hastings proved that if the Hamiltonian describing the system is gapped, the ground state entanglement has to follow an area law [1]. In 1D, the area law means that the entanglement entropy between any contiguous region and its complement is upper-bounded by a constant, independent of the size of the region. It was believed that even for non-gapped Hamiltonians, area-law violations would contribute at most log corrections in the system size. Such long-range correlations in a spin chain's ground state which scale with the system's size are a common indicator of criticality, i.e. they show that the system is close to a quantum phase transition. The entanglement entropy is then expected to scale logarithmically with the number of spins, since critical spin chains can often be related to a conformal field theory.

However, using Hamiltonian complexity techniques, Irani [2] constructed an example of a spin chain in 1D that exhibits violation of the area-law *beyond* logarithmic corrections, indicating that one cannot describe such behaviour by a conformal field theory. Irani's construction breaks translational invariance, so it cannot directly be compared to systems satisfying area laws. A later construction [3] can give a similar area-law violation while preserving translational invariance. However, the required local dimension, $O(10^6)$, is vast. It is therefore at best questionable whether this area-law violation could ever be observed in practice. Does this mean that such violations only occur for some peculiar theoretical models with non-translationally invariant couplings, or unrealistically large Hilbert space dimensions?

We now know that the answer to this question is negative. First, it was shown by [4] that, even for frustration-free spin-1 chains (i.e. local dimension 3), one can construct interactions that yield highly entangled ground states, indicating critical behaviour. In fact, this result delineates a strict dimension threshold for the presence ground state entanglement in frustration-free systems. For frustration-free spin-1/2 chains (i.e. local dimension 2) with translationally invariant nearest-neighbour interactions, it was already known that ground states are unentangled [5]. Building on this, Movassagh et al. [6] constructed models which give power-law violation of the area-law for translationally invariant spin-5/2 chains (i.e. local dimension¹ 6), significantly improving on the bound on the local dimension threshold for power-law area-law violation from Gottesman and Irani's result.

¹ [6] in fact prove their result for local dimension 5 but breaking strict translational invariance by adding boundary terms at the ends of the chain. Using a trick due to [3], the boundary terms can be removed at the cost of increasing the local dimension by 1.

Similar dimension-related physicality questions also surround Cubitt et al.'s result which proves that deciding whether a system is gapped or gapless in the thermodynamic limit is an undecidable problem, even for 2D spin lattices with translationally invariant local interactions [7]. Again, the local Hilbert space dimension in the model they describe is vast. Bravyi and Gosset recently derived necessary and sufficient conditions for a gapped or gapless phase for frustration-free spin-1/2 chains [8]. So at the other end of the local dimension scale, the spectral gap problem is decidable in some cases. However, there is evidence that an astronomical local dimension may not be a fundamental ingredient in the emergent behaviour that gives rise to undecidability of the spectral gap. The abrupt change in the spectrum at very large system sizes that is behind the undecidability, can also occur on 2D lattices of far lower-dimensional spins [9]. Again, this poses an immediate question of whether there is some local dimension threshold above which undecidability can occur, but below which it cannot.

The original and most widely studied question in Hamiltonian complexity theory, however, is that of estimating the ground state energy of a local Hamiltonian. Kitaev showed that this problem is QMA-hard [10] (i.e. at least as hard as every other problem in the complexity class QMA—the quantum generalization of NP). Similar to a spin glass, when cooled down these QMA-hard systems are predicted to get stuck in one of their many metastable configurations, and will take exponentially long (in the system size) to find their global minimum energy configuration. QMA-hardness-inspired constructions lie behind all the results mentioned above. Yet even though the parameters describing QMA hard ground state Hamiltonians have been improved successively [3, 11–14], a lower local dimension threshold below which systems cannot feature spin-glass-like frustration is not known; for non-translationally invariant systems we know that this bound can be at most 8. For the more physically relevant case of spin chains *with* translational symmetry, however, the best-known bound is $O(10^6)$, due to Gottesman and Irani [3], which is unphysically large. From a physical perspective, it makes a dramatic difference if the complexity threshold is e.g. 7, or 1000.

In this work, we improve the best-known upper bound on the local Hilbert space dimension required for QMA-hardness in translationally invariant spin chains by several orders of magnitude, showing that the question of estimating the ground state energy of a local translationally invariant Hamiltonian with nearest-neighbour interactions remains hard, even for spins on a chain with local dimension ≈ 40 .

2. Extended Introduction and Overview of Results

2.1. Historical Context

Hamiltonians are the one-stop shop for describing physical properties of multi-body quantum systems, and are of paramount interest for an array of disciplines ranging from experimental condensed matter physics to theoretical

computer science [3, 7, 11–13, 15–18]. While computer scientists are interested in the computational power of different models, for physicists it is important to calculate the structure of the low-energy spectrum of quantum systems, in particular to approximate the minimum energy of the system, i.e. the ground state energy.

The decision problem of determining whether such a local Hamiltonian operator has lowest energy—or eigenvalue—below some α or above some β , with $\beta > \alpha$, can be thought of as the quantum analogue of the maximum satisfiability problem MAX-SAT. Similar to the well-known 3-SAT, this asks for the maximum number of clauses of a Boolean formula in conjunctive normal form that can be satisfied simultaneously. In the quantum case, each local term \mathbf{h} of \mathbf{H} is analogous to a clause, while a global state $|\psi\rangle$ is analogous to a global variable assignment, and the smaller $\langle\psi|\mathbf{h}|\psi\rangle$ is, the closer $|\psi\rangle$ is to satisfying the corresponding clause \mathbf{h} . The LOCAL HAMILTONIAN problem formalizes the notion of maximizing the number of local terms of \mathbf{H} which can be simultaneously minimized by some global state $|\psi\rangle$, in the sense that $\langle\psi|\mathbf{H}|\psi\rangle$ is small. Physically, this minimum is equal to the lowest energy of the system.

Formally, we can state the LOCAL HAMILTONIAN problem as the following promise problem.

Definition 1 (*k-Local Hamiltonian*).

Input An integer n and a k -local Hamiltonian \mathbf{H} on a multipartite Hilbert space $(\mathbb{C}^d)^{\otimes n}$, and two real numbers $\beta > \alpha$ such that $\beta - \alpha \geq 1/p(n)$, for some fixed polynomial $p(n)$. The smallest eigenvalue λ_{\min} of \mathbf{H} is promised to be either smaller than α or greater than β .

Question Is $\lambda_{\min} < \alpha$, or $\lambda_{\min} > \beta$?

The k -LOCAL HAMILTONIAN problem has a track record of long-standing interest (cf. Table 1). The foundations were laid with Feynman’s paper

TABLE 1. Brief historic overview of QMA (QMA_{EXP} for [3]) completeness results in Hamiltonian complexity

	Locality	Local dimension	Geometry and symmetries
Kitaev (1999, see [39])	5	2	Arbitrary
Kempe et al. [11]	2	2	Arbitrary
Oliveira and Terhal [12]	2	2	2D, planar, nearest-neighbour interactions
Aharonov et al. [13]	2	12	Line, nearest-neighbour
Hallgren et al. [14]	2	8	Line, nearest-neighbour
Gottesman and Irani [3]	2	Huge ($\approx 10^6$)	Line, nearest-neighbour, translationally invariant

[19] on encoding quantum circuits into the ground state of a Hamiltonian, which motivated a whole series of interesting and increasingly sophisticated results showing that variants of this problem are QMA- or QMA_{EXP}-complete.²

On the other hand, just as in classical computer science 2-SAT is solvable in polynomial time, its quantum analogue—the QUANTUM 2-SAT, a special case of the 2-LOCAL HAMILTONIAN problem³—can also be solved deterministically in polynomial time: [20] proved an $O(n^4)$ runtime bound, and later a linear-time algorithm was discovered independently by [21] and [22]. Yet the resemblance with classical results goes further: QUANTUM 4-SAT and later QUANTUM 3-SAT were shown to be QMA₁-complete [20, 23]. In the same spirit, a recent result shows that in case of one-dimensional *gapped* local Hamiltonians, there exists an efficient randomized algorithm for approximating the ground state as a matrix product state [24] (this result is independent of the local dimension).

However, the LOCAL HAMILTONIAN problem, as defined in Definition 1, allows the Hamiltonian to be frustrated (going beyond local projectors), and encompasses Hamiltonians whose gap closes inverse polynomially in the system’s size. It is thus a natural question to ask whether this more general LOCAL HAMILTONIAN problem remains computationally hard, even under restrictions motivated on *physical* grounds (e.g. for translationally invariant interactions and for qubits), or whether there is a fundamental local dimension threshold below which it becomes tractable.

To motivate this further, it is crucial to note that Hamiltonian constructions in the spirit of [19] are a proof-of-concept and may not necessarily be natural, in the sense that we would not encounter them in nature describing an actual physical system. There are three fundamental criteria for judging the “physicality” of a Hamiltonian: the interactions should be geometrically local, the dimension of the interacting subsystems should be small, and the interactions should exhibit translational invariance. These properties apply to physical systems we typically encounter in nature. For example, translational invariance means that if the Hamiltonian is specified on a lattice, the interactions are the same independently of the location within the lattice.

Starting with Kitaev’s original proof of QMA-completeness of 5-LOCAL HAMILTONIAN [10], the locality and local dimension of the constructions were improved successively [11–13], cf. Table 1. For spins of local dimension 8 coupled by nearest-neighbour interactions on a chain, QMA-hardness was proven

² QMA_{EXP} is to QMA what NEXP is to NP. This is a necessary technicality whenever the input has to be specified in unary. The energy gap still scales inverse polynomially with system size n , and the physical implications are exactly the same as for QMA-completeness. We define these complexity classes rigorously in Sect. 3.2 and explain their difference in detail in Sect. 3.4.

³ More specifically, QUANTUM 2-SAT asks whether a sum of 2-local terms, where each term is a 2-qubit projector that acts on any pair of qubits, is frustration-free, i.e. has a 0-energy eigenstate or, equivalently, a state that simultaneously satisfies all local constraints.

by Hallgren et al. [14]. All of these results make heavy use of the non-translationally invariant nature of interactions, which vastly simplify the encoding of the problem instance and verifier circuit into the local structure of the Hamiltonian. The QMA_{EXP} -hardness result by Gottesman and Irani [3], which features a 2-local Hamiltonian on a line with translationally invariant nearest-neighbour interactions, shows that having translational symmetry does not change the complexity class of the local Hamiltonian problem. But one caveat remains: the local dimension is unphysically large, on the order of 10^6 .

2.2. Main Result

Our goal is to significantly improve on this best-known upper bound on the local dimension. We develop a set of new methods to prove that the complexity threshold above which the LOCAL HAMILTONIAN problem is computationally hard is at most 42, even under the strict physicality constraints outlined above. More precisely, we prove the following main theorem.

Theorem 2. *The LOCAL HAMILTONIAN problem with translationally invariant interactions between neighbouring spins on a chain with local dimension 42 is QMA_{EXP} -complete. This holds true even for Hamiltonians with local terms of the form $\mathbf{h} + p(n)\mathbf{b}$, where \mathbf{h} and \mathbf{b} are fixed 2-local interactions and $p(n)$ is a fixed polynomial in the chain length n .*

Following the notation in [3], we label this class of problems 2-TILH , for translationally invariant 2-local Hamiltonian. Analogous to all past hardness constructions, we prove our result by explicitly defining a family of QMA_{EXP} -hard instances of 2-TILH . More precisely, the instances we construct are so-called *history state* Hamiltonians: by choosing the local constraints in \mathbf{H} suitably, one can create a Hermitian operator with a ground state spanned by states that are a uniform superposition over the history of a computation, such that the state at step t is *entangled* with a corresponding state $|t\rangle$ in a separate time register (i.e. $\sum_t |t\rangle \otimes |\psi_t\rangle$). Measuring the time register at time t then yields the state of the computation at this step. This “program counter”, as Feynman describes it, can be thought of as a clock or a finite automaton driving the application of quantum gates. Originally, only linearly evolving clock constructions were used, since analysing the spectrum of a Hamiltonian with branching computational paths is more difficult. More recently, QMA -hardness constructions in 1D and 2D have used limited branching and cycles [14, 25]. These have also been exploited in the slightly different context of adiabatic and Hamiltonians quantum computation [26–28].

Whereas recent results [17] make use of perturbation gadgets—approximating higher-order interactions in the low-energy subspace of the system by an effective high-energy theory—it is known that this does not work in one-dimensional systems [13]. The improvements in [14] over [13] are possible, however, by approximating 4-local interactions by a sum of 2-local interactions, effectively introducing illegal transitions that have to be penalized. Perturbation gadgets and locality reduction both depend on introducing

a large energy scale to project out illegal subspaces. Our results, on the other hand, do *not* use perturbation theory.⁴

Our findings are based on the following three main technical contributions:

1. All previous constructions encode one of the standard models of quantum computation (almost always the circuit model, with the exception of [3] which encodes a quantum Turing machine), which are not optimized for this task. We design a new universal model of quantum computation—a *quantum ring machine* (QRM)—which we prove to be quantum Turing-complete. The periodicity of the QRM’s computational steps make it particularly well-suited for local Hamiltonian constructions.
2. We next introduce *unitary labelled graphs* and their associated Hamiltonians, which can accommodate a non-deterministic clock construction to drive quantum computation. This generalizes Feynman’s original clock construction [19], which corresponds to a path graph in our setup. Mirroring Kitaev’s analysis [10], our Hamiltonian is also equivalent to a Laplacian of the corresponding graph, which allows us to analyse its spectrum using a combination of spectral graph theory and matrix analysis techniques. These techniques let us analyse ground states of more complicated Hamiltonians than previously possible.
3. We define yet another computational model—a *quantum Thue system*, or a quantum string rewriting system—that on the one hand is particularly well-suited for embedding a computational model into local interactions of a Hamiltonian; and on the other hand, under simple *local* constraints on the rewriting rules, necessarily produces Hamiltonians that correspond to unitary labelled graphs. Quantum Thue systems can in a sense be thought of as an assembly language for compiling computational models into local translationally invariant Hamiltonians, which could also be used for adiabatic quantum computation, or Hamiltonian quantum computers (cf. [18, 26]).

In light of our result being rather involved and technical, we want to give a poor man’s overview of our findings, which—without any proofs—outline the technical contributions in this paper. We want to emphasize that we made an effort to keep each section largely self-contained; in particular the section on spectral analysis of graphs with unitary edge labels, quantum ring machines, and quantum Thue systems can be regarded independently of each other. The QMA_{EXP} -hardness proof in Sect. 7 of course utilizes all of our developed machinery, but in such a way that the proof of existence of QMA_{EXP} -hard instances themselves are given a separate section.

Since the latter part is somewhat technical and specific, we want to point out that one does not need to understand the construction itself to follow the

⁴ The polynomial in Theorem 2 is an artefact of the construction. A standard trick from [3] can reduce the Hamiltonian to fixed $O(1)$ interactions by slightly increasing the local dimension, see Remark 73.

idea behind the hardness proof, which hopefully facilitates an understanding of the result.⁵

2.3. Proof Ideas and Techniques

Spectral Analysis for Hamiltonians Encoding Non-Deterministic Computation.

As briefly explained in the introduction, the fundamental idea behind encoding quantum computation into the ground state of a Hamiltonian is based on the concept of *history states*, introduced by Feynman in 1986. For some quantum circuit represented by local gates $\mathbf{U}_1, \dots, \mathbf{U}_T$ on a Hilbert space \mathcal{H} , we define a Hamiltonian on the product space $\mathbb{C}^T \otimes \mathcal{H}$ as

$$\mathbf{H} := \sum_{t=1}^{T-1} (|t\rangle\langle t| \otimes \mathbb{1} + |t+1\rangle\langle t+1| \otimes \mathbb{1} - |t+1\rangle\langle t| \otimes \mathbf{U}_t - |t\rangle\langle t+1| \otimes \mathbf{U}_t^\dagger). \quad (1)$$

The ground state of this Hermitian operator is spanned by states of the form $\sum_t |t\rangle \otimes |\psi_t\rangle$, where $|\psi_t\rangle = \mathbf{U}_t \cdots \mathbf{U}_1 |\phi\rangle$ for some $|\phi\rangle \in \mathcal{H}$. For any $|\phi\rangle$, $\ker \mathbf{H}$ thus encodes the uniform superposition over the history of the quantum circuit acting on $|\phi\rangle$. An intuitive way of thinking about these ground states is that they represent quantum computation driven by a clock, i.e. for each increment of the clock register, the corresponding quantum gate is applied to the computational register.⁶

Essentially all past result employ such history state Hamiltonians with a linear clock, i.e. for every computational step, there exists precisely one unique forward and backward transition. For local Hamiltonian constructions—i.e. where \mathbf{H} is a sum of local terms—this implies that each *local* rule has to know the exact location within the overall computation.

To be more specific, consider a spin chain of length n as the Hilbert space $\mathcal{H}_{\text{loc}}^{\otimes n}$. The interactions on this chain then take the form of a set of local *rewriting rules* acting on neighbouring sections of spins: for $|\psi_i\rangle, |\phi_i\rangle \in \mathcal{H}_{\text{loc}}^{\otimes k}$ for some constant $k < n$, we encode the evolution $|\psi_i\rangle \mapsto |\phi_i\rangle$ by a local Hamiltonian term $\mathbf{h}_i = (|\psi_i\rangle - |\phi_i\rangle)(\langle\psi_i| - \langle\phi_i|)$. The overall Hamiltonian is then a sum of these local interactions over all spins, i.e.

$$\mathbf{H} = \sum_j \mathbb{1}_{1, \dots, j-1} \otimes \left(\sum_i \mathbf{h}_i \right)_{j, \dots, j+k-1} \otimes \mathbb{1}_{j+k, \dots, n}. \quad (2)$$

If the global evolution defined by the terms \mathbf{h}_i is unique, this implies that it is always possible to locally determine the global state of the computation.⁷ This means that locally, we have to store this state in one way or another: under this

⁵ The reason behind unhitching the explicit construction of QMA_{EXP}-hard instances in this way is to allow for further optimization of the local dimension to go through without having to reprove all of the claims; in fact, we encourage the interested reader to have a stab at finding a quantum Thue systems following the four properties given in Lemma 59, but with an alphabet that is smaller than ours.

⁶ The notion of time in this context is meaningless, but simplifies an intuitive understanding on how computation is embedded into the ground state of \mathbf{H} .

⁷ Hamiltonians such as in Eq. (2) are combined with a series of local projectors which single out a computationally valid ground state, so strictly speaking the local rules will only have

requirement it is difficult to push the limits of local Hilbert space dimension down, and much could be gained if we could e.g. allow the local computational state to be ambiguous to some extent (but such that if the wrong transition is applied, the computation does not proceed to tamper with the actual outcome of the embedded circuit).

In our work, we go beyond linear clock constructions, and prove a series of spectral graph-theoretic results which allow us to analyse more complicated history state Hamiltonians. We outline these novel techniques below.

If all $|\psi_i\rangle$ and $|\phi_i\rangle$ in Eq. (2) are standard basis vectors, then each rule corresponds to an edge in a graph G with vertices labelled by the canonical basis of the spin chain. \mathbf{H} thus equals the Laplacian of the graph G (whose spectrum is accessible) and the ground state of \mathbf{H} is given by the uniform superposition over connected graph components of G . We call ground states of such Hamiltonians as in Eq. (2) *history states*, since they encode the closure of states reachable under the given rewriting rules.

To analyse the spectrum of more general non-basis transitions $|\psi_i\rangle \mapsto |\phi_i\rangle$, one needs to prove that this choice still allows \mathbf{H} to be at least unitarily equivalent to a graph Laplacian Δ , e.g. by explicitly constructing a unitary similarity transform \mathbf{W} such that $\mathbf{W}^\dagger \mathbf{H} \mathbf{W} = \Delta \otimes \mathbf{1}$. Most if not all QMA-hard construction since Kitaev's go along this route; however, in the language of graphs, the unitary equivalence could only be proven if Δ is the Laplacian of a path graph.⁸ Just as in Eq. (2), this graph essentially corresponds to the finite state automaton “driving” the computation; if it is a path graph, the computational path is limited to a sequential application of transition rules $|\psi_i\rangle \mapsto |\phi_i\rangle$ or gate applications encoded therein.

We extend this notion to allow much more complicated branching in the computational path to occur. In particular, we prove a series of results which guarantee the existence of the partially diagonalizing unitary \mathbf{W} solely based on properties of the rewriting rules, without the need to explicitly analyse the overall evolution of the system. This has two major benefits: it allows more powerful state transitions which are not necessarily unique for every step, and it drastically simplifies the spectral analysis of \mathbf{H} for whichever construction we choose to work with, as we do not need to construct the equivalence between \mathbf{H} and Δ explicitly. As an important example, our model is the first to allow multiple threads of computation to run in parallel, which then join at some common state.

Footnote 7 continued

to discriminate the current computational state locally within this valid subspace—cf. [14], where this is exploited to break down 4-local interactions to 2-local ones.

⁸ A more complex construction with a local clock was considered in [25], where the authors consider a 2D surface and allow executing transitions in parallel, as long as the execution front behaves in a time-like fashion. To analyse the spectrum of the resulting Hamiltonian, they relate the propagation terms to the diffusion of a string on a torus, corresponding to a ferromagnetic Heisenberg model with partially twisted periodic boundary conditions. Their analysis, while elegant, is specific to their string diffusion-type execution order of quantum gates. These Hamiltonians cannot generally be translationally invariant, as the circuit must be laid out on the 2D surface.




FIGURE 1. Example of a unitary labelled graph (ULG) with vertices $\{1, 2, 3, 4, 5\}$ and three non-trivial unitaries $\mathbf{X}, \mathbf{Y}, \mathbf{Z} \in \mathbf{U}(\mathcal{H})$. The associated Hamiltonian for this ULG, as defined in Eq. (3), is unitarily equivalent to the Laplacian of the underlying graph if $\mathbf{ZYX} = \mathbf{1}$. We provide an explicit description of this change-of-basis unitary

In a bottom-up approach, we formalize the notion of a Hamiltonian associated with a graph. Starting from a simple directed graph $G = (V, E)$ we associate a Hilbert space \mathcal{H} to each vertex $v \in V$, and a unitary $\mathbf{U}_{(a,b)} : \mathcal{H} \rightarrow \mathcal{H}$ for every directed edge $(a, b) \in E$. We call such a graph with Hilbert space and family of unitaries a *unitary labelled graph*, or ULG for short. As an example, consider Fig. 1.

The *associated Hamiltonian* for the ULG is then defined as

$$\mathbf{H}(G) := \sum_{(a,b) \in E} \sum_i (|a\rangle \otimes |i\rangle - |b\rangle \otimes \mathbf{U}_{(a,b)} |i\rangle) (\text{ herm. conj. }), \quad (3)$$

where the $|i\rangle$ label a basis of \mathcal{H} . Observe that this construction is more general than a local Hamiltonian on a spin chain as in Eq. 2: $\mathbf{H}(G)$ is simply regarded as Hermitian operator on the overall Hilbert space $\mathbb{C}^V \otimes \mathcal{H}$ where the vertex labels are arbitrary, and not necessarily make $\mathbf{H}(G)$ local in any sense.

The associated Hamiltonian $\mathbf{H}(G)$ bears some structural resemblance with a graph Laplacian, as already mentioned. We prove the following theorem.

Theorem 3. *If the product of unitaries along any loop in the graph G is $\mathbf{1}$, a property we call simple, then $\mathbf{H}(G)$ is unitarily equivalent to $\Delta \otimes \mathbf{1}_n$, where $n = \dim \mathcal{H}$ and Δ is the Laplacian of G .*

Figure 1 satisfies this theorem if and only if the product of unitaries in the loop are $\mathbf{ZYX} = \mathbf{1}$. We provide an explicit expression for this diagonalizing unitary, which can be constructed in poly-time using a breadth-first search algorithm along a spanning tree of G .

Quantum String Rewriting. In order to reintroduce locality to our Hamiltonian construction, we further develop a notation which facilitates embedding transition rules as in Eqs. (2) and (3) into the ground state of a local Hamiltonian. This notation is heavily motivated by string rewriting models, and we extend this notion to introduce a new quantum Turing-complete model based on transitions able to perform quantum gates on part of the string’s alphabet.

As mentioned, past hardness constructions (summarized in Table 1) encode computation in *local* transition rules that act on spins connected by some

underlying graph of interactions. While some of these transitions are classical—i.e. basis-preserving—others act on the spin states with a non-diagonal unitary operator, performing the actual quantum computation. Inspired by classical string rewriting systems, we interpret these quantum interactions as local quantum rewriting rules, and introduce a new abstract rewriting system called *quantum Thue system*. This extends an already-existing model of string rewriting—semi-Thue systems⁹—which are well-studied classically [29].

A (classical) semi-Thue system consists of a finite alphabet Σ and length-preserving replacement rules for strings over this alphabet. Similar to the word problem, computation can be encoded in the question whether there exists a connecting path between some input and output strings s_i and s_f . It is straightforward to simulate universal classical Turing machines with a Thue system, which shows that the latter is a Turing-complete model for classical computation. But what about quantum computation?

For quantum Thue systems, we require that the alphabet splits into a classical and a quantum part, i.e. $\Sigma = \Sigma_{cl} \sqcup \Sigma_q$. Transition rules can be purely classical—between elements of Σ_{cl}^* , quantum—between elements of Σ_q^* , or a mixture thereof, in which case we require that the rule preserves the number of quantum symbols $|s|_q$ of a string $s \in \Sigma^*$. In addition, every rule r acting on at least some quantum symbols has a unitary $\mathbf{U}_r \in \mathbf{U}(\mathcal{H}^{\otimes |s|_q})$ attached, where \mathcal{H} is some fixed, finite-dimensional Hilbert space.

Starting on some string s and a state vector $|v\rangle \in \mathcal{H}^{\otimes |s|_q}$, we apply any matching string rewriting rule $s \xrightarrow{r} s'$ in turn. For every replacement, we *also* apply the corresponding unitary to the state vector, i.e. $\mathbf{U}_r |v\rangle = |v'\rangle$. In this fashion, we can model quantum computation, driven by a finite automaton: if we make the underlying classical Thue system implement a Turing machine that writes out a quantum circuit description on the string, and then perform this quantum circuit on a separate set of qubits attached to some quantum symbols, the final state vector will contain the output of a quantum computation.

One can then show that a quantum Thue system is itself a special case of a unitary labelled graph, which allows us to translate it into a Hamiltonian. We show that the locality of the resulting Hamiltonian only depends on the range of the largest replacement rule, e.g. if one at most replaces a 3-character string, the resulting Hamiltonian will also be 3-local and translationally invariant.

As replacement rules are not necessarily unique, the computation will have potential ambiguities. As such, we regard all strings connected to the initial starting string s_i via some arbitrary combination of rules, and the size of this set corresponds to the number of basis states that the corresponding *history state* (the ground state of the associated Hamiltonian of the unitary labelled graph defined by the quantum Thue system) is comprised of.

⁹ Named after the Norwegian mathematician Axel Thue. We require all rule sets for quantum string rewriting to be symmetric; a symmetric semi-Thue is simply called Thue system, explaining the name quantum Thue system.

A Simpler Computational Model. The complexity class QMA_{EXP} is usually defined in terms of the circuit model, i.e. as a uniform family of verifier circuits: a promise problem $\Pi = (\Pi_{\text{YES}}, \Pi_{\text{NO}})$ is in QMA_{EXP} if there exists a classical Turing machine, such that the verifier circuit for a problem instance $l \in \Pi_{\text{YES}} \cup \Pi_{\text{NO}}$ can be written out by the Turing machine in $O(\exp |l|)$ steps where $|l|$ is the instance size. Being used as an all-purpose computational model, Turing machines have significant downsides: they have complicated transition functions, need a lot of internal states (which translates to an enormous local dimension when encoded in a Hamiltonian) and are rarely written out explicitly (so it is hard to get tight bounds on the required dimension). On the other hand, in past constructions, embedding a circuit directly required the use of non-local clock states marking the position within the circuit, or non-translationally invariant terms that encode the circuit unambiguously.

We introduce a new computational model which allows us to circumvent the direct use of complicated Turing machines or quantum circuits. The so-called *quantum ring machine* consists of a cyclic ring of *qudits* (i.e. d -dimensional quantum systems) and a unitary \mathbf{R} describing a head that acts on two qudits at a time. At each time-step, the head moves in the same direction along the ring and cyclically acts on adjacent cells. We give the following definition (see Sect. 4 for more details).

Definition 4 (*Quantum ring machine*). A *quantum ring machine* consists of a ring of n qudits, each of dimension d , and a unitary operator \mathbf{R} acting on a pair of qudits. The n -qudit ring is initialized in state $|\psi_{\text{in}}\rangle$ and the machine proceeds by applying \mathbf{R} cyclically to pairs of adjacent qudits along the ring—see Fig. 3—until one of the qudits indicates halting: its reduced density matrix has support completely inside a certain halting subspace $\mathcal{H}_{\text{halt}}$, while the reduced states of all qudits up to this point were orthogonal to $\mathcal{H}_{\text{halt}}$.

To show that a quantum ring machine is computationally equivalent to a uniform family of quantum circuits, we encode a classical Turing machine’s transition function into \mathbf{R} , where the internal states, including the Turing machine’s halting flag, are stored as a classical information on the ring. Such ring machine can be used to write out and execute a quantum circuit “on-the-go”: it is universal for whichever uniform circuit class is encompassed by its allowed runtime. Quantum ring machines thus bridge the gap between circuits, which are particularly simple to specify locally but have a complex global structure, and Turing machines, which are difficult to specify locally due to a possibly large number of internal states, but have a straightforward global evolution as the tape only changes in at most one location at each step. A schematic of the ring machine can be found in Figs. 2 and 3.

The ring machine’s simple mechanism allows its evolution to be described by a set of local quantum rewriting rules. These rules operate at a *physical* level while the ring machine operates at a *logical* level—each application of ring machine’s head \mathbf{R} on a pair of logical qubits is implemented by a sequence of physical operations acting on a much larger number of qubits. At any given time, the ring machine’s head is positioned on a specific pair of logical qubits,




FIGURE 2. Schematic of a quantum ring machine. A fixed unitary \mathbf{R} is cyclically applied to a ring of qudits until one of the qudits indicates a halting configuration




FIGURE 3. Ring machine’s evolution implementing a uniform quantum circuit. Double lines carry classical while single lines carry quantum information. Classical wires encode where the next quantum gate from a small universal set will be applied

and after each application of \mathbf{R} this location is updated in a similar fashion as Turing machine’s head—it is shifted either up or down along the ring by one position. Overall, \mathbf{R} is a large controlled unitary that acts at a given logical location only if the ring machine’s internal state—stored as a classical bit on the physical tape—is in an active configuration.

QMAEXP-hardness of 2-TILH.

The final proof of Theorem 2 is based on the following lemma.

Lemma 5. *There exists a BQEXP-universal quantum Thue system with 39 symbols, 3 of which are quantum, with attached Hilbert space \mathbb{C}^2 and 2-local rules.*

We prove this by writing out a quantum Thue system which executes a BQEXP-universal quantum ring machine. The quantum Thue system makes heavy use of the new possibilities of ambiguous replacement rules, which allow the history state path to branch. For the QMAEXP hardness proof itself, we combine this Thue system with a series of local penalty terms, which allow us to single out the history state as lowest energy ground state for any encoded YES instance.

Furthermore, we prove that the quantum Thue system has a simple history state in the sense of Theorem 3, which allows us to analyse the spectrum of the resulting Hamiltonian. More specifically, we prove a variant of *Kitaev’s*

geometrical lemma (cf. Lemmas 30 and 44) which facilitates the spectral analysis of Hamiltonians that are sums of a unitary labelled graph Hamiltonian and local projectors. This finally allows us to prove our main result, Theorem 2, that 2-LOCAL-HAMILTONIAN is QMA_{EXP} -hard, even for translationally invariant nearest-neighbour interactions between spins of local dimension 42.

For completeness, we also want to give a brief overview over the family of hard QTS instances that we construct, but—as mentioned before—the QMA_{EXP} -hardness proof does not depend on the precise workings of it; assuming that Lemma 5 can be proven, Theorem 2 stands independently.

Treating the Hilbert space of the 2-TILH problem as a physical tape of length n —some symbols quantum, some classical—we write a set of transition rules to perform the following steps that simulate the quantum ring machine.

1. As in the construction by Gottesman and Irani [3], we use a counter to translate the chain length n into a *program* string of length $O(\log n)$ on the left hand side of the chain, while on the right-hand side we store the physical data qubits, i.e. the ring of qubits our ring machine is executed on.
2. The program on the left hand side contains a physical-level description of a quantum circuit (over a small, finite, universal gate set) for implementing one step of the quantum ring machine, i.e. one application of the ring machine’s head \mathbf{R} . The program’s rightmost bit always indicates the next gate in the circuit, and this gate is always applied to the two leftmost data qubits on the physical ring (see Fig. 4).




FIGURE 4. Illustration of *Turing’s wheelbarrow* construction (see Sect. 8 for more details). It consists of a tape that stores a program string on the left- and data qubits on the right-hand side. Two types of actions are supported: application of a quantum gate (left figure) and rewinding of the tape (right figure). The rightmost program bit always indicates the next action. For example, 0 indicates that a unitary gate \mathbf{U}_0 should be applied to the two leftmost data qubits, and the ring of qubits should then be cyclically rotated one position to the left. On the right, the action of the special symbol \diamond is depicted: it signals the rightmost qubit to move back to the left end of the tape. After each action, the program string is cyclically rotated one position to the right

3. Using the two types of basic commands—“apply gate” and “rewind tape”—shown in Fig. 4, the quantum circuit implementing \mathbf{R} can be executed cyclically on the physical data qubits, some of which are initialized to ancillary $|0\rangle$'s to be used in the computation.
4. The computation runs until a certain internal classical counter (stored on the ring) terminates. In our construction, we explicitly encode transitions for the gates SWAP, TOFFOLI and a controlled quantum-universal unitary; since SWAP and TOFFOLI are also universal for classical computation, the classical control machinery in the ring machine's head \mathbf{R} (i.e. the Turing machine used to write out the quantum circuit) can be executed exactly (without error). This means that the computation will halt deterministically (as otherwise there could be some overlap with a non-halting state). The transition rules for applying a gate as in Fig. 4 then have another control gate which only proceeds if the data bit to its right is in a specific configuration, terminating the machine's execution otherwise.
5. The length of the chain is chosen so that the program encodes a quantum ring machine equivalent to a BQEXP verifier circuit. It discriminates between YES and NO instances of the corresponding QMA_{EXP} language depending on whether the ring machine accepts or rejects, and a special symbol in the program description allows us to locally penalize a wrong initialization of ancillas and a NO output of the computation.

Our construction is universal in the sense that it can be used to implement an arbitrary quantum computation without the need to increase the local dimension (in the same spirit as a universal Turing machine can implement any computation without the need to increase the number of internal states). Since we leave parts of the input unconstrained, we conclude from BQEXP-completeness of these instances that they can be used as a QMA_{EXP} verifier, finalizing our claims.

2.4. Structure of the Paper

We summarize several standard definitions in Sect. 3. In Sect. 4, we define the aforementioned quantum ring machine and show that it is indeed Turing-complete for quantum computation. Section 6 formalizes the notion of quantum replacement rules and introduces the model of quantum Thue systems. Section 8 contains a constructive proof of a universal quantum Thue system, and Sect. 7 combines everything into our main hardness result.

3. Preliminaries

3.1. Reversible Turing Machines

We give the following standard definition of a (non-deterministic) Turing machine (for more background on Turing machines, see chapter 8 of [30]).

Definition 6 (*Turing machine*). A *Turing machine*—or TM for short—is a triple (Q, Σ, δ) , where Q is a finite set of internal states containing a distinct initial and halting state q_0 and q_f , respectively, and Σ is a finite set of tape symbols containing a designated blank symbol 0. Let $D := \{\text{left}, \text{right}\}$ be the two possible movement directions of the TM’s head. Then each element of the transition set $\delta \subseteq Q \times \Sigma \times \Sigma \times D \times Q$ is a quintuple of the form (q, s, s', d, q') , which means that if the Turing machine reads a symbol s under its head while in state q , it overwrites the symbol by s' , moves the head in direction $d \in D$ and transitions to state q' . At the beginning of the computation, the TM’s initial state is q_0 and the tape is initialized to all 0s, except for a finite block of consecutive cells containing the input. The machine *halts* once its internal state is q_f , for which there is no forward transition.

As we aim to implement TMs using quantum mechanics, we need them to be deterministic and reversible. The following is based on definition 10 from [31].

Definition 7 (*Deterministic and reversible Turing machine*). Consider a Turing machine (Q, Σ, δ) , and let $(q_1, s_1, s'_1, d_1, q'_1)$ and $(q_2, s_2, s'_2, d_2, q'_2)$ be any two distinct quintuples in δ . This TM is

- *deterministic* if $(q_1 = q_2) \implies (s_1 \neq s_2)$,
- *reversible* if $(q'_1 = q'_2) \implies (s'_1 \neq s'_2) \wedge (d_1 = d_2)$.

The first condition of Definition 7 rules out the possibility that $q_1 = q_2$ and $s_1 = s_2$, meaning that the current TM’s state and tape symbol should unambiguously determine the rest of the transition. Similarly, the second condition rules out the possibility that $q'_1 = q'_2$ and $s'_1 = s'_2$, as well as the possibility that $q'_1 = q'_2$ and $d_1 \neq d_2$, meaning that the reverse transition also is uniquely determined by the current state and tape symbol, and that the direction of the TM’s head movement in reverse is uniquely determined by the current state.

For a deterministic TM, one can regard δ as a partial function, namely $\delta : Q \times \Sigma \rightarrow \Sigma \times D \times Q$, since all combinations of internal state q and tape symbol s have at most one forward transition. For a reversible TM, δ is injective since all combinations of internal state q' and tape symbol s' have at most one backwards transition (whenever such transition exists, it uniquely determines the head movement direction d backwards). In fact, according to Definition 7, each state of a reversible TM can be entered only from one direction (this property is referred to as *unidirectionality* in [32]). In other words, it is sufficient to know only the TM’s current state (as opposed to both the state and the tape symbol) to answer the question “From which direction did the TM’s head arrive?”.

Due to unidirectionality, it is often natural to restrict the range of δ to $Q \times \Sigma$. In fact, the transition function δ of a deterministic reversible Turing machine can be replaced by a permutation matrix on $Q \times \Sigma$ without affecting the TM’s behaviour. For our convenience, we state this observation more formally (see also cor. B.2 and thm. 4.2 in [32]).

Lemma 8. *For any deterministic reversible Turing machine (Q, Σ, δ) , the partial transition function δ can be replaced by a pair (\mathbf{T}_δ, d) , where \mathbf{T}_δ is a permutation matrix on $Q \times \Sigma$ and $d : Q \rightarrow D$ is a function that determines, for each internal state $q \in Q$, the direction from which the TM’s head arrived in q . If we update the TM’s internal state and the current tape symbol according to \mathbf{T}_δ , and then move the TM’s head in the direction opposite to $d(q')$, where q' is the updated state, the behaviour is identical to the original transition function δ .*

Proof. The function d is readily obtained because of unidirectionality. A blueprint of \mathbf{T}_δ is obtained by restricting the range of δ to $Q \times \Sigma$ and describing δ ’s action on the elements of this set by a binary matrix. Since the TM is deterministic and reversible, this matrix contains at most one entry 1 in each row and column, so it can be easily extended to a permutation matrix. \square

From now on, we will consider only deterministic Turing machines and implicitly assume that they are reversible—this is justified by the following result due to Bennett [33] (see [31, 34] for more background on reversible computation).

Theorem 9 (Bennett [33]). *Any deterministic TM can be made reversible with at most polynomial overhead in terms of space and time.*

3.2. Quantum Complexity Classes

In this section, we formally define the quantum complexity classes BQP, BQEXP, QMA and QMA_{EXP} in terms of the circuit model, and refer reader to [32, 35, 36] for more details on quantum computational complexity.

In what follows, we fix some finite universal set of 2-qubit quantum gates, such as {HADAMARD, CNOT, $R(\pi/4)$ }—see [37, ch. 4.5]. We first define a uniform family of quantum circuits over this gate set.

Definition 10 (*Uniform family of quantum circuits*). Let $f : \mathbb{N} \rightarrow \mathbb{N}$ be a function and $(C_n)_{n \in \mathbb{N}}$ be a family of quantum circuits where each C_n

- acts on n qubits and has a distinct output qubit,
- requires at most $f(n)$ additional ancilla qubits initialized in $|0\rangle$,
- contains at most $f(n)$ gates from our universal set.

We say that $(C_n)_{n \in \mathbb{N}}$ is $f(n)$ -uniform if there exists a TM that on input 1^n produces an explicit description of C_n in less than $f(n)$ steps.

Let Σ be a finite set (*alphabet*), and let Σ^n and $\Sigma^* := \bigcup_{n \geq 0} \Sigma^n$ denote the sets of all length- n and all finite-length strings over Σ , respectively. A *promise problem* over alphabet Σ is a pair $\Pi = (\Pi_{\text{YES}}, \Pi_{\text{NO}})$ such that $\Pi_{\text{YES}} \cap \Pi_{\text{NO}} = \emptyset$, where $\Pi_{\text{YES}}, \Pi_{\text{NO}} \subseteq \Sigma^*$ are the sets of input strings corresponding to YES and NO instances, respectively. We will sometimes write $l \in \Pi$ meaning that $l \in \Pi_{\text{YES}} \cup \Pi_{\text{NO}}$.

Definition 11 (*Complexity class $BQ(f)$*). A promise problem $\Pi = (\Pi_{\text{YES}}, \Pi_{\text{NO}})$ is in $BQ(f)$, *bounded-error quantum $f(n)$ -time*, if there exists an $f(n)$ -uniform family of quantum circuits $(C_n)_{n \in \mathbb{N}}$ such that

$$\Pr(C_n(s) = \text{YES}) \geq \frac{2}{3} \quad \text{for } s \in \Pi_{\text{YES}} \quad \text{and}$$

$$\Pr(C_n(s) = \text{YES}) \leq \frac{1}{3} \quad \text{for } s \in \Pi_{\text{NO}},$$

where $C_n(s)$ denotes the random variable obtained by executing C_n on input $s \in \Pi$ of size $|s| = n$ and measuring the output qubit (the encoding of s as well as the measurement are performed in the computational basis).

We introduced the notation $BQ(f)$ to emphasize the fact that the definitions of classes BQP and BQEXP are essentially the same up to the bounding function:

$$\text{BQP} := \bigcup_{k \in \mathbb{N}} \text{BQ}(n^k) \quad \text{and} \quad \text{BQEXP} := \bigcup_{k \in \mathbb{N}} \text{BQ}(\exp(n^k)).$$

Trivially, $\text{BQP} \subseteq \text{BQEXP}$ since a longer runtime can only help.

It is well-known (see [35, Prop. 3]) that for BQP the probabilities of $2/3$ and $1/3$ in Definition 11 can be exponentially amplified while still remaining in the same complexity class. The same argument works for BQEXP as well, since we only need a polynomial number of repetitions to achieve the desired amplification.

Fact 12 (*Error-reduction for BQP and BQEXP*). *For any polynomial p , we can assume that $\Pr(C_n(s) = \text{YES}) \geq 1 - 2^{-p(n)}$ for $s \in \Pi_{\text{YES}}$ and $\Pr(C_n(s) = \text{YES}) \leq 2^{-p(n)}$ for $s \in \Pi_{\text{NO}}$ in the definitions of BQP and BQEXP.*

Intuitively, $\text{QMA}(f)$ is the class of promise problems for which the YES/NO answers can be verified by a $\text{BQ}(f)$ verifier.

Definition 13 (*Complexity class $\text{QMA}(f)$*). A promise problem $\Pi = (\Pi_{\text{YES}}, \Pi_{\text{NO}})$ is in $\text{QMA}(f)$, *$f(n)$ -time quantum Merlin-Arthur*, if there exists an $f(n)$ -uniform family of verifier quantum circuits¹⁰ $(C_n)_{n \in \mathbb{N}}$ such that

- if $s \in \Pi_{\text{YES}}$, \exists a witness state ρ on at most $f(n)$ qubits such that $C_n(s, \rho) = \text{YES}$ with probability at least $2/3$. This condition is known as *completeness*.
- if $s \in \Pi_{\text{NO}}$, \forall witness states ρ on at most $f(n)$ qubits $C_n(s, \rho) = \text{YES}$ with probability at most $1/3$. This condition is called *soundness*.

Observe that the witness size is implicitly constrained by the size of the quantum circuit family, cf. Definition 10, e.g. for BQP verifiers the witness is poly-sized while for BQEXP verifiers it can be exp-sized. As before, we define

$$\text{QMA} := \bigcup_{k \in \mathbb{N}} \text{QMA}(n^k) \quad \text{and} \quad \text{QMA}_{\text{EXP}} := \bigcup_{k \in \mathbb{N}} \text{QMA}(\exp(n^k)).$$

¹⁰ Here we use a slight variation of Definition 10: we also allow for at most $f(n)$ extra input qubits to store the witness state ρ (this is in addition to the n original input qubits and $f(n)$ ancillary qubits that are initialized in $|0\rangle$).

In particular, note that $\text{QMA} \subseteq \text{QMA}_{\text{EXP}}$ since a QMA verifier can be easily promoted to a QMA_{EXP} verifier. Indeed, while a QMA_{EXP} verifier gets an exponential-size witness and can run for an exponential amount of time, it does not have to (it can instead discard all witness qubits, except for a polynomial number, and verify them in polynomial time).

Soundness and completeness probabilities for QMA and QMA_{EXP} can also be amplified: see theorem 10 in [35], section 3.2 of [36], or lemma 14.1 in [10] (these techniques were originally devised for QMA, but they can be easily adapted also for QMA_{EXP}).

3.3. Geometrically k -Local Hamiltonians

In this section, we introduce basic notions relating to local Hamiltonians and formally state the `TILH` problem that will play central role. For more background on Hamiltonian complexity, see [35, 38, 39].

Definition 14. An n -qudit Hamiltonian is a Hermitian operator $\mathbf{H} = \mathbf{H}^\dagger$ acting on a multipartite Hilbert space $(\mathbb{C}^d)^{\otimes n}$ consisting of n systems (qudits), each of local dimension d .

We will label the individual systems by elements of $S := \{1, \dots, n\}$. Whenever we talk of a subset of systems $A \subseteq S$, we mean an ordered tuple of distinct elements of S . If \mathbf{h} is a k -qudit Hamiltonian for some $k \leq n$ and $A \subseteq S$ is a subset of $|A| = k$ systems, we write \mathbf{h}_A to denote the n -qudit Hamiltonian that acts as \mathbf{h} on qudits A and trivially (i.e. as $\mathbb{1}$) on the remaining qudits $S \setminus A$. We also write $A + i \subseteq S$ to denote A shifted by $i \in \mathbb{N}$ positions.

Definition 15. Let \mathbf{H} be an n -qudit Hamiltonian. Then

- \mathbf{H} is k -local if $\mathbf{H} = \sum_i \mathbf{h}(i)_{A_i}$ with $|A_i| \leq k \ \forall i$;
- \mathbf{H} is k -local and 1D if each $A_i \subseteq \{1, \dots, k\} + t_i$ for some shift t_i ;
- \mathbf{H} is translationally invariant if $\mathbf{H} = \sum_i \mathbf{h}_{A+i}$ for some $A \subseteq S$ where \mathbf{h} is fixed.

In particular, \mathbf{H} is a 1D translationally invariant k -local Hamiltonian if $\mathbf{H} = \sum_i \mathbf{h}_{\{1, \dots, k\}+i}$ for some fixed k -qudit Hamiltonian \mathbf{h} .

Our central problem of interest is deciding the ground energy of 1D translationally invariant k -local Hamiltonians of local dimension d . For brevity, we will refer to this as the `TILH` problem.

Definition 16 ($((k, d)\text{-TILH})$). Let $\mathbf{H} = \sum_i \mathbf{h}_{\{1, \dots, k\}+i}$ be a 1D translationally invariant k -local Hamiltonian on a qudit chain of length n , where each qudit has local dimension d and \mathbf{h} is some fixed k -qudit Hamiltonian.

Input The chain length n and the matrix entries of \mathbf{h} , as well as two real numbers α and β , all up to $\text{poly log } n$ bits of precision.

Promise The operator norm of each local term is bounded, $\|\mathbf{h}\| \leq 1$, and either $\lambda_{\min}(\mathbf{H}) \leq \alpha$ or $\lambda_{\min}(\mathbf{H}) \geq \beta$, where $\lambda_{\min}(\mathbf{H})$ denotes the smallest eigenvalue of \mathbf{H} and $\beta - \alpha \geq 1/p(n)$ for some fixed polynomial $p(n)$.

Output YES if $\lambda_{\min}(\mathbf{H}) \leq \alpha$, else NO.

We emphasize that the input in Definition 16 is just the description of the k -local term \mathbf{h} and the chain length n , not the entire (exponentially sized) Hamiltonian \mathbf{H} . An equivalent variant of the definition relaxes the norm bound to $\|\mathbf{h}\| \leq \text{poly } n$ and gives a promise that either $\lambda_{\min}(\mathbf{H}) \leq \alpha$ or $\lambda_{\min}(\mathbf{H}) \geq \beta$ for some fixed constants $\beta > \alpha$. We can always rescale the overall Hamiltonian by a polynomial factor to switch between the two definitions.

Theorem 17 (Kitaev [39]). *(k, d) -TILH is in QMA_{EXP} .*

Proof. This does not trivially follow from the inclusion $\text{QMA} \subseteq \text{QMA}_{\text{EXP}}$ since the input size for TILH is just $\text{poly } \log n$. However, Kitaev's QMA verifier for the standard LOCAL HAMILTONIAN problem runs in time $\text{poly } n$, which is not polynomial in the input size for TILH. However, the exponential-time verifier of QMA_{EXP} offsets the logarithmically small input size, so the same random sampling argument as presented for QMA in e.g. [39, prop 14.2] goes through. \square

3.4. QMA versus QMA_{EXP}

In this section, we clarify why QMA_{EXP} is the natural class when considering the LOCAL HAMILTONIAN problem with translationally invariant interactions on a system of size n . When specifying a k -local Hamiltonian $\mathbf{H} = \sum_{i \in I} \mathbf{h}_i$, for some set of interactions I with $|I| = \text{poly } n$, we have to specify each term \mathbf{h}_i individually. Since the locality k and the local dimension d are constant, the total input size in Definition 1 is thus $l = \text{poly } n$ bits. In contrast, specifying a *translationally invariant* Hamiltonian \mathbf{H} requires only a logarithmic number of bits: since all local terms \mathbf{h}_i are identical and do not vary with the system size n , the only part of the input that varies with n and can thus be used to encode different instances of the problem is the system size n itself.

A fact which we will discuss in great detail in Sect. 6 is that the gap of a Hamiltonian encoding computation as a superposition of basis states—a so-called *history state construction*—scales inversely polynomially in the runtime, i.e. $1/\text{poly}(f(l))$ for an input of size l and an $f(l)$ -time computation. Contrasting this with the $1/\text{poly } n$ gap required by Definitions 1 and 16 independently—inverse polynomially in the system size, *not* the input size—we conclude the following core differences between QMA and QMA_{EXP} in the context of the LOCAL HAMILTONIAN problem (recall that n denotes the length of the spin chain and l denotes the total size of the input).

QMA A BQP verifier has $\text{poly } l$ runtime on an input of size l , so the gap of the Hamiltonian that encodes the verifier scales as $1/\text{poly } l$. This agrees with $1/\text{poly } n$ in Definition 1 since l and n are poly-related. QMA is thus the natural class for the LOCAL HAMILTONIAN problem.

QMA_{EXP} The BQEXP verifier can run for $\text{exp poly } l$ steps in the input size l . The gap therefore scales as $1/\text{exp poly } l$, which agrees with $1/\text{poly } n$ in Definition 16 since $l = \text{poly } \log n$. QMA_{EXP} is thus the natural class for TILH.

One fact we have glossed over is that even though each instance of TILH is translationally invariant, we could still vary the local interaction for each

system size n . As an example, assume that the Hamiltonian \mathbf{H} is specified by a single local term,

$$\mathbf{H} := \sum_{i=1}^{N-1} \mathbf{h}_{i,i+1} \quad \text{where} \quad \mathbf{h} = \begin{pmatrix} 1 & 0 \\ 0 & \alpha(m) \end{pmatrix}$$

with $\alpha(m) = \underbrace{3.1415926\dots 42}_{m \text{ digits of } \pi}$.

Then the bit complexity of this input is $O(m)$, and the overall input size—i.e. the possible information specifiable using the two parameters, the system size n and a varying parameter m , is thus $O(m + \log n)$. In order not to overspecify a LOCAL HAMILTONIAN or TILH problem, in each case, we have to require both bit precision and size of the input parameter to be of the same order (within polynomial factors). We conclude with the following remark.

Remark 18. It is natural to allow poly n precision of the entries in the local terms of the Hamiltonian when working with QMA, whereas for QMA_{EXP} local terms need to be precision-limited by poly $\log n$.

However, we want to emphasize that *we will only make use of uniformly scaling local interaction terms*, as in [3]: this in particular allows us to use coupling constants that scale polynomially in n . We also want to note that the polynomially closing promise gap of history state constructions might not be the end of the story; at this point in time it is not known whether quantum computation can be encoded into the ground state of a local Hamiltonian for which the promise gap e.g. scales in a sub-linear fashion in the number of computational steps (cf. [40] for an extended discussion).

In [7], the authors use a phase estimation algorithm to extract $O(n)$ bits of information from a fixed Hamiltonian term. However, in their construction, the speed at which the gap closes is irrelevant, as long as it remains nonzero in the gapped phase.

With a poly(n)-bounded computation and a $1/\text{poly}(n)$ gap, however, it is not clear how to do this computation in a translationally invariant manner. For phase estimation of m bits, one requires gates of precision $O(\exp(-m))$, cf. [37]—the algorithm depends on being able to perform a unitary \mathbf{U} an exponential number of times, i.e. $\mathbf{U}, \mathbf{U}^2, \mathbf{U}^4, \dots, \mathbf{U}^{2^{m-1}}$. Without having direct access to all powers of this gate—which we do not, if we require bounded local dimension and locality—we need to approximate them in some way: using the *Solovay-Kitaev theorem* with the required exponential precision $O(\exp(-m))$ results in a circuit of size $\text{poly} \log 1/(\exp(-m)) = \text{poly } m$, which limits the amount of information we can extract to $m = O(\log n)$.

It is clear that this is a problem of bootstrapping. For TILH, we only have $O(\log n)$ information available to start the computation with, and again it is not known whether there exists a more direct way of extracting a phase without having to go through the *Solovay-Kitaev theorem*, which only gives a sufficient upper bound to approximate the phase estimation algorithm.

3.5. Laplacian Matrix and Algebraic Connectivity of Graphs

In this section, we revise general notation and basic results from graph theory. For more background, consult the standard references [41–43] on graph theory and algebraic graph theory, respectively.

Definition 19. An *undirected simple graph* $G = (V, E)$ consists of a set of vertices V and a set of edges E , each edge being an unordered pair of distinct elements of V (in particular, there are no self-loops and no multiple edges). If the number of vertices is $n = |V|$ and we label them as $V = \{v_1, \dots, v_n\}$, then the *adjacency matrix* of G is $\mathbf{A}(G) := (a_{ij})_{1 \leq i, j \leq n}$ where

$$a_{ij} := \begin{cases} 1 & \{v_i, v_j\} \in E, \\ 0 & \text{otherwise.} \end{cases}$$

We further define the *degree matrix* $\mathbf{D}(G) := \text{diag}((\deg v_i)_{1 \leq i \leq n})$ where $\deg v_i := \sum_{j=1}^n a_{ij}$.

We will usually omit the qualifiers “undirected” and “simple” in the rest of this paper. We proceed to introduce basic notions and facts from algebraic graph theory [43].

Definition 20 (*Laplacian matrix*). The *Laplacian matrix* of a graph G is defined as $\Delta(G) := \mathbf{D}(G) - \mathbf{A}(G)$.

Since $\mathbf{A}(G)$ and $\Delta(G)$ are linear operators on \mathbb{C}^n where $n = |V|$, it will often be convenient to label the basis vectors of this space by $|v\rangle$ where $v \in V$ and denote the space itself by \mathbb{C}^V .

Definition 21. We write $\lambda_{\min}(\mathbf{M})$ to denote the smallest eigenvalue of Hermitian operator \mathbf{M} . If $\mathbf{M} \geq 0$ then $\lambda_{\min}(\mathbf{M}|_{\text{supp } \mathbf{M}})$ denotes the smallest *nonzero* eigenvalue of \mathbf{M} .

Claim 22. For any graph G , $\Delta(G)$ is real symmetric. In fact, $\Delta(G)$ is positive semi-definite with smallest eigenvalue $\lambda_{\min}(\Delta(G)) = 0$ and corresponding eigenvector $(1, \dots, 1)$.

Proof. By construction, $\mathbf{A}(G)$ and $\mathbf{D}(G)$ are real symmetric and so is $\Delta(G)$. The second claim follows by observing that $\Delta(G)$ is symmetric and diagonally dominant. Alternatively, $\Delta(G)$ can be expressed as a sum of positive semi-definite matrices:

$$\Delta(G) = \sum_{\{a,b\} \in E} (|a\rangle - |b\rangle)(\langle a| - \langle b|), \tag{4}$$

where each term is a principal submatrix of the form

$$\begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$$

and encodes the Laplacian of a single edge. The last statement follows from the fact that the row sums of $\Delta(G)$ are zero. \square

Definition 23 (*Algebraic connectivity*). The second smallest eigenvalue of the Laplacian $\Delta(G)$ is denoted with $a(G)$ and called the *algebraic connectivity* of graph G . The corresponding eigenvector is known as the *Fiedler vector*.

Claim 24 (Fiedler [44]). *For any graph G , $a(G) > 0$ if and only if G is connected.*

Lemma 25. *If $G = G_1 \sqcup \dots \sqcup G_m$ is a disjoint union of connected components G_i then $\Delta(G)$ has eigenvalue 0 with multiplicity exactly m and the next smallest eigenvalue is $\lambda_{\min}(\Delta(G)|_{\text{supp } \Delta(G)}) = \min_i a(G_i)$. Furthermore, $\{|\Phi_1\rangle, \dots, |\Phi_m\rangle\}$ with*


$$|\Phi_i\rangle := \frac{1}{\sqrt{|V_i|}} \sum_{v \in V_i} |v\rangle$$

is an orthonormal basis of the 0-eigenspace (ground space) of Δ .

Proof. Note that $\Delta = \Delta_1 \oplus \dots \oplus \Delta_m$ where Δ_i is the Laplacian of G_i . Recall from Claim 22 that $\Delta_i \geq 0$ and $\lambda_{\min}(\Delta_i) = 0$, hence the m smallest eigenvalues of Δ are equal to 0. Since each G_i is connected, $a(G_i) > 0$ for every i by Claim 24. Hence, the multiplicity of eigenvalue 0 must be m and the $(m + 1)$ -st smallest eigenvalue of Δ is positive and equal to $a(G_i)$ for some i . Finally, recall from Claim 22 that the uniform superposition over all vertices V_i of G_i is a 0-eigenvector of Δ_i , thus $\Delta |\Phi_i\rangle = 0$ for each $i \in \{1, \dots, m\}$. There are no further vectors in the ground space of Δ since Δ has eigenvalue 0 with multiplicity m . □

Corollary 26. *If Δ is the Laplacian of graph $G = (V, E)$ and $U \subseteq V$ is some connected component of G , then $|\Phi_U\rangle := \sum_{v \in U} |v\rangle / \sqrt{|U|}$ is a 0-eigenvector of Δ . In fact, any 0-eigenvector of Δ is a linear combination of such vectors.*

Claim 27 (Fiedler [44]). *Let G_L be the path graph on L vertices:*



Then $a(G_L) = 2(1 - \cos(\pi/L)) \sim \pi^2/L^2$. In particular, $a(G_L) = \Theta(1/L^2)$.

Corollary 28. *Let G be a connected graph with L vertices. Then $a(G) = \Omega(1/L^2)$.*

Proof. The algebraic connectivity is non-decreasing under adding edges [44, corollary 3.2], so for any connected graph on L vertices it is lower-bounded by that of a path graph on L vertices, which is given by Claim 27. □

3.6. Kitaev’s Geometrical Lemma for Graphs

We will need *Kitaev’s geometrical lemma* (see Lemma 14.4 in [39]) whose proof is reproduced below.

Lemma 29 ([39], p. 147). *Using notation from Definition 21, assume $\mathbf{A}, \mathbf{B} \geq 0$ are such that $\lambda_{\min}(\mathbf{A}|_{\text{supp } \mathbf{A}}) \geq \mu$ and $\lambda_{\min}(\mathbf{B}|_{\text{supp } \mathbf{B}}) \geq \mu$, and the null spaces of \mathbf{A} and \mathbf{B} have no vector in common other than 0, i.e. $\ker \mathbf{A} \cap \ker \mathbf{B} = \{0\}$.*

Then $\lambda_{\min}(\mathbf{A} + \mathbf{B}) \geq 2\mu \sin^2 \frac{\theta}{2}$, where θ is the angle between subspaces $\ker \mathbf{A}$ and $\ker \mathbf{B}$, i.e.

$$\cos \theta := \max_{\substack{|\alpha\rangle \in \ker \mathbf{A} \\ |\beta\rangle \in \ker \mathbf{B}}} |\langle \alpha | \beta \rangle|$$

where $|\alpha\rangle$ and $|\beta\rangle$ are unit vectors.

Proof. We define $\Pi_{\mathbf{A}}$ to be the projector onto $\ker \mathbf{A}$, and analogously for $\Pi_{\mathbf{B}}$. It follows from $\lambda_{\min}(\mathbf{A}|_{\text{supp } \mathbf{A}}) \geq \mu$ that $\mathbf{A} \geq \mu(\mathbf{1} - \Pi_{\mathbf{A}})$ and similarly for \mathbf{B} . It is hence enough to show that $(\mathbf{1} - \Pi_{\mathbf{A}}) + (\mathbf{1} - \Pi_{\mathbf{B}}) \geq (2 \sin^2 \frac{\theta}{2})\mathbf{1}$, which is equivalent to $(1 + \cos \theta)\mathbf{1} \geq \Pi_{\mathbf{A}} + \Pi_{\mathbf{B}}$. In other words, we want to show that every eigenvalue λ of $\Pi_{\mathbf{A}} + \Pi_{\mathbf{B}}$ satisfies

$$1 + \cos \theta \geq \lambda. \tag{5}$$

Let $|\psi\rangle$ be a normalized eigenvector of $\Pi_{\mathbf{A}} + \Pi_{\mathbf{B}}$ with eigenvalue $\lambda \geq 0$. Since Eq. (5) holds trivially for $\lambda = 0$, we can assume $\lambda > 0$. Since $\Pi_{\mathbf{A}}$ projects onto $\ker \mathbf{A}$, we can find a unit vector $|\psi_{\mathbf{A}}\rangle \in \ker \mathbf{A}$ such that $\Pi_{\mathbf{A}} |\psi\rangle = a |\psi_{\mathbf{A}}\rangle$ for some $a \in \mathbb{C}$; we can adjust the global phase of $|\psi_{\mathbf{A}}\rangle$ to guarantee that $a \geq 0$. Similarly, $\Pi_{\mathbf{B}} |\psi\rangle = b |\psi_{\mathbf{B}}\rangle$ for some unit vector $|\psi_{\mathbf{B}}\rangle \in \ker \mathbf{B}$ and $b \geq 0$. Since $\Pi_{\mathbf{A}}$ and $\Pi_{\mathbf{B}}$ are projectors, $\langle \psi | \Pi_{\mathbf{A}} | \psi \rangle = \langle \psi | \Pi_{\mathbf{A}}^\dagger \Pi_{\mathbf{A}} | \psi \rangle = a^2 \langle \psi_{\mathbf{A}} | \psi_{\mathbf{A}} \rangle = a^2$ and $\langle \psi | \Pi_{\mathbf{B}} | \psi \rangle = b^2$. From $\lambda |\psi\rangle = (\Pi_{\mathbf{A}} + \Pi_{\mathbf{B}}) |\psi\rangle$ we get by linearity that

$$\lambda = \langle \psi | (\Pi_{\mathbf{A}} + \Pi_{\mathbf{B}}) | \psi \rangle = a^2 + b^2.$$

Furthermore,

$$\begin{aligned} \lambda^2 &= \langle \psi | (\Pi_{\mathbf{A}} + \Pi_{\mathbf{B}})^2 | \psi \rangle = a^2 + b^2 + 2ab \operatorname{Re} \langle \psi_{\mathbf{A}} | \psi_{\mathbf{B}} \rangle \\ &\leq \lambda + 2ab |\langle \psi_{\mathbf{A}} | \psi_{\mathbf{B}} \rangle| \leq \lambda + (a^2 + b^2) |\langle \psi_{\mathbf{A}} | \psi_{\mathbf{B}} \rangle| = \lambda(1 + |\langle \psi_{\mathbf{A}} | \psi_{\mathbf{B}} \rangle|) \\ &\leq \lambda \left[1 + \max_{\substack{|\alpha\rangle \in \ker \mathbf{A} \\ |\beta\rangle \in \ker \mathbf{B}}} |\langle \alpha | \beta \rangle| \right] = \lambda(1 + \cos \theta), \end{aligned}$$

and hence $\lambda \leq 1 + \cos \theta$, which proves Eq. (5). □

We want to use Kitaev’s geometrical lemma to lower-bound the smallest eigenvalue of a graph Laplacian when certain vertices are penalized. To be more specific, for a graph $G = (V, E)$ and a set of vertices $P \subsetneq V$, we write a *penalizing matrix*

$$\mathbf{P}(G, P) := \sum_{v \in P} |v\rangle\langle v|.$$

A priori, it is not clear at all what the spectrum of the *penalized Laplacian* $\Delta + \mathbf{P}$ is, however we can obtain a lower bound on the smallest eigenvalue.

Lemma 30 (Kitaev’s geometrical lemma for graphs). *Let $G = (V, E)$ be a connected graph. Pick a non-empty subset of penalized vertices $P \subsetneq V$ and write the penalized Laplacian as $\Delta_P(G) := \Delta(G) + \mathbf{P}(G, P)$. Then $\lambda_{\min}(\Delta_P) = \Omega(1/|V|^3)$.*

Proof. Let us first verify that Δ and \mathbf{P} satisfy the prerequisites of Lemma 29. Since G is connected, $\lambda_{\min}(\Delta|_{\text{supp } \Delta}) = a(G) > 0$ by Claim 24. Moreover, $\ker \Delta$ is spanned by the all-ones vector $|\Phi_V\rangle := \sum_{v \in V} |v\rangle / \sqrt{|V|}$ according to Corollary 26. Clearly, $\lambda_{\min}(\mathbf{P}|_{\text{supp } \mathbf{P}}) = 1$ and $|\Phi_V\rangle \notin \ker \mathbf{P}$ since $\mathbf{P} \neq \emptyset$, so $\ker \Delta \cap \ker \mathbf{P} = \{0\}$. We can take the constant in Lemma 29 to be $\mu := \min\{a(G), 1\} = \Omega(1/|V|^2)$, where we used the lower bound $a(G) = \Omega(1/|V|^2)$ from Corollary 28 on the algebraic connectivity of G . It remains to compute the angle θ between $\ker \Delta = \text{span}\{|\Phi_V\rangle\}$ and $\ker \mathbf{P} = \text{span}\{|v\rangle : v \notin \mathbf{P}\}$. We have:

$$\cos \theta = \langle \Phi_V | \left(\frac{1}{\sqrt{|V| - |\mathbf{P}|}} \sum_{v \notin \mathbf{P}} |v\rangle \right) \rangle = \frac{|V| - |\mathbf{P}|}{\sqrt{|V|(|V| - |\mathbf{P}|)}} = \sqrt{1 - \frac{|\mathbf{P}|}{|V|}}$$

and hence

$$2 \sin^2 \frac{\theta}{2} = 1 - \cos \theta = 1 - \sqrt{1 - \frac{|\mathbf{P}|}{|V|}} \geq \frac{1}{2} \frac{|\mathbf{P}|}{|V|} \geq \frac{1}{2|V|}.$$

We conclude by Lemma 29 that $\lambda_{\min}(\Delta_{\mathbf{P}}) \geq 2\mu \sin^2 \frac{\theta}{2} = \Omega(1/|V|^3)$. □

4. Quantum Ring Machine

4.1. Definition

We define a new computational model, a *quantum ring machine* (QRM), and show that it is poly-time equivalent to a uniform class of quantum circuits. Recall that any uniform class of circuits—such as poly-time circuits or exponential-time circuits—inherits its uniformity condition from the corresponding class of classical Turing machines producing these circuit families. To prove that QRMs are quantum-universal, we will encode the given Turing machine into a specific instance of a QRM whose inner workings correspond to those of the original Turing machine, but with an additional quantum tape. In other words, the local Hilbert space of our QRM will be partitioned into two parts: a classical part, storing individual cells of the TM’s tape and the internal state of the TM, and a quantum part, storing one qubit per cell. However, since a general QRM does not need to have this specific internal structure, we first give an abstract definition.

Definition 31. A *quantum ring machine* (QRM) is a tuple $(\mathbf{R}, n, |\psi_{\text{in}}\rangle, \mathcal{H}_{\text{halt}})$, where

- $\mathbf{R} \in U(\mathbb{C}^d \otimes \mathbb{C}^d)$ is a unitary operator on a pair of qudits, each of dimension d ,
- $n \in \mathbb{N}$ is the total number of qudits on the ring,
- $|\psi_{\text{in}}\rangle \in \mathcal{H}$ is the *initial state* where $\mathcal{H} := (\mathbb{C}^d)^{\otimes n}$ denotes the joint Hilbert space,
- $\mathcal{H}_{\text{halt}} \subseteq \mathbb{C}^d$ is the *halting subspace* of each qudit.

Starting from a ring of n qudits initialized in $|\psi_{\text{in}}\rangle$, the operation \mathbf{R} is applied cyclically to adjacent pairs of qudits—see Fig. 5—until some qudit indicates




FIGURE 5. Quantum ring machine (QRM). Starting from a ring of qudits $\mathcal{H} = (\mathbb{C}^d)^{\otimes n}$ in an initial configuration $|\psi_{\text{in}}\rangle \in \mathcal{H}$, a unitary $\mathbf{R} \in U(\mathbb{C}^d \otimes \mathbb{C}^d)$ is applied to pairs of adjacent qudits until one of them is completely in some halting subspace $\mathcal{H}_{\text{halt}} \subseteq \mathbb{C}^d$




FIGURE 6. Circuit diagram of a QRM with a ring of size 4. The double lines indicate classical wires that are used to store the TM’s internal states and tape, as well as a flag indicating either the TM’s halting or the direction of its next head movement (see the proof of Lemma 33 for more details). The internal details of the QRM’s unitary operation \mathbf{R} are shown in Fig. 7

halting: its reduced density matrix has support completely within the halting subspace $\mathcal{H}_{\text{halt}}$; up until that point, the probability of finding any qudit within $\mathcal{H}_{\text{halt}}$ is zero.¹¹

Figure 6 visualizes a QRM as a quantum circuit. Because the ring is cyclic, we can arbitrarily mark a starting position on the ring. Starting at this position, part of the initial state $|\psi_{\text{in}}\rangle$ contains the input while the rest will be used as a workspace. The *input size* is thus upper-bounded by the ring size.

In the following definition, we consider a slight extension of QRMs from Definition 31 where $|\psi_{\text{in}}\rangle$ is replaced by a family of input states $\{|\psi_{\text{in}}(x)\rangle\}_{x \in I}$ for some index set I .

¹¹ In particular this means that if, after every application of \mathbf{R} , the corresponding two qudits are measured, each with respect to $\mathcal{H}_{\text{halt}}$ and its orthogonal complement $\mathcal{H}_{\text{halt}}^\perp$, then the probability of finding the reduced state in $\mathcal{H}_{\text{halt}}$ should always be either zero or one, with the latter case indicating halting.

Definition 32. A QRM *terminates* on $\{|\psi_{\text{in}}(x)\rangle\}_{x \in I}$ if it halts, in finitely many steps, on any initial state $|\psi_{\text{in}}(x)\rangle$ for $x \in I$. Let $(M_n)_n$ be a family of QRMs where M_n has a ring of size n . This family is *poly-time terminating* if there exists a polynomial p such that M_n terminates in $p(n)$ steps on all states $|\psi_{\text{in}}(x)\rangle$; similarly, it is *exponential-time terminating* if there exists an exponential function $f(n) = O(\exp(cn))$, for some $c > 0$, such that M_n terminates in $f(n)$ steps on all states $|\psi_{\text{in}}(x)\rangle$.

4.2. Universality

Lemma 33. Let $\Pi = (\Pi_{\text{YES}}, \Pi_{\text{NO}})$ be a promise problem in BQP. Then there exists a polynomial p and a poly-time terminating family of QRMs $(M_n)_n$,

$$M_n = (\mathbf{R}, n, \{|\psi_{\text{in}}(x)\rangle\}_{x \in I_n}, \mathcal{H}_{\text{halt}}),$$

with the following properties:

1. All M_n share the same unitary \mathbf{R} and the same terminating subspace $\mathcal{H}_{\text{halt}}$. The ring size of M_n is n .
2. The input states $\{|\psi_{\text{in}}(x)\rangle\}_{x \in I_n}$ of each M_n consist of *trivial*¹² encodings of instances $I_n := \{x \in \Pi : p(|x|) = n\}$, so that the whole computation fits on a ring of size n .
3. If $x \in \Pi_{\text{YES}}$, the reduced density matrix of the cell that signals halting satisfies an extra constraint: if measured, it collapses to an accepting subspace $\mathcal{H}_{\text{acc}} \subseteq \mathcal{H}_{\text{halt}}$ with probability $\geq 2/3$.
4. If $x \in \Pi_{\text{NO}}$, it collapses to $\mathcal{H}_{\text{acc}} \subseteq \mathcal{H}_{\text{halt}}$ with probability $\leq 1/3$.

Proof. Our goal is to construct a QRM for simulating a classical Turing machine (TM) that produces a description of a uniform quantum circuit. In addition to computing the circuit’s description, the QRM also executes it one gate at a time. More formally, the QRM simulates a deterministic and reversible TM (see Definition 7) augmented with the following quantum features:

- in addition to the classical data, each cell of the TM’s tape stores one qubit,
- a special subset of the TM’s states is associated with a universal set of two-qubit quantum gates; whenever the TM enters one of these states, the corresponding gate is applied on the two adjacent qubits that are stored in the pair of cells between which the TM’s head just moved.

It is straightforward to verify that such quantum-enhanced TM is equivalent to a uniform family of quantum circuits.

Let us now describe the simulation procedure more formally. We write the complex linear span of a finite set S as $\mathbb{C}^S := \text{span}\{|s\rangle \in \mathbb{C}^{|S|} : s \in S\}$ and refer to $(\mathbb{C}^S)^{\otimes n}$ as a *ring* of size $n \in \mathbb{N}$, where each copy of \mathbb{C}^S represents one *cell* of the ring. Each cell further consists of three registers: a *quantum bit* (labelled by $\{0, 1\}$), a *classical data* register (labelled by elements of some

¹² One must be able to produce $|\psi_{\text{in}}(x)\rangle$ from $x \in \Pi$ with a constant-depth quantum circuit (in particular, one cannot cheat by allowing the input to contain the answer to the problem), e.g. see Eq. (6). This is similar to the types of input encodings one would allow for a poly-time classical TM.

finite set Γ), and a *flag* register (labelled by another set F). The standard basis of each ring cell is thus labelled by triples of the form

$$S := \{0, 1\} \times \Gamma \times F.$$

Using the notation from Definition 7, let (Q, Σ, δ) be the deterministic TM we want to simulate (it is reversible without loss of generality, see Theorem 9). The first register of the QRM stores the quantum state obtained by executing the quantum circuit produced by the TM. The second register Γ stores the TM’s internal state and tape, so $\Gamma := Q \times \Sigma$ where Q is the set of internal states and Σ is the TM’s alphabet. The flag symbols F in the third register are used to mark the location of the TM’s head.

The flag register’s alphabet is given by

$$F := \{\leftarrow, \rightarrow, -, h\}$$

and is used as follows. At any time, exactly one cell on the ring contains an *active flag* (either “ \leftarrow ”, “ \rightarrow ”, or “ h ”), while the rest are padded with “ $-$ ”. The TM’s internal state is always stored in the active cell. Unless the TM has halted (indicated by flag “ h ”), the active flag shows in which direction (“ \leftarrow ” for *left* and “ \rightarrow ” for *right*) the TM’s head must be moved before the simulation of the next step can begin. Every time the TM’s head moves or its internal state changes, the QRM updates the flag registers and the description of the TM’s internal state accordingly. Whenever the TM enters one of the special “quantum” states, the QRM applies the corresponding two-qubit unitary.

Recall from Definition 31 that QRM operates by cyclically applying a fixed unitary \mathbf{R} on pairs of consecutive cells along the ring (see Fig. 6). Most of the time \mathbf{R} acts trivially, since a non-trivial action is triggered only when either of the two active flags “ \leftarrow ” or “ \rightarrow ” is encountered. Note that \mathbf{R} acts on *two* adjacent cells, one of them marked by the active flag and the other indicated by the direction of the flag’s arrow:



It is crucial that \mathbf{R} is two-local for the following two reasons. First, updating the active location requires changing two symbols (e.g. when the TM’s head moves left, we need to replace “ $-\leftarrow$ ” by “ $\leftarrow-$ ” or “ $\rightarrow-$ ”, depending on the direction the head will move next). For applying a two-qubit gate, we clearly also need a two-local interaction (we use the same convention as above to determine on which two qubits the gate is applied).

Recall from Lemma 8 that, instead of quintuples δ , we can work with a permutation matrix \mathbf{T}_δ on Γ and a function $d : Q \rightarrow \{\text{left, right}\}$ telling us where the TM’s head came from. For convenience, we include a special *dummy* state “ \perp ” in Q and a designated *blank* symbol “ $_$ ” in Σ : the dummy state is stored in all cells (except the active cell which stores the actual state of the TM), while the blank symbol is used to initialize the TM’s tape. We accordingly extend \mathbf{T}_δ so that it acts trivially on $|\perp, \sigma\rangle$ for any $\sigma \in \Sigma$, and we define $d(\perp) := -$ so that dummy states do not trigger any action in our simulation.

We take the ring size to be $n = p(|x|)$ for an instance $x \in \Pi$, since the TM can access at most that many tape cells. We require that the ring starts out in a well-formed state, i.e. for some binary representation $x = x_1x_2 \cdots x_l$ and $l = |x|$, a state of the form

$$|\psi_{\text{in}}(x)\rangle := \bigotimes_{j=1}^{n-l-1} (|0\rangle \otimes |\perp, _ \rangle \otimes |-\rangle) \otimes \bigotimes_{i=1}^l (|x_i\rangle \otimes |\perp, _ \rangle \otimes |-\rangle) \otimes (|0\rangle \otimes |q_0, _ \rangle \otimes |-\rangle), \quad (6)$$

i.e. where all cells but the last are initialized as follows: the TM is in the dummy state “ \perp ”, the TM’s tape is initialized to a designated blank symbol “ $_$ ”, and the flag is set to “ $-$ ”. The last cell contains the TM’s initial state q_0 and the “ \rightarrow ” flag. The input $x \in \Pi$ is written on the qubit part of the tape, i.e. the first register of each cell.

We can now describe in more detail the steps involved in our simulation, and how to perform them reversibly (see Figs. 7 and 8 for more details):

1. If the active cell has the halting flag “ h ”, the TM has halted so nothing happens.
2. If the active cell has one of the other two flags “ \leftarrow ” or “ \rightarrow ”:
 - (a) The Q part of the Γ registers of the active cell and its neighbour—indicated by the flag—are exchanged, thus simulating the movement of the TM’s head.
 - (b) The flag register of the active cell is uncomputed using the function d .




FIGURE 7. Circuit diagram for implementing the QRM’s head unitary \mathbf{R} (double wires are classical while single wires are quantum). All computation is classical except for a single classically controlled quantum gate \mathbf{U} that can be triggered by either of the two Γ registers. A classical circuit for implementing \mathbf{C} is shown in Fig. 8




FIGURE 8. Circuit diagram for implementing the classical permutation \mathbf{C} in Fig. 7 (all wires are classical and all gates are reversible). Conditioned on the flag registers being either “ \rightarrow -” or “- \leftarrow ”, the controlled-controlled-SWAP gate exchanges the internal state registers of the two cells. The CNOT gates in the first layer are conditioned on the value of $d(q)$, for state q , and they uncompute the flag register of the opposite cell (the cell where the TM’s head came from). The permutation \mathbf{T}_δ acts on Γ registers of both cells to update the TM’s internal state and the current tape symbol. Recall that \mathbf{T}_δ acts trivially if the state is dummy (at most one of the cells is in a non-dummy state). The final layer of CNOT gates again condition on $d(q')$, where q' is the new state, and update the flag registers to indicate where the TM’s head will move next. These flags will be uncomputed by the next iteration

- (c) Description of the TM’s internal state and the current tape symbol is updated using \mathbf{T}_δ .
- (d) Based on the updated internal state, a new flag register is computed using d (it belongs to the same cell where the TM’s new state is stored, and it indicates in which direction the TM’s head will move before the next iteration begins).
- (e) If the TM is in one of the special states indicating a quantum gate, the corresponding unitary is applied on the two data registers.

We now describe the unitary operator \mathbf{R} that acts on two adjacent QRM’s cells:

- For $a = (\psi_a, \gamma_a, f_a) \in \mathcal{S}$, write the corresponding basis state as $|a\rangle := |\psi_a, \gamma_a, f_a\rangle \in \mathbb{C}^{\mathcal{S}}$ where $\psi_a \in \{0, 1\}$, $\gamma_a \in \Gamma$, $f_a \in F$, and analogously for $|b\rangle$. Then $|a\rangle \otimes |b\rangle \in \mathbb{C}^{\mathcal{S} \times \mathcal{S}}$ is also a basis state and we require, up to reordering the registers (see Fig. 7), that

$$\mathbf{R}(|a\rangle \otimes |b\rangle) = |\Psi\rangle \otimes |\gamma'_a, f'_a\rangle \otimes |\gamma'_b, f'_b\rangle$$

for some $\gamma'_a, \gamma'_b \in \Gamma$ and $f'_a, f'_b \in F$, i.e. \mathbf{R} acts classically on each register except for the quantum data registers $|\psi_a\rangle$ and $|\psi_b\rangle$ (in particular, we allow $|\Psi\rangle \in \mathbb{C}^2 \otimes \mathbb{C}^2$ to be entangled).

- Using the same notation, if $f_a \neq \rightarrow$ and $f_b \neq \leftarrow$, we further demand $|\gamma'_a, f'_a\rangle = |\gamma_a, f_a\rangle$, $|\gamma'_b, f'_b\rangle = |\gamma_b, f_b\rangle$, and $|\Psi\rangle = |\psi_a\rangle \otimes |\psi_b\rangle$, i.e. if neither f_a nor f_b signal “apply head here”, \mathbf{R} acts as the identity operator on all registers.
- The active flag always moves in the direction indicated by the arrow. If $f_a = \rightarrow$ then $f'_a = -$ and $f'_b \neq -$, meaning that the head has moved right. Similarly, if $f_b = \leftarrow$ then $f'_b = -$ and $f'_a \neq -$, meaning that the head has moved left. In each case there are three possible transitions—they indicate whether the TM has halted or in which direction its head has to move next:

$$\begin{array}{cc}
 f_a f_b & f_a f_b \\
 \rightarrow - & - \leftarrow \\
 \Downarrow & \Downarrow \\
 f'_a f'_b & f'_a f'_b \\
 - \rightarrow & \rightarrow - \\
 - \leftarrow & \leftarrow - \\
 - h & h -
 \end{array}$$

Figure 7 shows how \mathbf{R} acts on two adjacent cells. For each cell, the halting subspace $\mathcal{H}_{\text{halt}}$ is spanned by all standard basis vectors with the last register in the halting state $|h\rangle$:

$$\mathcal{H}_{\text{halt}} := \mathbb{C}^2 \otimes \mathbb{C}^\Gamma \otimes |h\rangle.$$

Figure 8 provides details on how to implement \mathbf{C} reversibly.

We construct the desired family of QRMs $(M_n)_n$ in the special form described above. It is straightforward to verify that this ring machine executes the circuit written out by the TM, and the runtime overhead of M_n as compared to the circuit is at most quadratic. \square

Corollary 34. *Using an exponential-time terminating family of QRMs, Lemma 33 holds for BQEXP as well.*

5. Unitary Labelled Graphs

5.1. Definitions

The following definition introduces graphs whose vertices are labelled by Hilbert spaces and whose edges are labelled by unitaries between these spaces.

Definition 35. Given an undirected graph $G = (S, E)$ without self-loops, a *unitary labelled graph* (ULG) is a triple $(G, (\mathcal{H}_v)_{v \in S}, g)$ where

- $(\mathcal{H}_v)_{v \in S}$ is a family of Hilbert spaces, one space \mathcal{H}_v for each vertex $v \in S$,
- g is a function that assigns to each directed¹³ edge $ab \in E$ some unitary operator $g(ab) \in U(\mathcal{H}_a)$ so that $g(ab) = g(ba)^\dagger$ (this requires that $\mathcal{H}_a \cong \mathcal{H}_b$ whenever $ab \in E$).

¹³ While G is an undirected graph, we need to arbitrarily direct its edges so that we can discern between labels \mathbf{U} and \mathbf{U}^\dagger assigned to edges ab and ba , respectively.




FIGURE 9. Unitary labelled graph from Example 36. Observe that we mark the direction for the unitaries with an arrow, despite working with undirected graphs

To facilitate notation, we will write an edge and its associated unitary jointly as $(a \leftrightarrow b, \mathbf{U})$ and call it a *rule*. By definition, the rule $(a \leftrightarrow b, \mathbf{U})$ is equivalent to the rule $(b \leftrightarrow a, \mathbf{U}^\dagger)$. With this notation, it is convenient to specify a unitary labelled graph by $G = (S, \mathbf{R})$ where $\mathbf{R} := \{(a \leftrightarrow b, g(ab)) : ab \in E\}$ is the corresponding set of rules.

Example 36. Let $S := \{1, 2, 3, 4, 5\}$, $\mathcal{H} = \mathbb{C}^2$ for all vertices, and consider the following set of rules:

$$\mathbf{R} := \{(1 \leftrightarrow 2, \mathbf{U}), (2 \leftrightarrow 3, \mathbf{V}), (3 \leftrightarrow 4, \mathbf{W}), (4 \leftrightarrow 1, \mathbf{1}), (4 \leftrightarrow 5, \mathbf{1})\}.$$

The underlying graph for this example is shown in Fig. 9.

Definition 37. Let $G = (S, \mathbf{R})$ be a ULG. If the product of unitaries along any directed path connecting a and b is equal, and this property holds for all $a, b \in S$, we call the ULG *simple*. Equivalently, for a ULG to be simple, the product of unitaries along any directed cycle should be $\mathbf{1}$.

The ULG in Example 36 is simple if and only if $\mathbf{WVU} = \mathbf{1}$.

The following definition assigns a Hamiltonian to each ULG. This Hamiltonian extends the notion of a graph Laplacian, see Definition 20, to ULGs (while this might not be immediately obvious from the definition, it will be made more clear in Lemma 41 below).

Definition 38. Let $G = (S, \mathbf{R})$ be a connected ULG, \mathcal{H} denote the Hilbert space attached to each of its vertices, $n := \dim \mathcal{H}$ be the dimension of \mathcal{H} , and let $\{|e_i\rangle\}_{i=1}^n$ be some orthonormal basis of \mathcal{H} . The *Hamiltonian associated with G* is the following Hermitian operator on $\mathbb{C}^S \otimes \mathcal{H}$:

$$\mathbf{H}(G) := \sum_{(a \leftrightarrow b, \mathbf{U}) \in \mathbf{R}} \sum_{i=1}^n (|a\rangle \otimes |e_i\rangle - |b\rangle \otimes \mathbf{U}|e_i\rangle)(\langle a| \otimes \langle e_i| - \langle b| \otimes \langle e_i| \mathbf{U}^\dagger). \tag{7}$$

This is reminiscent of Eq. (4) for $\Delta(G)$, the Laplacian of graph G . Furthermore, it also explains why we excluded self-loops in Definition 35: just as they have no effect on the graph Laplacian, they also impose no changes in the associated Hamiltonian of a simple ULG—the only possible self-loop unitary for such ULG is $\mathbf{1}$, making the corresponding term in Eq. (7) vanish.

Proposition 39. *The Hamiltonian $\mathbf{H} = \mathbf{H}(G)$ of a connected UGL $G = (S, R)$, see Definition 38, is invariant under replacing any rule $(a \leftrightarrow b, \mathbf{U}) \in R$ with the corresponding inverse rule $(b \leftrightarrow a, \mathbf{U}^\dagger)$. Moreover, the matrix entries of \mathbf{H} do not depend on the choice of the basis $\{|e_i\rangle\}_{i=1}^n$.*

Proof. Since G is connected, the Hilbert spaces attached to all its vertices are isomorphic. Observe further that

$$\mathbf{H} \equiv \sum_{(a \leftrightarrow b, \mathbf{U}) \in R} (|a\rangle\langle a| \otimes \mathbb{1}_n + |b\rangle\langle b| \otimes \mathbb{1}_n - |a\rangle\langle b| \otimes \mathbf{U}^\dagger - |b\rangle\langle a| \otimes \mathbf{U}), \tag{8}$$

hence the claim follows. □

Note that an alternative way of writing \mathbf{H} is as follows:

$$\mathbf{H} \equiv \sum_{(a \leftrightarrow b, \mathbf{U}) \in R} \sum_i (|a\rangle \otimes |i\rangle - |b\rangle \otimes \mathbf{U}|i\rangle)(\langle a| \otimes \langle i| - \langle b| \otimes \langle i| \mathbf{U}^\dagger),$$

which emphasizes the fact that each term is positive semi-definite.

One can extend the notion of an associated Hamiltonian to a non-connected ULG as well by taking a direct sum of the Hamiltonians for each component of the graph (equivalently, one can assume that the Hilbert spaces associated to different components of the graph are mutually orthogonal and take the new Hilbert space to be their direct sum). Either way, such extension yields a block-diagonal associated Hamiltonian.

5.2. Semi-classical Unitary Labelled Graphs

A ULG is semi-classical if its Hamiltonian is equal to a graph Laplacian (see Definition 20), after a unitary change of basis.

Definition 40. A ULG G is *semi-classical* if its associated Hamiltonian can be expressed as $\mathbf{H} = \mathbf{W}(\Delta \otimes \mathbb{1}_n)\mathbf{W}^\dagger$, where \mathbf{W} is some unitary operator, Δ is the Laplacian of G , and $\mathbb{1}_n$ acts on the n -dimensional Hilbert space attached to each vertex of G .

This definition can be easily extended also to non-connected ULGs as well.

The following lemma is important for analysing the spectrum of any Hamiltonian coming from a simple ULG. It reduces the problem to analysing instead the spectrum of the corresponding graph Laplacian.

Lemma 41. *Any simple ULG is semi-classical.*

Proof. Denote the UGL by $G = (S, R)$ where S and R are the sets of vertices and rules, respectively. If G has disjoint components, \mathbf{H} is block-diagonal and we can deal with each block separately, hence we can assume without loss of generality that G is connected and all its vertices have isomorphic attached Hilbert spaces \mathcal{H} .

Pick an arbitrary vertex $a \in S$ and denote its set of neighbours by S_a . Using proposition 39, rewrite R in a form where a only has outgoing edges. Following Eq. (8), define the term that encodes rule $(a \leftrightarrow b, \mathbf{U}_b) \in R$ as follows:




FIGURE 10. Terms of \mathbf{H} , see Eq. 10, grouped according to how far the corresponding edges are from the chosen vertex a . For this ULG to be simple, the labels of edges forming the triangle abc must satisfy $\mathbf{U}_c^\dagger \mathbf{U} \mathbf{U}_b = \mathbf{1}$. We have omitted the labels of all other edges

$$\mathbf{h}_{ab} := |a\rangle\langle a| \otimes \mathbf{1} + |b\rangle\langle b| \otimes \mathbf{1} - |a\rangle\langle b| \otimes \mathbf{U}_b^\dagger - |b\rangle\langle a| \otimes \mathbf{U}_b, \tag{9}$$

where the subscript of \mathbf{U}_b identifies the vertex with incoming edge. Then the terms of \mathbf{H} can be grouped as follows (see Fig. 10):

$$\begin{aligned} \mathbf{H} &= \sum_{\substack{(a \leftrightarrow b, \mathbf{U}_b) \\ b \in S_a}} \mathbf{h}_{ab} + \sum_{\substack{(b \leftrightarrow c, \mathbf{U}_c) \\ b, c \in S_a}} \mathbf{h}_{bc} + \sum_{\substack{(b \leftrightarrow d, \mathbf{U}_d) \\ b \in S_a \wedge d \notin S_a \cup \{a\}}} \mathbf{h}_{bd} + \mathbf{F} \\ &=: \mathbf{H}_a + \mathbf{H}_n + \mathbf{H}_{nn} + \mathbf{F}, \end{aligned} \tag{10}$$

where \mathbf{F} denotes the rest of the terms and all sums range over \mathbb{R} , with some restrictions on the endpoints of the edges. Our strategy now is to apply a sequence of unitary transformations to bring the Hamiltonian \mathbf{H} to the desired form, a few terms at a time.

First, for the given vertex $a \in S$, define the following unitary:

$$\mathbf{W}_a := \prod_{\substack{(a \leftrightarrow b, \mathbf{U}_b) \\ b \in S_a}} (|b\rangle\langle b| \otimes \mathbf{U}_b + (\mathbf{1} - |b\rangle\langle b|) \otimes \mathbf{1}).$$

Observe that all terms in the product commute. Moreover, $\mathbf{W}_a^\dagger \mathbf{F} \mathbf{W}_a = \mathbf{F}$ and, by Eq. (9),

$$\mathbf{W}_a^\dagger \mathbf{h}_{ab} \mathbf{W}_a = (|a\rangle\langle a| + |b\rangle\langle b| - |a\rangle\langle b| - |b\rangle\langle a|) \otimes \mathbf{1}$$

for all $b \in S_a$, so \mathbf{W}_a takes care of all terms of \mathbf{H}_a simultaneously.

For the terms in \mathbf{H}_n , pick any edge $(b \leftrightarrow c, \mathbf{U}) \in \mathbb{R}$, with $b, c \in S_a$, and note from Eq. (9) that

$$\mathbf{W}_a^\dagger \mathbf{h}_{bc} \mathbf{W}_a = |b\rangle\langle b| \otimes \mathbf{1} + |c\rangle\langle c| \otimes \mathbf{1} - |b\rangle\langle c| \otimes \mathbf{U}_b^\dagger \mathbf{U}^\dagger \mathbf{U}_c - |c\rangle\langle b| \otimes \mathbf{U}_c^\dagger \mathbf{U} \mathbf{U}_b.$$

However, since abc is a cycle and the ULG is simple, the product of unitaries along the cycle must be $\mathbb{1}$, i.e. $\mathbf{U}_c^\dagger \mathbf{U} \mathbf{U}_b = \mathbb{1}$ (see Fig. 10), so the formula simplifies to

$$\begin{aligned} \mathbf{W}_a^\dagger \mathbf{h}_{bc} \mathbf{W}_a &= (|b\rangle\langle b| + |c\rangle\langle c| - |b\rangle\langle c| - |c\rangle\langle b|) \otimes \mathbb{1} \\ &= (|b\rangle - |c\rangle)(\langle b| - \langle c|) \otimes \mathbb{1}. \end{aligned}$$

By a similar argument, we can show that the rules in \mathbf{H}_{nn} change their unitary by a factor of \mathbf{U}_b when outgoing, or \mathbf{U}_b^\dagger when incoming, respectively. We are left with a ULG with a new set of rules, namely, every edge that is either attached to a or between two different neighbours of a is trivial, i.e. the edge unitary is the identity operator $\mathbb{1}$.

We will apply the same procedure to different vertices until all edges become trivial. More specifically, we consider an arbitrary sequence of subsets $S_1 \subset S_2 \subset \dots \subset S_m \subset S$, starting from any vertex $\{a_1\} = S_1$ and ending with the set of all vertices S , such that each subsequent S_{k+1} can be obtained from S_k by including all neighbours of some vertex $a_k \in S_k$. The overall unitary is then $\mathbf{W} = \prod_{k=1}^m \mathbf{W}_{a_k}$, where the product is over the sequence of vertices $a_1, a_2, \dots, a_m \in S$. Each successive unitary \mathbf{W}_{a_k} is obtained from the current set of rules R_k , where $R_1 = R$ is the original set, while all rules in the final set are trivial. Our goal is to show that, at every step k , we can guarantee that each rule in R_k has a trivial unitary whenever both endpoints of the corresponding edge are in S_k .

We proceed by induction. Since there are no edges between vertices in $S_1 = \{a_1\}$, the induction basis is trivial. Assuming the inductive hypothesis holds for k , we take $a_k \in S_k$ and apply the unitary \mathbf{W}_{a_k} that acts non-trivially to all neighbours of a_k (recall that S_{k+1} is formed by S_k together with the neighbours of a_k). As discussed above, \mathbf{W}_{a_k} trivializes all edges between a_k and any of its neighbours. By simplicity of the ULG, it trivializes also all edges between any two different neighbours of a_k . Moreover, it does not affect any edges within S_k (they are trivial already by the inductive assumption). In other words, all edges between vertices in S_{k+1} are trivial, thus completing the induction.

Since all edge unitaries have now been transformed to the trivial unitary $\mathbb{1}$, all terms in $\mathbf{W}^\dagger \mathbf{H} \mathbf{W}$ are of the form

$$(|a\rangle\langle a| + |b\rangle\langle b| - |a\rangle\langle b| - |b\rangle\langle a|) \otimes \mathbb{1} = (|a\rangle - |b\rangle)(\langle a| - \langle b|) \otimes \mathbb{1}, \tag{11}$$

for some $a, b \in S$. Comparing this to Eq. (4), the overall Hamiltonian is in fact equivalent to the Laplacian Δ of G , i.e. $\mathbf{W}^\dagger \mathbf{H} \mathbf{W} = \Delta \otimes \mathbb{1}$. \square

Lemma 42. *Let G be a simple ULG with vertices S and rules R . Write $G = G_1 \oplus \dots \oplus G_N$, $G_i = (V_i, E_i)$ for the associated Laplacian of the induced classical graph (S, R') with $R' := \{a \leftrightarrow b : (a \leftrightarrow b, \mathbf{U}) \in R\}$, and pick an arbitrary $v_i \in V_i \forall i$. Let further $n_i := \dim \mathcal{H}_i$ and choose a basis $\{|e_{i,j}\rangle\}_j$ of \mathcal{H}_i for all connected graph components G_i . Then the ground space $\ker \mathbf{H}$ is spanned by the set $\{|\Psi_{i,0}\rangle, \dots, |\Psi_{i,n_i}\rangle\}_{i=1}^N$, where*

$$|\Psi_{i,j}\rangle := \frac{1}{\sqrt{|V_i|}} \sum_{s \in V_i} |s\rangle \otimes |q_s\rangle \quad \text{and} \quad |q_s\rangle = \begin{cases} |e_{i,j}\rangle & \text{if } s = v_i, \\ \mathbf{U} |q_r\rangle & \text{if } (r \leftrightarrow s, \mathbf{U}) \in E_i. \end{cases}$$

Furthermore, the $|\Psi_{i,j}\rangle$ form a basis of $\ker \mathbf{H}$.

Proof. Because the ULG is simple, by Lemma 41, there exists a unitary \mathbf{W} and a classical Laplacian Δ such that $\mathbf{W}^\dagger \mathbf{H} \mathbf{W} = \Delta \otimes \mathbf{1}_n$, and hence $\ker \mathbf{H} = \ker \Delta \otimes \mathbf{1}_n$. By Lemma 25, the ground space of $\Delta \otimes \mathbf{1}_n$ has a basis given by

$$|\Phi_{i,j}\rangle := \frac{1}{\sqrt{|V_i|}} \sum_{s \in V_i} |s\rangle \otimes |e_i\rangle, \quad i = 1, \dots, N \quad \text{and} \quad j = 1, \dots, n.$$

Observe that

$$\mathbf{W} |\Phi_{i,j}\rangle = \frac{1}{\sqrt{|V_i|}} \sum_{s \in V_i} \prod_{a \in V_i} \mathbf{W}_a |s\rangle \otimes |e_i\rangle = |\Psi_{i,j}\rangle,$$

which can be easily verified. □

5.3. Kitaev’s Geometrical Lemma for Unitary Labelled Graphs

Analogous to Sect. 3.6, we extend the notion of penalizing vertices to ULGs. First we state an immediate corollary from Lemma 30.

Corollary 43. *Take a connected simple ULG with Hilbert space \mathcal{H} for all vertices $s \in S$. Pick a non-empty subset of vertices $P \subsetneq S$ and write the penalized associated Hamiltonian $\mathbf{H}_P(G) := \mathbf{H}(G) + \mathbf{P}(G, P) \otimes \mathbf{1}_{\dim \mathcal{H}}$. Then $\lambda_{\min}(\mathbf{H}_P(G)) = \Omega(1/|S|^3)$.*

Proof. G is simple and $\mathbf{H}(G)$ has the same spectrum as $\Delta(G)$, up to multiplicity. Now use Lemma 30 on $\Delta(G) + \mathbf{P}(G, P)$. □

A more interesting case is when one does not want to penalize the entire Hilbert space attached to a vertex, but only a subspace. This is captured in the following lemma.

Lemma 44. *Take a connected simple ULG with Hilbert space \mathcal{H} for all vertices $s \in S$. Pick a non-empty subset of vertices $P \subsetneq S$ and a set of projectors $\Pi = \{\Pi_p\}_{p \in P}$ on \mathcal{H} . For $\mathbf{H}_P(G, \Pi) := \mathbf{H}(G) + \sum_{p \in P} |p\rangle\langle p| \otimes \Pi_p$, we have $\lambda_{\min}(\mathbf{H}_P(G, \Pi)) \geq \mu \Omega(1/|S|^3)$, where $\mu = 1 - \max\{|\lambda_{\max}(\Pi_i^c \mathbf{U}_{ij} \Pi_j^c)| : p_i, p_j \in P, i \neq j\}$ and \mathbf{U}_{ij} is the product of unitaries of a path connecting vertices p_i and p_j .*

Proof. First note that the \mathbf{U}_{ij} are well-defined, since the ULG is simple and connected. Construct \mathbf{W} such that the root of the spanning tree is one of the penalized vertices, namely r with projector Π_r . Then

$$\begin{aligned} \mathbf{W}^\dagger \mathbf{H}_P(G, \Pi) \mathbf{W} &= \Delta(G) \otimes \mathbf{1} + |r\rangle\langle r| \otimes \Pi_r + \sum_{p \neq r} |p\rangle\langle p| \otimes \mathbf{R}_p \Pi_p \mathbf{R}_p^\dagger \\ &=: \Delta \otimes \mathbf{1} + \mathbf{A} \end{aligned}$$

where the \mathbf{R}_p are the product of unitaries connecting vertex p with the root r . Following the notation of Lemma 30, we want to calculate the angle between

the kernels of the Laplacian and the penalty terms. We write $\Pi_{\mathbf{A}} := \mathbb{1} \otimes \mathbb{1} - \sum_p |p\rangle\langle p| \Pi_p$ for the projector onto the kernel of the penalty terms. Then

$$\mathbf{W}^\dagger \Pi_{\mathbf{A}} \mathbf{W} = |r\rangle\langle r| \otimes \Pi_r^c + \sum_{p \neq r} |p\rangle\langle p| \otimes \mathbf{R}_p \Pi_p^c \mathbf{R}_p^\dagger + \sum_{v \notin \mathcal{P}} |v\rangle\langle v| \otimes \mathbb{1}.$$

Noting that the kernel of $\Delta(G) \otimes \mathbb{1}$ is spanned by $\{|\Psi_V\rangle \otimes |\phi\rangle : |\phi\rangle \in \mathcal{H}\}$, we get

$$\begin{aligned} \cos \theta &= \max_{|\phi\rangle} \langle \Phi_V | \langle \phi | \mathbf{W}^\dagger \Pi_{\mathbf{A}} \mathbf{W} | \Phi_V \rangle | \phi \rangle \\ &= \frac{1}{|V|} \max_{|\phi\rangle} \sum_{|v\rangle, |v'\rangle} \langle v | \langle \phi | \mathbf{W}^\dagger \Pi_{\mathbf{A}} \mathbf{W} | \Phi_V \rangle | v' \rangle | \phi \rangle \\ &= \frac{|V| - |P|}{|V|} + \frac{1}{|V|} \max_{|\phi\rangle} \langle \phi | \Pi_r^c + \sum_{p \neq r} \mathbf{R}_p \Pi_p^c \mathbf{R}_p^\dagger | \phi \rangle \\ &\leq \frac{|V| - |P|}{|V|} + \frac{|P| - 1}{|V|} (1 + \cos \vartheta) \leq 1 - \frac{1}{|V|} (1 - \cos \vartheta), \end{aligned}$$

where we used Eq. (5) in Lemma 29 in the last line with a bound on the angle between subspaces $\cos \vartheta = \max_{p \neq r} \angle(\text{supp } \Pi_r^c, \text{supp } \mathbf{R}_p \Pi_p^c \mathbf{R}_p^\dagger)$ and $|P| \geq 2$. We can bound this further by

$$\begin{aligned} \cos^2 \vartheta &= \max_{p \neq r} \max_{\substack{|\xi\rangle \in \text{supp } \Pi_r^c \\ |\eta\rangle \in \text{supp } \mathbf{R}_p \Pi_p^c \mathbf{R}_p^\dagger}} |\langle \xi | \eta \rangle|^2 = \max_{p \neq r} \max_{|\xi\rangle, |\eta\rangle} |\langle \xi | \Pi_r^c \mathbf{R}_p \Pi_p^c \mathbf{R}_p^\dagger | \eta \rangle|^2 \\ &\equiv \max_{p \neq r} \max_{|\xi\rangle, |\eta\rangle} |\langle \xi | \Pi_r^c \mathbf{R}_p \Pi_p^c | \eta \rangle|^2 \leq \max_{p \neq r} |\lambda_{\max}(\Pi_r^c \mathbf{R}_p \Pi_p^c)|^2 =: \lambda_{\max}^2. \end{aligned}$$

The rest follows Lemma 30:

$$2 \sin^2 \frac{\theta}{2} = 1 - \cos \theta \geq 1 - 1 + \frac{1}{|V|} (1 - \lambda_{\max}) = \frac{1 - \lambda_{\max}}{|V|},$$

and the claim follows. □

6. Quantum Thue Systems

6.1. Thue Systems

Let us briefly recall the idea behind classical Thue systems, also known as string rewriting systems.

Definition 45. A *Thue system*—TS for short—is a tuple (Σ, R) of a finite alphabet Σ and a finite symmetric binary relation $R \subset \Sigma^* \times \Sigma^*$, where $\Sigma^* := \bigcup_{i=0}^\infty \Sigma^i$ denotes the set of all *strings* over the alphabet Σ .

The binary relation R is usually written as a set of rewrite rules $a \leftrightarrow b$ and they are naturally extended to other strings in Σ^* : if $a \leftrightarrow b \in R$, then $c \leftrightarrow d$ in R if there exist $u, v \in \Sigma^*$ such that $c = uav$ and $d = ubv$. Thue systems with this extension—denoted by R^* —are a special case of abstract reduction systems and well-studied as computational models—see [29].




FIGURE 11. Undirected graph associated to strings of length 3 for the Thue system in Example 47

Thue systems are multiway systems, i.e. starting from an initial string $a \in \Sigma^*$, exactly one substring is replaced at a time—in particular, this means that there might be branching points when the substitution is not unique. For our purpose, it is enough to consider length-preserving substitutions, i.e. any space Σ^i should be invariant under R ; we denote the length of any string s with $|s|$. In this case, there exists a natural representation of the Thue system over strings of length N as a finite, undirected and not necessarily connected graph.

Definition 46. Let (Σ, R) be a TS and $N \in \mathbb{N}$. The *associated graph* $G = (V, E)$ for strings of length N has vertices $V := \Sigma^N$ and edges $E := \{(a, b) : a \leftrightarrow b \in R^*\}$. The *Laplacian* of the TS is defined as the discrete Laplacian of the associated graph G . For brevity we just write $G = (\Sigma, R)$.

Example 47. Take the alphabet $\Sigma := \{a, b, c\}$ and $R := \{c \leftrightarrow b, ab \leftrightarrow cc\}$. For example, starting from string aab , we can obtain a chain $aab \leftrightarrow acc \leftrightarrow acb \leftrightarrow abb \leftrightarrow ccb \leftrightarrow bcb \leftrightarrow bbb$. The entire graph for strings of length 3 for this example is shown in Fig. 11.

Definition 48. Let (Σ, R) be a TS. We call a non-empty subset $U \subseteq \Sigma^*$ a *valid evolution* if U is closed under the transition rules R^* . We call U *irreducible* if there exists no valid evolution $U' \subsetneq U$.

Example 49. In Example 47, the set $\{caa, baa, aaa\}$ is a valid evolution, but it is not irreducible. The only irreducible evolutions for strings of length 3 in this example are the sets formed from the connected components.

It is immediate to see this one-to-one correspondence between connected associated graph components and irreducible evolutions.

We want to introduce a sense of locality to TS relations.

Definition 50. A TS (Σ, R) is k -local, where $k := \min\{i : \exists r \in R : r \in \Sigma^i \times \Sigma^i\}$.

Observe how this definition is well-defined, as we required R to be finite, cf. Definition 45.

6.2. Quantum Thue Systems and Their Hamiltonian

We begin by generalizing the notion of Thue systems to the case where our alphabet has special quantum symbols, with rewriting rules being unitary operators between them. To work with these two alphabets, consider $\Sigma = \Sigma_{cl} \sqcup \Sigma_q$ as the union of two disjoint—classical and quantum—alphabets. For a string $s \in \Sigma^*$, write $|s|_q$ for the number of letters from Σ_q in s . This allows the following definition.

Definition 51. A *quantum Thue system* (QTS) is a quadruple $(\Sigma, R, \{\mathbf{U}_r\}_{r \in R}, \mathcal{H})$ of a bipartite alphabet $\Sigma = \Sigma_q \sqcup \Sigma_{cl}$, a relation R , a unitary operator \mathbf{U}_r for each rule $r \in R$ and a finite-dimensional Hilbert space \mathcal{H} with the following properties:

- (Σ, R) is a TS,
- $|\cdot|_q$ is invariant under any rule $r \in R$,
- $\mathbf{U}_r \in U(\mathcal{H}^{\otimes |r|_q})$ for all $r \in R$.

The invariance of $|\cdot|_q$ under a rule $r = s_1 \leftrightarrow s_2$ allows to abbreviate $|r|_q := |s_1|_q = |s_2|_q$, which indicates the number of quantum letters the rule r acts on.

We can again use the QTS to form sequences of strings: starting from a string $s \in \Sigma^*$, apply rules consecutively as for TSSs. In addition, to each string $s \in \Sigma^*$, we attach a Hilbert space $\mathcal{H}_s := \mathcal{H}^{\otimes |s|_q}$: starting from some vector $v \in \mathcal{H}_s$, each time a rule r is applied to a substring, the corresponding unitary acts on the subspace wherever the rule matches, acting as identity everywhere else.

Analogous to Fig. 11, we can build a graph for strings of length L for any QTS—where each edge is labelled by the acting unitary. The following lemma should therefore not come as a surprise.

Lemma 52. *Any k -local QTS $(\Sigma, R, \{\mathbf{U}_r\}, \mathcal{H})$ restricted to strings of a certain length $N \geq k$ is also a ULG. Furthermore, the associated Hamiltonian for strings of length N is isomorphic to a geometrically k -local and translationally invariant Hamiltonian on a chain $(\mathbb{C}^\Sigma \otimes \mathcal{H})^{\otimes N}$ with the same spectrum.*

Proof. We explicitly define the ULG (S, R') for strings of length $N \geq k$. The vertex set $S := \Sigma^N$ is straightforward. For every $r \in R$ denoted $s_1 \leftrightarrow s_2$, define the ULG edges $(us_1v \leftrightarrow us_2v, \mathbf{1}_{\mathcal{H}}^{\otimes |u|} \otimes \mathbf{U}_r \otimes \mathbf{1}_{\mathcal{H}}^{\otimes |v|})$ for any $u, v \in \Sigma^*$ —potentially extending \mathbf{U}_r to $\mathcal{H}^{\otimes N-k}$ acting trivially on classical substrings—such that $us_1v \in \Sigma^N$. It is straightforward to verify that this defines a valid ULG.

The second claim follows from the canonical isomorphism between the two Hilbert spaces $(\mathbb{C}^\Sigma)^{\otimes N} \otimes \mathcal{H}^{\otimes N} \xrightarrow{\sim} (\mathbb{C}^\Sigma \otimes \mathcal{H})^{\otimes N}$, i.e. a simple rearrangement. Conjugating the associated Hamiltonian of the ULG with this isomorphism proves the second claim. □

As QTSs are also ULGs, we will be—without always specifying the string length restriction explicitly—using ULG terminology for QTSs, e.g. associated Hamiltonians, irreducible evolutions or speak of QTSs being simple.

Lemma 53. *Let the setup be as in Lemma 52. Then the isomorphism extends to a Hamiltonian on the chain $(\mathbb{C}^{\Sigma_{cl}} \oplus (\mathbb{C}^{\Sigma_q} \otimes \mathcal{H}))^{\otimes N}$ with the same spectrum up to multiplicities.*

Proof. Any rule on a classical substring acts identically on \mathcal{H} , hence on the set of Hamiltonians with this property a conjugation of the isomorphic Hamiltonian in Lemma 52 **H** with the projector $(\mathbb{C}^{\Sigma} \otimes \mathcal{H})^{\otimes N} \hookrightarrow (\mathbb{C}^{\Sigma_{cl}} \oplus (\mathbb{C}^{\Sigma_q} \otimes \mathcal{H}))^{\otimes N}$ is an isomorphism. The projector preserves the spectrum, up to multiplicities. \square

In the following, we often gloss over the fact and simply assume that the associated Hamiltonian of a QTS is local in the sense of Lemma 53. Observe however that a ULG induced from a QTS is not necessarily simple, and it is easy to find a counterexample.

6.3. Quantum Thue Systems as a Computational Model

To use a QTS for computation, we need to mark some strings that have special meaning, e.g. are input or output of the computation.

Definition 54. For a QTS with alphabet Σ and Hilbert space \mathcal{H} , a *marker* is any tuple (s, π) where $s \in \Sigma^*$ and π is a projector on some subspace of $\mathcal{H}^{\otimes |s|_q}$. The set of markers on strings of length k —called k -local markers—is denoted $\mathcal{M}^{(k)}$, and $\mathcal{M} := \bigcup_k \mathcal{M}^{(k)}$.

That is, we can specify a string s and a configuration of the quantum part of this string as a specific state in the computation. It is useful to think of using one marker s_{inp} to mark a string as the start of the computation, and a second one s_{out} to mark the end; the quantum parts of the markers— Π_{inp} and Π_{out} —then define the valid input and output of the computation.

Definition 55. Let $(\Sigma, R, \{\mathbf{U}_r\}, \mathcal{H})$ be a QTS and $\Pi = (\Pi_{\text{YES}}, \Pi_{\text{NO}})$ a promise problem. We introduce an *encoding function* $\text{enc} : \Pi \rightarrow \Sigma^*$, input and output markers $(s_{\text{inp}}, \Pi_{\text{inp}}), (s_{\text{out}}, \Pi_{\text{out}}) \in \mathcal{M}^{(n)}$ for some $n \in \mathbb{N}$. Then the QTS

- *rejects* an instance $l \in \Pi$ if there exists a chain of rules in R connecting $\text{enc}(l)$ with two strings containing s_{inp} and s_{out} , respectively, and $\langle \psi | \Pi_{\text{inp}} + \mathbf{U}^\dagger \Pi_{\text{out}} \mathbf{U} | \psi \rangle \geq \epsilon$ for all $|\psi\rangle \in \mathcal{H}$ —here $\mathbf{U} = \mathbf{U}(l)$ stands for the product of unitaries along this chain, and Π_{inp} and Π_{out} are extended trivially to the entire chain in case $|s_{\text{inp}}|_q > |\text{enc}(l)|_q$ or $|s_{\text{out}}|_q > |\text{enc}(l)|_q$.
- *accepts* l if there exists a $|\psi\rangle \in \mathcal{H}$ such that $\langle \psi | \Pi_{\text{inp}} + \mathbf{U}^\dagger \Pi_{\text{out}} \mathbf{U} | \psi \rangle \leq \epsilon/2$.
- *decides* Π if for all $l \in \Pi$, l is accepted if $l \in \Pi_{\text{YES}}$, and rejected for $l \in \Pi_{\text{NO}}$.

The rejection and acceptance threshold ϵ will depend on the class of promise problems that we want to decide. In particular, we want to allow this threshold to scale with the problem instance size, i.e. $\epsilon = \epsilon(|l|)$, and

thus indirectly with the time that a computation can take, as specified in the following definition.

Definition 56. Let Q be a QTS that decides Π . For an instance $l \in \Pi$, the *history state* is defined as the irreducible evolution of the ULG containing $\text{enc}(l)$.

For a QTS with unambiguous transition rules—i.e. where the history state is a line—the size of the history state simply corresponds to the runtime of the underlying computation.

We now want to describe a simple example for a QTS which can decide the following simple promise problem.

Definition 57 (*Even Natural Number*).

Instance. Natural number $n \in \mathbb{N}$.

Output. YES if n even, otherwise NO.

Example 58. Let the alphabet $\Sigma := \{-, \star, \|\}$, where \star is the only quantum symbol with Hilbert space \mathbb{C}^2 . We define $s_{\text{out}} = \star\|$ and $\Pi_{\text{out}} = |1\rangle\langle 1|$. Let further

$$\text{enc} : \mathbb{N} \rightarrow \Sigma^* \quad \text{where} \quad \text{enc}(n) := \underbrace{\star - - \dots -}_{n \text{ times}} \|\.$$

$s_{\text{inp}} = \text{enc}(l)$, $s_{\text{out}} = \star\|$, and $\Pi_{\text{inp}} = \Pi_{\text{out}} = |1\rangle\langle 1|$. We have a single rule $(\star - \leftrightarrow - \star, \mathbf{R})$ where \mathbf{R} is a rotation by $\pi/2$, i.e. $\mathbf{R} := -|1\rangle\langle 0| + |0\rangle\langle 1|$. Then this QTS decides EVEN NATURAL NUMBER.

Proof. The proof is straightforward. Starting on the encoded input $\text{enc}(n)$, the TS generates a sequence

$$\star - - \dots - \|\ \mapsto \ - \star - - \dots - \|\ \mapsto \ \dots \mapsto \ - - \dots - \star\|,$$

so there always exists a chain of rules that connects $\text{enc}(n)$ with a string containing s_{out} . The decision is thus made by the content of the quantum part: for n applications of the rule, starting from a vector $|v\rangle \in \mathbb{C}^2$, we apply \mathbf{R} n times. Now take any state $|\psi\rangle \in \mathbb{C}^2$ and write $|\psi\rangle = a|0\rangle + b|1\rangle$. Then

$$\begin{aligned} &\langle 0| \left(\Pi_{\text{inp}} + (\mathbf{R}^\dagger)^n \Pi_{\text{out}} \mathbf{R}^n \right) |0\rangle \\ &= |\langle 0| \mathbf{R}^n |1\rangle|^2 = |\langle 0| \mathbf{R}^{n \bmod 2} |1\rangle|^2 = \begin{cases} 0 & \text{if } n \text{ is odd,} \\ 1 & \text{otherwise,} \end{cases} \end{aligned}$$

and

$$\begin{aligned} &\langle 1| \left(\Pi_{\text{inp}} + (\mathbf{R}^\dagger)^n \Pi_{\text{out}} \mathbf{R}^n \right) |1\rangle \\ &= 1 + |\langle 1| \mathbf{R}^n |1\rangle|^2 \geq 1. \end{aligned}$$

Therefore, if n is even,

$$\langle \psi| \left(\Pi_{\text{inp}} + (\mathbf{R}^\dagger)^n \Pi_{\text{out}} \mathbf{R}^n \right) |\psi\rangle \geq |a|^2 + |b|^2 = 1$$

and $|0\rangle$ is an accepting state for odd n . The claim follows. □

7. Hardness Results

7.1. A Special Kind of Quantum Thue System

We have seen that QTS can be used to answer simple problems. On the other hand, a more interesting question is whether there exists a universal QTS which can run any computation of a certain class of promise problems \mathcal{C} , i.e. is complete for \mathcal{C} . Of particular interest in this setting is the question about scaling of the defining parameters for such a QTS: how big is the alphabet, what is the locality and how does the string length of the vertices in the history state scale, i.e. for a promise problem $\Pi \in \mathcal{C}$, does there exist a function f such that for $l \in \Pi$, $|\text{enc}(l)| = O(f(|l|))$? And what about the size of the history state?

For the complexity class BQEXP, we have the following lemma.

Lemma 59. *For any BQEXP promise problem Π , there exists a 2-local QTS $(\Sigma, R, \{\mathbf{U}_r\}, \mathbb{C}^2)$ which decides Π , and has the following uniform properties:*

- W1** *The alphabet has special characters $H \subset \Sigma$ —heads—and $B \subset \Sigma$ —boundaries, and a set of allowed pairs $\mathfrak{A} \subset \Sigma \times \Sigma$. All transition rules preserve any symbols from B and the number of symbols in H (denoted $|s|_h$ for a string s).*
- W2** *Let $j : \Pi \rightarrow \mathbb{N}$ be a map with $j(l) = O(\text{exp poly } |l|)$, where $|l|$ denotes the size of instance $l \in \Pi$. The QTS decides instance l on strings of this length, i.e. $|\text{enc}(l)| = j(l)$. Both input and output penalty are 2-local markers containing precisely one head symbol, and $s_{\text{inp}}, s_{\text{out}} \in \mathfrak{A}$.*
- W3** *For any $l \in \Pi$, the history state \mathcal{M}_l is simple. All strings $s \in \mathcal{M}_l$ are of the form $s \in B \times (\Sigma \setminus B)^* \times B$ (bracketed), and have one head $|s|_h = 1$. Furthermore, all length-2 substrings of s are in \mathfrak{A} , and the size $|\mathcal{M}_l| = \text{poly}(j(l))$.*
- W4** *For all other irreducible evolutions $\mathcal{M} \neq \mathcal{M}_l$, at least one of the following is true:*
 - $|s|_h = 0 \ \forall s \in \mathcal{M}$,
 - \mathcal{M} is not bracketed (i.e. with a boundary symbol on the left and right ends),
 - \mathcal{M} can be broken up into $O(g(l))$ -sized connected parts—where $g(l) = \text{poly}(j(l))$ —each of which containing at least one string containing an invalid character tuple not in \mathfrak{A} .

7.2. Hardness Result

Theorem 60. $(2, 42)$ -HAMILTONIAN is QMA_{EXP} -hard.

Proof. Following Definitions 13 and 16, we need to show that there exists a 1D translationally invariant 2-local Hamiltonian \mathbf{H} on $\mathcal{H} = (\mathbb{C}^{42})^{\otimes n}$ with $O(1)$ local terms, such that either (a) $\lambda_{\min}(\mathbf{H}) \leq \alpha$ or (b) $\lambda_{\min}(\mathbf{H}) \geq \beta$ with a polynomial promise gap $\beta - \alpha = \Omega(1/\text{poly } n)$, and deciding between (a) and (b) is at least as hard as some QMA_{EXP} -hard promise problem. The proof will be a simple combination of our previously collected results.

1. Let Π be a promise problem in QMA_{EXP} . By Definition 13, the verification of Π is a BQEXP problem. By Fact 12, we can assume without loss of generality that the accept and reject probabilities in Definition 55 are $1 - \epsilon$ and ϵ , respectively, where $\epsilon = 1/3^{\text{poly}(|l|)}$ to be specified below, where $|l|$ denotes the length of the problem input.
2. By Corollary 72, we can thus create a QTS with properties as in Lemma 59 that verifies Π : more specifically, for an instance $l \in \Pi$ and by item **W2**, we know that this QTS verifies l on strings of length $j(l)$.
3. By Lemma 52, the QTS restricted to strings of length $j(l)$ is also a ULG. Denote the Hamiltonian associated to this ULG by \mathbf{H}_l , block-diagonal in the irreducible evolutions.

With Γ denoting the alphabet from Definition 62 and $j(l)$ denoting the number of systems, we define a Hamiltonian on the Hilbert space $(\mathbb{C}^\Gamma)^{\otimes j(l)}$ as follows:

$$\mathbf{H} := \mathbf{H}_l + \mathbf{B}_{\text{heads}} + p(l)(\mathbf{P}_{\text{boundaries}} + \mathbf{P}) + \mathbf{P}_{\text{in/out}}, \tag{12}$$

where

- $\mathbf{P}_{\text{boundaries}}$ penalizes any non-bracketed string (i.e. strings without a boundary symbol on at least one end)—cf. item **W3**,
- $\mathbf{B}_{\text{heads}}$ acting on a string $|s\rangle \in (\mathbb{C}^\Gamma)^{\otimes j(l)}$ gives a bonus of $|s|_h$, according to how many head symbols there are in s ,
- \mathbf{P} penalizes any character tuple not in \mathfrak{A} ,
- $p(l)$ is a function used to scale the penalties, which will be specified later, but—keeping Remark 18 in mind—must not exceed $p(l) = \text{poly } j(l)$.

Penalizing non-bracketed strings follows an idea by [3]. With $\mathbf{P}_{\text{boundaries}}$, we give a 1-local bonus of size 1 to brackets appearing anywhere, but a penalty of $1/2$ to them appearing next to any other symbol; since no transition rule ever moves the boundaries, this gives a uniform energy shift to all strings with brackets. The unique highest-bonus string will have a bracket appearing at the start and end with a bonus of size 1.

The encoding and output penalties $\Pi_{\text{inp}} = \Pi_{\text{out}} = |1\rangle\langle 1|$ are translationally invariantly extended to the entire chain, i.e. on Hilbert space $(\mathcal{H}_{cl} \oplus \mathcal{H}_q)^{\otimes j(l)}$, we act with the 2-local projector

$$\mathbf{P}_{\text{in}} := \sum_{i=1}^{j(l)-1} (|s_{\text{inp}}\rangle\langle s_{\text{inp}}| \oplus \Pi_{\text{inp}})_{i,i+1},$$

and analogously for \mathbf{P}_{out} .

Completeness.

Assume l is a YES-instance, and denote the history state as an eigenvector of \mathbf{H}_l with $|\Psi_l\rangle$, which by item **W3** is also an eigenstate of $\mathbf{B}_{\text{heads}}$, $\mathbf{P}_{\text{boundaries}}$ and \mathbf{P} . A direct calculation yields

$$\begin{aligned} \langle \Psi_l | \mathbf{H} | \Psi_l \rangle &= \overbrace{\langle \Psi_l | \mathbf{H}_l | \Psi_l \rangle}^{=0} + \overbrace{\langle \Psi_l | \mathbf{B}_{\text{heads}} | \Psi_l \rangle}^{=-1} + p(l) \overbrace{(\langle \Psi_l | \mathbf{P}_{\text{boundaries}} | \Psi_l \rangle)}^{=-1} \\ &\quad + \overbrace{\langle \Psi_l | \mathbf{P} | \Psi_l \rangle}^{=0} + \langle \Psi_l | \mathbf{P}_{\text{in/out}} | \Psi_l \rangle \\ &= -2 + \langle \Psi_l | \mathbf{P}_{\text{in/out}} | \Psi_l \rangle. \end{aligned}$$

By item **W2**, we further know that at least one vertex in $|\Psi_l\rangle$ has the in- and output substrings s_{inp} , s_{out} , and because $|s_{\text{inp}}|_h = |s_{\text{out}}|_h = 1$, there is at most one such substring match for every vertex. As an upper bound, we can thus assume that the penalty applies exactly once in every vertex—i.e. $\langle \Psi_l | \mathbf{P}_{\text{in/out}} | \Psi_l \rangle \leq \epsilon$, and conclude $\langle \Psi_l | \mathbf{H} | \Psi_l \rangle \leq -2 + \epsilon$.

Soundness. Assume l is a **NO**-instance. We need to lower-bound the lowest energy eigenvalue of \mathbf{H} , and since we know that \mathbf{H} is block-diagonal in the irreducible evolutions, we can bound each block separately—the history state block given in item **W3** and any other irreducible evolution block characterized by item **W4**. Without loss of generality, we can therefore assume that $|\psi\rangle$ is completely supported on a single block of \mathbf{H} (but not necessarily an eigenvector).

Take any $|\psi\rangle$ with support constrained to the history state block. As in the completeness part, a direct calculation allows the estimate

$$\begin{aligned} \langle \psi | \mathbf{H} | \psi \rangle &= \langle \psi | \mathbf{H}_l | \psi \rangle + \overbrace{\langle \psi | \mathbf{B}_{\text{heads}} | \psi \rangle}^{=-1} + p(l) \overbrace{(\langle \psi | \mathbf{P}_{\text{boundaries}} | \psi \rangle)}^{\geq -1} + \overbrace{\langle \psi | \mathbf{P} | \psi \rangle}^{\geq 0} \\ &\quad + \langle \psi | \mathbf{P}_{\text{in/out}} | \psi \rangle \\ &\geq -2 + \langle \psi | \mathbf{H}_l + \mathbf{P}_{\text{in/out}} | \psi \rangle. \end{aligned}$$

We can now apply Lemma 44 to the last expression. By item **W2** and Definition 55, we obtain a bound $\langle \psi | \mathbf{H}_l + \mathbf{P}_{\text{in/out}} | \psi \rangle \geq (1 - \epsilon)/|\mathcal{M}_l|^3$. Observe how this lower bound scales $\propto 1/|\mathcal{M}_l|^3$, whereas for **YES**-instances the upper bound scales constant $\propto \epsilon$. Since we want the lower bound for **NO**-instances— β —and the upper bound for **YES**-instances— α —to be separated by at least some $\beta - \alpha = \Omega(1/\text{poly } j(l))$, cf. Definition 16, we need to amplify the accepting probability to $\epsilon = O(1/|\mathcal{M}_l|^4) = O(1/\text{poly } j(l)^4)$. Observe that this does not exceed the allowed amplification, which is only limited to $O(1/3^{\text{poly } |l|})$.

We proceed to show lower bounds for all other minimum valid evolutions, following item **W4**. Assume we are in a block with 0 heads, which is well-defined by item **W1**. The bonus term $\mathbf{B}_{\text{heads}}$ vanishes on this subspace while all other operators in Eq. (12) are positive semi-definite, so we obtain a lower bound of $\langle \psi | \mathbf{H} | \psi \rangle \geq 0$ for any state solely supported there.

Analogously, non-bracketed blocks can be bounded by a direct calculation, as $\mathbf{P}_{\text{boundaries}}$ penalizes all vertices equally: any non-bracketed state $|\psi\rangle$ for a block with h heads satisfies $\langle \psi | \mathbf{H} | \psi \rangle \geq -h + p(l)$. It thus suffices to set $p(l) \geq j(l)$, as the number of possible heads on a string is limited by its length, i.e. $h \leq j(l)$.

The last blocks remaining are the ones with $g(j(l))$ -sized connected parts with invalid tuples, where $g(n) = \text{poly } n$ as defined in item **W4**. First observe

that this part of the ULG is not necessarily simple, so we remove the transitions which allow non-trivial loops without breaking the graph up into multiple parts. We *then* split this graph into $g(j(l))$ -sized connected components by temporarily removing further edges from it, which yields a Hamiltonian for a sparser graph \mathbf{H}' . Since adding any edges back in corresponds to adding a positive semi-definite matrix to \mathbf{H}' , it suffices to lower-bound the spectrum of \mathbf{H}' on this subspace. Note that we do not remove vertices or change any penalties, so in particular all the diagonal operators in Eq. (12) remain untouched.

Hence, assume $|\psi\rangle$ has support in one of the slices of size upper-bounded by $g(j(l))$ with h heads, such that at least one vertex picks up a penalty from \mathbf{P} . Again applying Lemma 44, we obtain a bound $\langle\psi|\mathbf{H}|\psi\rangle \geq \langle\psi|\mathbf{H}'|\psi\rangle \geq -h - 1 + p(l) \times \Omega(1/g(j(l))^3)$. We therefore have to scale p to e.g. $p(l) \geq g(j(l))^5$, which is still allowed by Remark 18 (namely, p is polynomial in l). This concludes the proof. \square

What remains to be shown is the existence of a QTS as in Lemma 59. The next section will provide an explicit construction, finalizing the proof of our main result. This construction is meant as a proof-of-concept—the model we present can be modified in numerous ways and is likely not optimal. It does, however, make heavy use of our newly developed methods such as branching, thus reducing the local dimension of the underlying Hamiltonian to 42, as compared to the hitherto best result by [3] which is larger by at least several orders of magnitude.

8. Turing's Wheelbarrow

8.1. Introduction

Turing's Wheelbarrow is our constructive proof of a QTS with properties as mentioned in Lemma 59. The QTS will be optimized for local dimension and locality—every transition rule will be 2-local and act on strings from an alphabet Γ with 48 characters. We describe the QTS by explicitly writing out all transition rules of the QTS and then prove the properties from Lemma 59. Finally, in Sect. 8.7, we reduce its local dimension down to 42.

The conceptual idea of the Wheelbarrow QTS is the following. To build a QTS which can decide a promise problem $\Pi \in \text{BQEXP}$, we first prefix the original circuit C_l deciding an instance $l \in \Pi$ by another circuit which verifies that a number of ancillas necessary for C_l are correctly initialized to $|0\rangle$. On some extra ancillas, we write out the problem instance l , and also leave an unconstrained section of qubits available for C_l . This witness section, problem instance and the leftover ancillas are then fed into C_l , and the output wire contains $|\text{out}\rangle = \cos((p_a + p_{\text{out}})/3)|0\rangle + \sin((p_a + p_{\text{out}})/3)|1\rangle$ for the amplitudes p_a —all ancillas being 0—and p_{out} —the circuit output of C_l on the ancillas and problem instance. This overall circuit, denoted C'_l , is shown in Fig. 12.

It is clear that this augmented circuit family $(C'_l)_{l \in \Pi}$ is in the same uniformity class as the original circuit family $(C_l)_{l \in \Pi}$, and we can thus define these circuits with output $|\text{out}\rangle$ to be a separate BQEXP problem Π' . By Lemma 33




FIGURE 12. An augmented quantum verifier circuit. The circuit uses one ancilla $|0\rangle$ to verify that as many ancillas as necessary for the computation are set to 0, rotating the single guaranteed $|0\rangle$ ancilla by $\pi/3$ if this is not the case. On some ancillas, the problem instance l is written out. Another rotation by $\pi/3$ is applied depending on the output of the verifier circuit. The overall output state then takes the form $|\text{out}\rangle = \cos((p_a + p_{\text{out}})/3) |0\rangle + \sin((p_a + p_{\text{out}})/3) |1\rangle$

and its proof, this new promise problem can be decided by a family of BQEXP QRMs with the special property that the head motion and all internal QRM states are classical—cf. Fig. 7.

Using the *Solovay-Kitaev theorem* [37, appx. 3], the head unitary of such a QRM can in turn be efficiently rewritten as a circuit R using the following small set of gates.

Remark 61. TOFFOLI, SWAP and a classically controlled quantum-universal unitary together with at least one classical and quantum ancilla is universal for quantum computation and exactly universal for classical computation.

In particular, S-K tells us that since the head circuit $R = R(l)$ depends on the problem instance l —as it needs to write the instance out—and the size of this circuit is $|R(l)| = \text{poly } |l|$. The Wheelbarrow QTS which we construct will then be able to execute this head cyclically on a ring of qubits, where the execution is halted once the QRM terminates: as the QRM motion is deterministic, the runtime will be $\text{exp poly } |l|$ -bounded, as required for a BQEXP computation.

The first step is to bootstrap the QRM head \mathbf{U} . Starting from an initially empty string, we use a number of rules to translate the string length N into a circuit description of \mathbf{U} on the left side of the string. This section will have size $\approx \log_6 N$, as we need 6 instruction symbols—a classically controlled unitary $\bar{\mathbf{U}}$, TOFFOLI \mathbf{T} , ancilla-checking symbol \mathbf{A} , swap \mathbf{S} , left-shift symbol \diamond and halt \mathbf{H} . The remaining right side of the string will act as classical and quantum tape that the computation runs on. Figure 13 outlines how a circuit can be translated into such a 6-ary circuit description.

The QRM is then executed: for every round, a program bit is taken from the left side of the string, moved towards the tape and then applied




FIGURE 13. Example on how to translate a sample circuit section into a program description, where U_0 , U_1 and $U_?$ can stand for any unitary gate. Starting from the top left, the description here is $1\ 0\ 0\ 1\ \diamond\ 0\ 0\ \diamond\ 0\ 0\ \diamond\ \diamond\ \diamond\ \diamond\ \diamond\ \diamond\ 1\ \diamond\ ?$. The dashed line stands for a normal identity $0\ \diamond\ 0$, as $CNOT^2 = \mathbb{1}$. The rhombus \diamond is a special symbol that shifts the current gate position up by one; as in each successive step the position moves down by one by default, it suffices to only have this one special shift symbol. The last gate is a special identity to be used to initialize ancillas and penalize a section of the output




FIGURE 14. The two actions that can be performed by the wheelbarrow construction. On the left, we apply a gate U_0 corresponding to the rightmost program bit 0 . The ring of qubits is then rotated by one, which is the default downwards shift as mentioned in Fig. 13. On the right, the special action of the \diamond symbol is depicted: it signals the rightmost qubit to move back to the left side. After either action, the program string is rotated by one

to the leftmost two data qubits. The leftmost data qubit is then picked up and carried to the right, where it is deposited. The revert action is similar, only that the rightmost data qubit is picked up and moved to the left of the tape. Figure 14 illustrates both operations. The execution runs until the underlying ring machine terminates, which can be determined using a special halt operation H which only proceeds if the tape data is not in a halting configuration.

This also explains the choice of *Turing's Wheelbarrow* as name for this QTS: qubits and program symbols are moving across the tape in two cyclic

motions, mimicking a busy worker carrying and depositing information in a wheelbarrow.

8.2. Notation

For convenience, we define a special notation to describe the construction of Turing’s Wheelbarrow. We begin by introducing the alphabet and tape.

Definition 62. Let $\Gamma := \Gamma_{cl} \sqcup \Gamma_q$ denote the alphabet consisting of 48 symbols where

$$\Gamma_{cl} := \left\{ \begin{array}{l} \parallel, \blacksquare, \boxtimes, \boxplus, \curvearrowright, \curvearrowleft, \curvearrowup, \curvearrowdown, \blacktriangleright, \blacktriangleleft, \blacktriangleright, \blacktriangleleft, \blacktriangleleft, \blacktriangleright, \\ \mathbf{U}, \mathbf{T}, \mathbf{A}, \mathbf{S}, \mathbf{H}, \diamond, \mathbf{U}, \mathbf{T}, \mathbf{A}, \mathbf{S}, \mathbf{H}, \blacklozenge, \square, \mathbf{U}, \mathbf{T}, \mathbf{A}, \mathbf{S}, \mathbf{H}, \blacklozenge, \blacklozenge, \\ \mathbf{!}, \vec{0}, \vec{1}, \dot{1}, \ddot{1}, 0, 1, \mathbf{0}, \mathbf{1} \end{array} \right\},$$

$$\Gamma_q := \left\{ \begin{array}{l} \mathbf{a}, \mathbf{a}, \mathbf{a}, \mathbf{a}, \mathbf{a} \end{array} \right\}.$$

These two sets correspond to the classical and quantum symbols, respectively, and are of size $|\Gamma_{cl}| = 43$ and $|\Gamma_q| = 5$. The set of head characters is

$$\mathbf{H} := \Gamma \setminus \{ \parallel, \blacksquare, 0, 1, \mathbf{U}, \mathbf{T}, \mathbf{A}, \mathbf{S}, \mathbf{H}, \diamond, \square, \blacklozenge, \blacklozenge, \blacklozenge, \mathbf{a} \},$$

and the boundary characters are $\mathbf{B} := \{ \parallel, \boxtimes, \blacksquare \}$.

The number of alphabet characters can be further reduced to 39, 3 of which are quantum, see Corollary 72. For reasons of clarity, we use a slightly larger alphabet in this construction.

We will generally use the letters x, y, z as placeholders for program symbols—denoting any of the symbols $\mathbf{U}, \mathbf{T}, \mathbf{A}, \mathbf{S}, \mathbf{H}, \diamond$ as x , or alternatively $\mathbf{U}, \mathbf{T}, \mathbf{A}, \mathbf{S}, \mathbf{H}, \blacklozenge$ as \mathbf{x} , which is always clear from the context. The symbol \mathbf{U} encodes a classically controlled unitary, \mathbf{T} a TOFFOLI, \mathbf{A} an ancilla, \mathbf{S} a SWAP, \mathbf{H} a halt and \diamond a special tape revert symbol.

We now introduce the notation for transition rules.

Definition 63. We write a transition rule $xy \leftrightarrow zw$ of a quantum Thue system as

$$\begin{array}{c} \mathbf{x} \mathbf{y} \\ \mathbf{z} \mathbf{w} \end{array} .$$

The blue shading is used to indicate the location on the tape where the transition rule is applied. Note that, by construction, transition rules are symmetric, i.e.

$$\begin{array}{c} \mathbf{x} \mathbf{y} \\ \mathbf{z} \mathbf{w} \end{array} \text{ is equivalent to } \begin{array}{c} \mathbf{z} \mathbf{w} \\ \mathbf{x} \mathbf{y} \end{array} .$$

If the first rule is associated with a non-trivial unitary \mathbf{U} , the inverse rule is associated with the adjoint \mathbf{U}^\dagger .

As in Definition 45, we never need to write out the values of the qubits anywhere. In fact, the only place where the associated Hilbert space comes in is when we want to apply a quantum gate to the qubits (see Sect. 8.3.5). As an example, consider the action of swapping two neighbouring qubits. The Thue system itself does not notice this, e.g. we would have a transition with an explicit comment on the Hilbert space unitary, i.e. $\begin{matrix} a & a \\ a & a \end{matrix}$ where the associated Hilbert spaces are swapped with $U = \text{SWAP}$.

To emphasize that the subspaces are in fact swapped, we generally use the letters a, b, c, d to label different quantum subsystems. *This is only to facilitate notation!* In principle, we could stick to the letter a and write out the swap action for every transition rule where this is relevant. But because we believe it is easier to read and most of the non-trivial unitaries that we use are swaps, we simply write

$$\begin{matrix} a & b \\ b & a \end{matrix}$$

which is self-explanatory.


8.3. Transition Rules and History State

The following table contains a list of all transition rules, visualized from a starting string of the form $\|\exists a c \cdots a c c \ominus\|$. The horizontal direction corresponds to space while the time flows from top to bottom. By default, the unitary associated to any rule—if not mentioned otherwise—is the identity. Apart from SWAP operations, the only non-trivial unitary appears in the computation step in Sect. 8.3.5.





Observe that there are many possible local ambiguities within the history state, which we analyse in detail in Sect. 8.4.



8.3.1. Initialization. The initialization is done by moving a special symbol, a “sweeper” \exists , from one end of the tape to the other side. This ensures that the tape is actually correctly initialized, since any symbol apart from a , c or the ghost \ominus would result in a penalized configuration, cf. Table 2.

Left hand side has a sweeper \exists , right hand side an inactive ghost \ominus , and all middle symbols are qubits a or data bits $c \in \{0, 1\}$, which are opaque for the ghosts. We let the sweeper move through all middle symbols. This allows a dynamic “initialization” of the tape: if the sweeper \exists bumps into any symbol that is not a qubit, data bit or inactive ghost, we can penalize the configuration, singling out the proper history state. Once the sweeper reaches the ghost \ominus at the right boundary, it activates the ghost to \ominus and transitions to the box \square .



8.3.2. Ghost. The ghost symbols act as general “carriage return” symbols: this saves having different return variants for each head symbol used, and is solely a way of saving local dimension. The ghost can thus be seen as a particle to the right side of any other head symbol, and which diffuses freely on the tape (i.e. randomly moves left or right). Only if the ghost is “activated”—i.e. carries a “head flag”—can it interact non-trivially with the symbols around it.

Generally, if there is an extra head symbol on the tape, the ghost is inactive ( and ). The ghost can itself carry the head flag, in which case we call it *active* and denote it by  or . The white active ghost can either turn itself into a head symbol on the left hand side, or activate the boundary. On the right boundary, it oscillates between white and black. This construction saves us a lot of symbols, since we only ever need to specify special right-moving heads, whereas the left movement of the head state is done generically by the ghost. We will often gloss over inactive ghost transitions and assume the ghost just “moves out of the way” as necessary.

Ghosts can change color on the right boundary. Since the black ghosts  and  are static, any incoming head from the left can detect when it has reached the boundary as it will encounter a black ghost.




White ghosts can move through all static symbols, but *not* through heads.




8.3.3. Base-6 Counter. From a high-level perspective, the base-6 counter and unary counter (next section) work together to translate the tape length into a base-6 big endian number on the left side of the tape. This base-6 number then encodes the program which we execute afterwards: we count in base 6 through the sequence **T**, **U**, **A**, **S**, **H**, **◇**—encoding a TOFFOLI, classically controlled unitary, ancilla, swap, halt or tape revert operation, respectively.


Active ghost \mathcal{A} hits the left boundary and turns into the incrementer \blacktriangleright .



If the incrementer encounters \diamond —the highest-valued digit—it flips it to \mathbf{T} —the lowest-valued digit. This results in an overflow that is carried over to the next digit to the right. If the next digit is again \diamond , the same procedure repeats until a different symbol or a qubit is encountered.



If the incrementer \blacktriangleright encounters \mathbf{T} or classical zero $\mathbf{0}$ —both of which are treated as lowest-value symbol—it increments it to the next higher-valued symbol \mathbf{U} . To uniquely distinguish to which symbol to decrement when run in reverse, the incrementer \blacktriangleright has to transition to the checking symbol \blacktriangleright or \blacktriangleright , verifying that the symbol to its right is another counter or tape symbol, respectively. We never encounter the configuration $\blacktriangleright \mathbf{1}$ or $\blacktriangleright a$, as this is penalized.



The incrementation ends with the reverter symbol \blacktriangleleft .


If the incrementer encounters \mathbf{U} , \mathbf{A} , \mathbf{S} or \mathbf{H} , it increments the symbol to the next higher one— \mathbf{A} , \mathbf{S} , \mathbf{H} or \diamond , respectively—turning into the reverter symbol \blacktriangleleft .



The reverter \blacktriangleleft moves through the lowest-valued digits \mathbf{T} all the way to the left boundary where it turns into the right mover \blacktriangleright . Note that no digits other than \mathbf{T} are possible to the left of \blacktriangleleft since incrementation proceeds to the next digit only in case of an overflow.



The right mover \blacktriangleright proceeds to the right through all digits of the base-6 counter.



Eventually \blacktriangleright encounters a qubit a or classical bit c . It turns into an inactive ghost \mathcal{A} and picks up the qubit a or classical bit c . Afterwards the ghost \mathcal{A} moves out of the way and we proceed to the unary counter.



For the configuration $\blacktriangleright \square$ there is no forward transition, which means that once we entered the computation phase, this counting does not continue.

8.3.4. Unary Counter. The unary counter is necessary so that the base-6 counter knows when to stop. We use a block symbol \square to denote the position of the unary counter on the tape, starting from the right and moving to the left at each increment. Whenever this block is moved left once, the base-6 counter has been incremented by one as well. In this way, once the unary counter has run out of space, we have translated the tape length into a base-6 number on the left side of the tape.

A qubit a or classical bit c is carried to the right past all other qubits or classical bits.



The position of \square indicates the value of the unary counter. As the qubit a moves through it, the block \square is pushed one position to the left—this increments the unary counter by one.



Once the moving qubit a reaches the black ghost \blacksquare at the right boundary, the qubit a is dropped and the ghost is activated to \blacklozenge .



8.3.5. Computation. The tapenow has the form $\| \mathbb{H} x \cdots x \square c a a c \cdots a \blacklozenge \|$, i.e. the counting is complete and by our choice of the chain length, the program description starts with a halt symbol \mathbb{H} . The rest of the program string does not contain any \mathbb{H} 's.

The idea behind the computation is depicted in Fig. 4. We first take the base 6 symbol from the left end of the program description and move it to the right end (e.g. $x 2345$ would become $2345x$). This symbol x can then be picked up by the box \square , which becomes activated to \blacksquare . The active box is now followed by a set of rules which applies this program action to the (qu)bits right next to it. Afterwards, the leftmost (qu)bit is carried to the right end and the procedure repeats.

The content of the tape symbols is checked on the fly using the ancilla program bit symbol \mathbb{A} . If it appears next to a qubit, a penalty is given for the qubit marginal being $|1\rangle\langle 1|$; for a classical bit, we penalize 0 (this is because we do not have a NOT gate, but NOT can be implemented with TOFFOLI, which maps $111 \mapsto 110$). The implementation details of all the different program bits are explained in the following table.

The computation halts once a halting program bit \mathbb{H} is next to a classical 1.


The boundary \mathbb{H} is flipped to \mathbb{I} and can only revert to \mathbb{H} next to a halt symbol \mathbb{H} —this ensures that we can only transition back and forth between counting and computation if the program bits are in their original order. An active ghost \mathbb{Q} can hit this left boundary \mathbb{I} and activate it to \mathbb{H} .




A program bit $x \in \{\mathbb{T}, \mathbb{U}, \mathbb{A}, \mathbb{S}, \mathbb{H}, \mathbb{O}\}$ is picked up as \mathbb{D} and carried to the right of the program string.



These transition rules apply the TOFFOLI gate to three classical bits $c c c$. If either $0 c c$ or $1 0 c$, the last bit remains unchanged. Only for the configuration $1 1 c$ we perform a bit flip on the last bit, i.e. $d = -c$. The first bit is then picked up with a carrier \mathbb{D} or \mathbb{D} .




These transition rules apply the classically-controlled unitary operation to a pair of qubits $a a$, but only for the configuration $1 a a$. The control bit is then picked up with a carrier \mathbb{D} or \mathbb{D} . Observe that this is the *only* position where we apply a unitary operation to the quantum symbols.



$\mathbb{A} a$ acts as identity on the qubit, but is used later on to penalize when the attached Hilbert space is $|1\rangle$, giving us the possibility for ancillas. $\mathbb{A} 1$ acts as identity, and we will penalize $\mathbb{A} 0$, giving us a classical ancilla bit. Observe that we choose 1 here, as we can create 0s out of nothing but 1s with the TOFFOLI gate, but not vice versa.




\mathbb{S} implements the SWAP gate.



\mathbb{H} acts as identity on $\mathbb{0}$, but has no forward transition for $\mathbb{H} 1$, i.e. the operation explicitly halts the computation.



\blacklozenge implements a tape revert, i.e. moving the current tape position up by one. \blacklozenge acts like an activated boundary on the left hand side, i.e. it blocks ghosts \mathbb{Q} or \mathbb{Q} . An incoming ghost from the right can activate from \mathbb{Q} to \mathbb{Q} , after which it proceeds back to the right end. Outlined in section 8.3.4 but run backwards, the activated ghost will move through to the right hand side and deactivate at the boundary, while picking up a qubit \mathbb{D} . This qubit will move to the left until it encounters \blacklozenge . It drops the qubit and deactivates \blacklozenge to \blacklozenge . As soon as the inactive ghost \mathbb{Q} encounters this symbol, the ghost is reactivated and the box \square is restored.



Definition 64 (*Turing’s Wheelbarrow*). *Turing’s Wheelbarrow* is the Quantum Thue System $(\Gamma, R, \{\mathbb{U}_r\}_{r \in R}, \mathbb{C}^2)$, where Γ is given in Definition 62 and the

relation R is defined by the transition rules in Sect. 8.3 (with the conventions on notation from Definition 63).

One can verify that Turing’s Wheelbarrow, when applied to an initial string of the form $\| \exists a c c a \cdots c a \mathbb{M} \|$, where the sequence of c and a is such that they match the counting and computation phase, first translates the string length into a program description on the left string side, which is then executed cyclically on the tape. We call an initial configuration of this type *valid initial configuration*.

There are, however, ambiguous transitions, which lead to branching in the graph—we discuss all possible branching points for the irreducible evolution containing this initial configuration.

8.4. Branching in the History State

We make extensive use of branching and ambiguous transitions to compress the number of symbols necessary to implement the Wheelbarrow. Therefore we need to show two things.

1. The size of the history state is poly-bounded.
2. There are no ambiguous transitions which lead to a penalized configuration.

We take Fig. 15 as a point of reference.

Ghosts. Whenever there is a ghost on the tape, it can either be active— \mathbb{M} or \mathbb{M} , or inactive— \mathbb{M} or \mathbb{M} .


Fact 65. *Inactive ghosts never change non-head symbols or pass through heads.*

Therefore, we will disregard any branching due to inactive ghosts, which happens because we can always move either the head or ghost at each step. This increases the history state size by an at most quadratic factor.

Counting Phase.

Fact 66. *Initialization and 6-ary counter are not ambiguous in either direction.*

We have two ambiguities to analyse. Starting from an intermediary counting stage where the program string starts with \mathbb{H} , we can prematurely transition to the computation phase:







FIGURE 15. Evolution of the history state without branching

As there is no forward transition for $\mathbb{B}a$, this branch is just a leg, increasing the history state by a small constant factor ≤ 2 . The same argument holds for transitioning to \mathbb{I} during incrementation, i.e.

$$\begin{array}{c} \mathbb{H}x \triangleright x \cdots x a \cdots c \square a \cdots a \\ \mathbb{I}x \triangleright x \cdots x a \cdots c \square a \cdots a \end{array}$$

Run forward, there is no transition for $\mathbb{I} \blacktriangleleft$, and run backwards there is none for $\mathbb{I} \blacktriangleright$.

Running the carrier a or c backwards for the unary counter at any point before counting is completed leads to another ambiguity, e.g.




There is no backwards transition for $c \blacktriangleright$ or $a \blacktriangleright$ though. If there is no box $a \blacklozenge$, the branch dies off even before that. This ambiguity hence increases the history state size by another small constant factor.

Computation Phase. A similar argument as in the last section shows that a late transition into the counting phase once we are in the middle of the computation does not proceed, as there is no forward transition for a configuration $\blacktriangleright \square$. Furthermore, the same ambiguity running a carrier $c \blacklozenge$ or $a \blacklozenge$ backwards holds, which we have already discussed.

Fact 67. *The application of gates \mathbb{H} , \mathbb{U} , \mathbb{T} , \mathbb{A} and \mathbb{S} does not introduce any branching.*

It remains to analyse the revert command, where we have a branching point for a configuration



Observe, however, that all that could happen is that the tape symbol is carried to the right, where it is dropped next to the boundary ghost \blacklozenge . The ghost is activated and moves back to $\blacklozenge \blacklozenge$, where it deactivates. The branch does not continue further, as there is no transition out of $a \blacklozenge$ or $c \blacklozenge$. This increases the history state size by some small constant ≤ 2 .

We define the set of tuples \mathfrak{A} as all the possible character pairs that appear in this history state—including all branches—in Table 2.

This exhaustive analysis of all possible branching points in the history state allows us to conclude the following corollary.

Corollary 68. *For strings of length n , the size of the irreducible evolution containing a valid initial configuration of the form $\| \exists a c c a \cdots c a \blacklozenge \|$ —the history state—is of size $O(n^3)$, and contains no forbidden character pairs.*

TABLE 2. All possible character tuples occurring in the history state of Turing’s Wheelbarrow

	non-heads				heads																														
			x			c	a																	$!$	\bar{c}	\bar{i}	\bar{i}	\bar{c}	\bar{a}	\bar{a}	\bar{a}	\bar{a}			
							✓																												
								✓																											
x			✓	✓	✓	✓	✓	✓	✓				✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	
	✓							✓	✓	✓	✓	✓																							
									✓	✓	✓																								
										✓	✓																								
c	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓																							
a	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓																							
	✓	✓	✓																																
	✓																																		
	✓	✓	✓																																
	✓																																		
	✓																																		
	✓																																		
	✓																																		
	✓																																		
	✓																																		
	✓																																		
	✓																																		
	✓																																		
	✓																																		
	✓																																		
	✓																																		
	✓																																		
$!$																																			
\bar{c}																																			
\bar{i}																																			
\bar{i}																																			
\bar{c}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}																																			
\bar{a}						</																													

8.5. Simplicity of Turing’s Wheelbarrow

Let us briefly recall the idea behind simplicity in the context of QTSs. A QTS is called simple, if, for any two strings connected by more than one chain of transitions, the product of unitaries along this chain is identical. Equivalently, we can show that there are no loops in the graph connecting any strings. Regarding the QTS transition rules for Turing’s Wheelbarrow, as constructed in the last section, it is easy to see that it will not be simple. However, for our purposes, it suffices to prove the following lemma.

Lemma 69. *Each bracketed string in Turing’s Wheelbarrow with at least one head either belongs to the history state, which is simple, or—by removing edges—can be broken up into poly n -sized valid evolutions with illegal pairs.*

Proof. As no transition rule ever changes the number of heads or position of brackets, the distinction is well-defined. We can analyse each separately.

One Head. We can exclude strings with illegal pairs right away. Furthermore, we can disregard configurations of non-head characters which are just allowed because there is a head symbol or ghost between them, such as $a \text{⊠} c$, as moving the head either way (which is possible, since there is only one of them) transitions to an illegal pair.

So, disregarding any head and ghost state on the string for now, the most general non-head-non-boundary string compatible with Table 2 is

$$\underbrace{x^*}_A \text{⊠}^? \underbrace{(c \text{⊠} | a \text{⊠} | c | a)^+}_B.$$

It is straightforward to see that evolving this configuration backwards will transition to an illegal pair, if either

- B has multiple ⊠ , or $\text{⊠}^?$ and at least one ⊠ , or neither $\text{⊠}^?$ or ⊠ .
- A does not match the string length $\log_6 n$. Decrementation can only start if the substring A starts with the halt symbol H , so it cannot happen that we start decrementing a rotated number, e.g. U H A A S T instead of H A A S T U , which would translate into different lengths.

Evolving this string forward then reaches the computation part, and in case the pattern of classical and qubit states in B does not match the one required for the encoded gates in A we again have an illegal pair.

We are left with the history state, and it suffices to check any transition rule containing a non-trivial unitary attached, which by construction is the computational step only, i.e.

$$\begin{matrix} \dot{a} & a \\ \hline \dot{a} & a \end{matrix}.$$

Following the transitions forward to the next such transition, by construction, the encoded Turing machine evolution is reversible; hence, there is no loop as the Turing machine changes the classical content of the tape in section B .

Multiple Heads. None of the heads can pass through each other. As further boundary markers such as $\mathbb{1}$, $\mathbb{1}$ and $\mathbb{3}$ are immobile and opaque and there exists no transition out of \mathbb{A} or \mathbb{A} if not left of a boundary, we can without loss of generality assume that the tape is bracketed by either of $\mathbb{1}$, $\mathbb{3}$, \mathbb{A} , \mathbb{A} , or possibly no opaque symbol if our subsection lies at the tape ends.

If there are $h \geq 2$ head symbols on the tape, a simple argument allows us to slice the graph up into $\text{poly}(j(l))$ -sized parts: first observe that following any of the heads—with potential intermediate transitions—sweeps the entire width of the string. For any configuration of the first $h - 1$ heads, the last head will thus necessarily bump into the $h - 1^{\text{st}}$ within $O(j(l))$ steps. The same argument shows that there can be at most one ghost on the tape. \square

8.6. Special Properties

Proof of lemma 59. We will check the properties of Lemma 59 one-by-one.

The deciding property follows by construction. The projectors $\Pi_{\text{inp}} = \Pi_{\text{out}} = |1\rangle\langle 1|$ are supposed to act on checking the first ancilla and output as seen in Fig. 12, i.e. they apply to the qubit after the special identity symbol $\mathbb{1}$.

Item **W1** is readily verified.

Item **W2**. The encoding is given by the valid initial configuration

$$\text{enc}(l) := \mathbb{1} \underbrace{\mathbb{3} \ a \ c \ a \ c \ \dots \ a \ \mathbb{A} \ \mathbb{1}}_{N \text{ times}},$$

where N is a unary encoding of the QRM head circuit executing Fig. 12 rewritten as depicted in Fig. 13, and the sequence of a s and c s is such that they match the counting and computation phase. By construction, we therefore obtain, $N + 4 = \text{poly}(|l|) =: j(l)$.

Also by construction and as outlined in Sect. 8.1, the program string on the left side of enc describes the head of a QRM writing out the circuit Fig. 12. This QRM is in the same uniformity class as the original verifier's, and a constant in the size of $l \in \Pi$. We can hence pad it—using identity gates—to get the space and runtime for the QRM right, which can be as large as $\text{poly } j(l)$, as required for a BQEXP computation.

Both input and output markers $s_{\text{inp}} = s_{\text{out}} = \mathbb{A}a$ are 2-local, contain one head \mathbb{A} , and $\Pi_{\text{inp}} = \Pi_{\text{out}} = |1\rangle\langle 1|$.

Item **W3** We have shown the first claim in Lemma 69 and Corollary 68. The rest follows by direct verification.

Item **W4** We can immediately sort out the no-head and not-bracketed cases. The rest follows from Lemma 69.

This concludes the proof. \square

8.7. Final Dimension Reduction

We want to make a few final remarks, and suggest an immediate optimization of the Wheelbarrow construction.

The distinction between the quantum and classical tape symbols \bar{c} and \bar{a} is unnecessary, if we can ensure that there is never a quantum operation on classical symbols and vice versa. This is already proven.

The reason why we can merge these symbols is that while the QTS requires the ULG vertices to comprise only the classical alphabet symbols, we do not need to make this distinction for a ULG—as long as we can ensure that the Hilbert space dimension on each vertex in a connected component is the same. It is also clear that this does not break simplicity in Lemma 69, as we always *know* which tape symbols are classical (the ones appearing next to classical operations, e.g. $\mathbb{1}$) and which ones are quantum (e.g. the one next to $\mathbb{!}$). This observation allows the following optimization.

Remark 70. The Wheelbarrow construction works exactly the same when merging \bar{a} with \bar{c} , \bar{a} with \bar{c} , and \bar{a} with \bar{c} .

Once we have merged the symbols, there is another merge possible. We know that TOFFOLI and some basis-changing unitary \mathbf{U} are quantum-universal, see e.g. [37, ch. 4.5]. This means that we can replace the classically controlled unitary with such a one-qubit unitary, and apply TOFFOLI gates to quantum symbols as well. A similar argument as before shows that this does not break Lemma 69, and we phrase the following remark.

Remark 71. The Wheelbarrow construction works when replacing the controlled unitary with a single-qubit basis-changing unitary, and extending TOFFOLI to work on classical and quantum tape content. This makes the symbols $\mathbb{!}$, \bar{a} and \bar{a} obsolete.

Including the saved symbols from the last two remarks— \bar{c} , \bar{c} , \bar{c} , $\mathbb{!}$, \bar{a} and \bar{a} —we conclude with the existence proof of Lemma 59.

Corollary 72. *There exists a family of simple QTSs with 2-local rules on an alphabet of size $39-3$ of which are quantum with a Hilbert space \mathbb{C}^2 —and all properties given in Lemma 59.*

Remark 73. It is straightforward to get $O(1)$ -interactions, i.e. removing the scaling polynomial $p(l)$ in Theorem 60 if we can locally distinguish the history state at all times. This is possible e.g. by using distinct non-head symbols on the left- and right-hand side of the head and penalizing invalid configuration using regular expressions as in [3]. This would increase our dimension by roughly 15.

9. Conclusion

This work was motivated by the idea of finding a simple, translationally invariant and physically interesting system, for which the ground state energy problem is QMA_{EXP} -hard. In [3], Gottesman and Irani concluded that their construction is not “particularly natural”, due to the large local dimension

necessary, but that the existence of some very simple QMA_{EXP} -hard local Hamiltonian problems seems quite possible.

Our results bring us another step closer to this goal: we reprove the hardness result in [3] but with a local dimension of 42, whereas in [3]—though not explicitly specified—it was several orders of magnitude larger. To prove this result, we develop new tools and computational models which we believe are applicable to a wider range of problems.

At this point, it would be interesting to see where the threshold for the translationally invariant local Hamiltonian problem lies: does there exist a local dimension d_{\min} , for which the problem is in BQP , or BQEXP ? We have shown that $d_{\min} < 42$, but do not believe this to be a strict bound. We therefore encourage the interested reader to construct their own version of the Wheelbarrow, which might yield an even lower local dimension, and thus tighten our bound.

Furthermore, a lot of work recently has been done to analyse non-translationally invariant systems, and to classify interactions with locally-varying interaction strengths, e.g. [16, 17]. In contrast to our construction, the hardness results in [17] resemble more a tiling construction, a subject also addressed in [3]. It would be an interesting approach to see if these two—fundamentally quite different—results can be combined, or if there exists yet another, completely different, method of encoding computation into the ground state of a local Hamiltonian.

Finally, we want to mention that while the research focus—as outlined in Table 1—quickly shifted towards the 1D variant of the problem, from a physical perspective both 2D and 3D versions of this result are still of great interest, and apart from a trivial extension of our result to higher spatial dimensions, both remain open problems.

Acknowledgements

J.B. acknowledges support from the German National Academic Foundation and the Engineering and Physical Sciences Research Council (EPSRC Grant No. 1600123). T.S.C. is supported by the Royal Society. This work was made possible through the support of Grant #48322 from the John Templeton Foundation. M.O. acknowledges Leverhulme Trust Early Career Fellowship (ECF-2015-256) and European Union project QALGO (Grant Agreement No. 600700) for financial support. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the John Templeton Foundation.

Open Access. This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- [1] Hastings, M.B.: An area law for one-dimensional quantum systems. *J. Stat. Mech.: Theory Exp.* **08**, P08024 (2007)
- [2] Irani, S.: The complexity of quantum systems on a one-dimensional chain (2007). [arXiv:0705.4067](https://arxiv.org/abs/0705.4067)
- [3] Gottesman, D., Irani, S.: The quantum and classical complexity of translationally invariant tiling and Hamiltonian problems. *Theory Comput.* **9**(2), 31–116 (2013)
- [4] Bravyi, S., Caha, L., Movassagh, R., Nagaj, D., Shor, P.W.: Criticality without frustration for quantum spin-1 chains. *Phys. Rev. Lett.* **109**(20), 207202 (2012)
- [5] Chen, J., Chen, X., Duan, R., Ji, Z., Zeng, B.: No-go theorem for one-way quantum computing on naturally occurring two-level systems. *Phys. Rev. A* **83**(5), 050301 (2011)
- [6] Movassagh, R., Shor, P.W.: Supercritical entanglement in local systems: Counterexample to the area law for quantum matter. *PNAS* **113**(47), 13278–13282 (2016)
- [7] Cubitt, T., Perez-Garcia, D., Wolf, M.M.: Undecidability of the spectral gap. *Nature* **528**(7581), 207–211 (2015)
- [8] Bravyi, S., Gosset, D.: Gapped and gapless phases of frustration-free spin-1/2 chains. *J. Math. Phys.* **56**(6), 061902 (2015)
- [9] Bausch, J., Cubitt, T., Lucia, A., Perez-Garcia, D., Wolf, M.M.: Size-driven quantum phase transitions (Dec. 2016). [arXiv:1512.05687](https://arxiv.org/abs/1512.05687)
- [10] Kitaev, A.Y., Shen, A., Vyalıy, M.N.: Classical and quantum computing. In: *Quantum Information*, pp. 203–217. Springer, New York (2002). https://doi.org/10.1007/978-0-387-36944-0_13
- [11] Kempe, J., Kitaev, A.Y., Regev, O.: The complexity of the local Hamiltonian problem. *SIAM J. Comput.* **35**(5), 1070–1097 (2006)
- [12] Oliveira, R.I., Terhal, B.M.: The complexity of quantum spin systems on a two-dimensional square lattice. *Quantum Inf. Comput.* **8**(10), 0900–0924 (2008)
- [13] Aharonov, D., Gottesman, D., Irani, S., Kempe, J.: The power of quantum systems on a line. *Commun. Math. Phys.* **287**(1), 41–65 (2009)
- [14] Hallgren, S., Nagaj, D., Narayanaswami, S.: The local Hamiltonian problem on a line with eight states is QMA-complete. *Quantum Inf. Comput.* **13**(9&10), 0721–0750 (2013)
- [15] Barthel, T., Hübener, R.: Solving condensed-matter ground-state problems by semidefinite relaxations. *Phys. Rev. Lett.* **108**(20), 200404 (2012)
- [16] Cubitt, T., Montanaro, A.: Complexity classification of local Hamiltonian problems. *SIAM J. Comput.* **45**(2), 268–316 (2016)
- [17] Piddock, S., Montanaro, A.: The complexity of antiferromagnetic interactions and 2D lattices (June 2015). [arXiv:1506.04014](https://arxiv.org/abs/1506.04014)
- [18] Wei, T.-C., Liang, J.C.: Hamiltonian quantum computer in one dimension. *Phys. Rev. A* **92**(6), 062334 (2015)
- [19] Feynman, R.P.: Quantum mechanical computers. *Found. Phys.* **16**(6), 507–531 (1986)

- [20] Bravyi, S.: Efficient algorithm for a quantum analogue of 2-SAT. In: Mahdavi, K., et al. (eds.) *Cross Disciplinary Advances in Quantum Computing*, vol. 536, pp. 33–48. AMS (2011). <https://doi.org/10.1090/conm/536/10552>
- [21] Arad, I., Santha, M., Sundaram, A., Zhang, S.: Linear time algorithm for quantum 2SAT. In: Chatzigiannakis, I., et al. (eds) *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016) Leibniz International Proceedings in Informatics (LIPIcs)*, vol. 55, pp. 15:1–15:14. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl (2016)
- [22] de Beaudrap, N., Gharibian, S.: A linear time algorithm for quantum 2-SAT. In: Raz, R. (ed.) *31st Conference on Computational Complexity (CCC 2016)*. Leibniz International Proceedings in Informatics (LIPIcs), vol. 50, pp. 27:1–27:21. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Germany (2016)
- [23] Gosset, D., Nagaj, D.: Quantum 3-SAT is QMA1-complete. *SIAM J. Comput.* **45**(3), 1080–1128 (2016)
- [24] Landau, Z., Vazirani, U., Vidick, T.: A polynomial time algorithm for the ground state of one-dimensional gapped local Hamiltonians. *Nat. Phys.* **11**(7), 566–569 (2015)
- [25] Breuckmann, N.P., Terhal, B.M.: Space-time circuit-to-Hamiltonian construction and its applications. *J. Phys. A: Math. Theor.* **47**(19), 195304 (2014)
- [26] Nagaj, D., Wocjan, P.: Hamiltonian quantum cellular automata in one dimension. *Phys. Rev. A* **78**(3), 032311 (2008)
- [27] Nagaj, D.: Universal two-body-Hamiltonian quantum computing. *Phys. Rev. A* **85**(3), 032330 (2012)
- [28] Gosset, D., Terhal, B.M., Vershynina, A.: Universal adiabatic quantum computation via the space-time circuit-to-Hamiltonian construction. *Phys. Rev. Lett.* **114**(14), 140501 (2015)
- [29] Thomas, W.: When nobody else dreamed of these things—Axel Thue und die Termersetzung. *Informatik-Spektrum* **33**(5), 504–508 (2010)
- [30] Hopcroft, J.E., Ullman, J.D.: *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, Boston (1979)
- [31] Morita, K.: Reversible computing and cellular automata—a survey. *Theor. Comput. Sci.* **395**(1), 101–131 (2008)
- [32] Bernstein, E., Vazirani, U.: Quantum complexity theory. *SIAM J. Comput.* **26**(5), 1411–1473 (1997)
- [33] Bennett, C.H.: Logical reversibility of computation. *IBM J. Res. Dev.* **17**(6), 525–532 (1973)
- [34] Perumalla, K.S.: *Introduction to Reversible Computing* (Chapman & Hall/CRC Computational Science Series). CRC Press, Boca Raton (2013)
- [35] Watrous, J.: Quantum computational complexity. In: Meyers, R.A. (ed.) *Computational Complexity: Theory, Techniques, and Applications*, pp. 2361–2387. Springer, New York (2012)
- [36] Vidick, T., Watrous, J.: Quantum proofs. *Found. Trends Theor. Comput. Sci.* **11**(1–2), 1–215 (2016)
- [37] Nielsen, M.A., Chuang, I.L.: *Quantum Computation and Quantum Information*, 10th edn. Cambridge University Press, Cambridge (2010)
- [38] Gharibian, S., Huang, Y., Landau, Z., Woo Shin, S.: Quantum Hamiltonian complexity. *Found. Trends Theor. Comput. Sci.* **10**(3), 159–282 (2014)

- [39] Kitaev, A.Y., Shen, A., Vyalıy, M.N.: Classical and Quantum Computing (Graduate Studies in Mathematics), vol. 47. American Mathematical Society, Providence (2002)
- [40] Bausch, J., Crosson, E.: Increasing the quantum UNSAT penalty of the circuit-to-Hamiltonian construction, Sept 2016. [arXiv:1609.08571](https://arxiv.org/abs/1609.08571)
- [41] Trudeau, R.J.: Introduction to Graph Theory (Dover Books on Mathematics). Dover Publications, New York (2013)
- [42] Diestel, R.: Graph Theory (Graduate Texts in Mathematics), vol. 173, 5th edn. Springer, New York (2016)
- [43] Godsil, C., Royle, G.: Algebraic Graph Theory (Graduate Texts in Mathematics), vol. 207. Springer, New York (2001)
- [44] Fiedler, M.: Algebraic connectivity of graphs. Czechoslov. Math. J. **23**(2), 298–305 (1973)

Johannes Bausch and Maris Ozols
DAMTP, Centre for Mathematical Sciences
University of Cambridge
Cambridge CB3 0WB
UK
e-mail: jkrb2@cam.ac.uk

Maris Ozols
e-mail: marozols@gmail.com

Toby Cubitt
Department of Computer Science
University College London
London WC1E 6BT
UK
e-mail: t.cubitt@ucl.ac.uk

Communicated by David Perez-García.
Received: January 16, 2017.
Accepted: June 29, 2017.