Complexity of commuting Hamiltonians on a square lattice of qubits

Norbert Schuch
California Institute of Technology, Institute for Quantum Information, MC 305-16, Pasadena CA 91125, U.S.A.

We consider the computational complexity of Hamiltonians which are sums of commuting terms acting on plaquettes in a square lattice of qubits, and we show that deciding whether the ground state minimizes the energy of each local term individually is in the complexity class NP. That is, if the ground states has this property, this can be proven using a classical certificate which can be efficiently verified on a classical computer. Different to previous results on commuting Hamiltonians, our certificate proves the existence of such a state without giving instructions on how to prepare it.

Full version available at arXiv:1105.2843

Background.—The study of Hamiltonian complexity, which in essence captures the difficulty of determining the ground state energy of a local Hamiltonian, has attracted considerable attention in the last years. For classical systems, the problem of deciding whether there exists a state with energy below a certain threshold is NP–hard: Clearly, the ground state can be described classically, and its energy can be computed efficiently; and NP–hardness has been proven e.g. for Ising spin glasses [1]. For quantum Hamiltonians, the best one can say a priori is that it is possible for a prover to provide the ground state as a quantum state, and that its energy can be efficiently estimated by a quantum computer to polynomial accuracy, putting the problem in the class QMA; and it was indeed proven by Kitaev that the general problem is hard for the class QMA [2]. This result has later been generalized to various setups, including two-dimensional lattices of qubits [3] and (in strong contrast to classical Hamiltonians) one-dimensional chains [4].

It is an interesting question where the additional complexity of quantum spin systems as compared to classical systems arises. To this end, one can study restricted classes of Hamiltonians which are "more classical" than general quantum Hamiltonians. One such class is formed by *commuting Hamiltonians*, that is, Hamiltonians which can be written as a sum of mutually commuting few-body terms. For those systems, all terms can be simultaneously diagonalized, just as for classical systems; however, the corresponding eigenbasis can be highly entangled, making it unclear whether a useful classical description of the ground state can be provided. In fact, commuting Hamiltonians encompass e.g. models with topological order, such as Kitaev's toric code [5] or Levin and Wen's string net models [6]. Commuting Hamiltonians are also of interest from the point of view of condensed matter physics, since the fixed points of renormalization flows in gapped phases are expected to be commuting Hamiltonians, and thus understanding their structure should give insight into the structure of gapped quantum phases. Finally, understanding the complexity of commuting Hamiltonians might form a step towards a quantum PCP theorem, which would assess how the difficulty of estimating the ground state energy is related to the desired accuracy which is integer for commuting projectors.

The problem we will consider is COMMUTING HAMILTONIAN: Given a Hamiltonian with commuting terms, does there exist a state which minimizes the energy of each term in the Hamiltonian individually? For lattices in two and higher dimensions, COMMUTING HAMILTONIAN is an NP–hard problem, as it encompasses NP–hard classical models [1]. On the other hand, it is not clear whether the general COMMUTING HAMILTONIAN problem is *inside* NP, since it is not clear in general how to provide an efficiently checkable description of the ground state. For two-local (i.e., two-body) Hamiltonians, Bravyi and Vyalyi [7] have shown that the problem is in NP (their result also implies that one-dimensional commuting Hamiltonians are efficiently solvable); subsequently, Aharonov and Eldar [8] have proven containment in NP for Hamiltonians with three-body interactions both for qubits on arbitrary graphs, and qutrits on nearly-Euclidean interaction graphs. In all of these cases, the classical certificates do not only prove the existence of a such a ground state, but can in fact be used to construct constant depth quantum circuits which *prepare* the ground state. This implies that the corresponding Hamiltonians – including qutrits with three-body interactions – cannot exhibit topological order [8, 9]. On the other hand, Kitaev's toric code, which is the ground state of a commuting Hamiltonian with four-body interactions of qubits, does have topological order, and thus, we cannot expect any approach which yields a low-depth circuit to work beyond three qutrits.

In this work, we study the COMMUTING HAMILTONIAN problem on a square lattice of qubits with plaquette-wise interactions, and prove that it is in NP. That is, we consider a square lattice of qubits, with a Hamiltonian with mutually commuting terms acting on the four qubits adjacent to each plaquette, and show that the problem of deciding whether its ground state minimizes the energy of each local term in the Hamiltonian is in NP: i.e., in case the ground state has this property, a classical certificate exists which can be checked efficiently by a quantum computer. Our approach differs considerably from the aforementioned approaches in that the certificate cannot

be used to devise a quantum circuit for preparing the ground state, and is thus also applicable to systems with topological order.

The setup.—We consider a 2D square lattice with spins on the vertices. The Hamiltonian $H = \sum_p h_p$, $[h_p, h_q] = 0$, consists of plaquette terms h_p which act on the four spins adjacent to the plaquette p. We are interested in the complexity of the following problem, called COMMUTING HAMILTONIAN: Is there an eigenstate $|\psi\rangle$ of H which minimizes the energies for all h_p individually, i.e., are the ground states of H also ground states of each h_p? As we will show, in the case of qubits the existence of such a state can be proven within NP, i.e., there is a classical certificate which proves the existence of such a $|\psi\rangle$, and which can be checked efficiently classically. (Note that NP–hardness follows from the NP–hardness results for Ising spin glasses.)

In the following, we will use the following reformulation of COMMUTING HAMILTONIAN: Define Π_p as the projector onto the ground state subspace of h_p; again, $[\Pi_p, \Pi_q] = 0$. Then, $\Pi_{\mathrm{GS}} = \prod_p \Pi_p$ is the projector onto the subspace spanned by the states which are ground states of all h_p. COMMUTING HAMILTONIAN asks whether such states exist, i.e., whether $\Pi_{\mathrm{GS}} \neq 0$.

The proof that COMMUTING HAMILTONIAN is in NP consists of two steps: First, we show that the Hamiltonian can be split into two layers which both can be solved in NP, and the problem reduces to computing the overlap of the two solutions, and second, we show that this overlap can indeed be computed efficiently, which puts the overall problem in NP.

Step I: Splitting into two layers.—We start by coloring the plaquettes of the square lattice black and white in a checkerboard pattern, and denote the set of black and white plaquettes by $\mathcal{P}_B$ and $\mathcal{P}_W$, respectively. Let

$$\Pi_B = \prod_{p\in\mathcal{P}_B} \Pi_p \quad \text{and} \quad \Pi_W = \prod_{p\in\mathcal{P}_W} \Pi_p \ ;$$

then, COMMUTING HAMILTONIAN corresponds to determining whether $\Pi_B \Pi_W \neq 0$, or equivalently

$$\mathrm{tr}[\Pi_B \Pi_W] \neq 0 \ . \tag{1}$$

To prove COMMUTING HAMILTONIAN is contained in NP, we therefore need to show that a classical certificate for the validity of (1) can be provided.

A helpful example to keep in mind is Kitaev's toric code [5]: There, $\Pi_p = \frac{1}{2}(\mathbb{1} + Z^{\otimes 4})$ for $p \in \mathcal{P}_B$, and $\Pi_p = \frac{1}{2}(\mathbb{1} + X^{\otimes 4})$ for $p \in \mathcal{P}_W$.

Let us next study the structure of Π_B. Using C*-algebraic techniques introduced to the problem in [7], one can show that there exists a decomposition of the Hilbert space at vertex v into subspaces $\sum_{\alpha_v} \pi^v_{\alpha_v} = \mathbb{1}$, with the $\pi^v_{\alpha_v}$ projectors, such that Π_B projected onto any "slice" $\vec{\alpha} = (\alpha_v)_{v\in V}$ (with projector $\otimes_v \pi^v_{\alpha_v}$) factorizes, i.e., Π_B can be written as

$$\Pi_B = \bigoplus_{\vec{\alpha}} \bigotimes_{p\in\mathcal{P}_B} \Pi_p^{\vec{\alpha}} \equiv \sum_{\vec{\alpha}} \bigotimes_{p\in\mathcal{P}_B} \Pi_p^{\vec{\alpha}} \ . \tag{2}$$

Here, the projectors $\Pi_p^{\vec{\alpha}}$ are obtained by restricing Π_p to slice $\vec{\alpha}$. each $\Pi_p^{\vec{\alpha}}$ acts only on plaquette p, but the $\Pi_p^{\vec{\alpha}}$ for different plaquettes now factorize; note that the the expression on the r.h.s. should be understood with respect to the natural embedding of the "slices" into the full Hilbert space.

The same decomposition for the white sublattice gives a different decomposition into subspaces $\bar{\pi}^v_{\beta_v}$ at each vertex, which induce a factorizing decomposition

$$\Pi_W = \sum_{\vec{\beta}} \bigotimes_{p\in\mathcal{P}_W} \Pi_p^{\vec{\beta}} \ . \tag{3}$$

E.g., for Kitaev's toric code, the $\pi^v_{\alpha_v}$ are projectors onto the Z eigenstates, and the $\bar{\pi}^v_{\beta_v}$ onto the X eigenstates.

Using Eqs. (2) and (3), we can rewrite the COMMUTING HAMILTONIAN problem, Eq. (1), as

$$0 \neq \sum_{\vec{\alpha},\vec{\beta}} \underbrace{\mathrm{tr}\left[\Big(\bigotimes_{p\in\mathcal{P}_B} \Pi_p^{\vec{\alpha}}\Big)\Big(\bigotimes_{p\in\mathcal{P}_W} \Pi_p^{\vec{\beta}}\Big) \right]}_{=:\Omega(\vec{\alpha},\vec{\beta})} \ . \tag{4}$$

Since each of the traces is non-negative, we can ask the prover to provide us with $\vec{\alpha}$ and $\vec{\beta}$ such that $\Omega(\vec{\alpha},\vec{\beta}) \neq 0$ – clearly, such $\vec{\alpha}$ and $\vec{\beta}$ exist only for yes-instances. In order to prove containment in NP, it remains to show that the overlap $\Omega(\vec{\alpha},\vec{\beta})$ can be computed efficiently (or at least in NP).

For Kitaev's toric code, e.g., we could choose $\vec{\alpha} = \vec{\beta} = (0,\ldots,0)$: This yields

$$\Omega(\vec{\alpha},\vec{\beta}) = \mathrm{tr}\big[(|0\rangle\langle 0|)^{\otimes N}(|+\rangle\langle +|)^{\otimes N}\big] = 2^{-N} \neq 0 \ ,$$

which is efficiently computable and proves the existence of a zero-energy ground state; note that this certificate does not carry any information how to prepare the state.

Step II: Computing the overlap.—The preceding discussion holds for any local Hilbert space dimension. In the following, we will restrict to the case of qubits to show how to compute the overlap $\Omega(\vec{\alpha},\vec{\beta})$. We start by tracing over all sites on which $\Pi_p^{\vec{\alpha}}$ or $\Pi_p^{\vec{\beta}}$ act on a one-dimensional subspace only, i.e., the corresponding projection $\pi^v_{\alpha_v}$ or $\bar{\pi}^v_{\beta_v}$ is one-dimensional. Tracing out those qubits gives new states ρ_p for each plaquette (we omit $\vec{\alpha}$ and $\vec{\beta}$ from now on), moreover, all traced out qubits are now missing from the lattice. Thus, we can reexpress

$$\Omega(\vec{\alpha},\vec{\beta}) = T_{1\mathrm{D}}\, \Theta$$

where $T_{1\mathrm{D}}$ is the prefactor obtained from tracing out the aforementioned qubits, and Θ is the overlap of the states ρ_p supported on the white and on the black plaquettes; by construction, the ρ_p are now such that for two diagonally adjacent plaquettes, only one ρ_p acts on the shared qubit.

Clearly, if $T_{1\mathrm{D}} = 0$, we can reject the proof. It remains to show that we can check efficiently whether $\Theta \neq 0$.

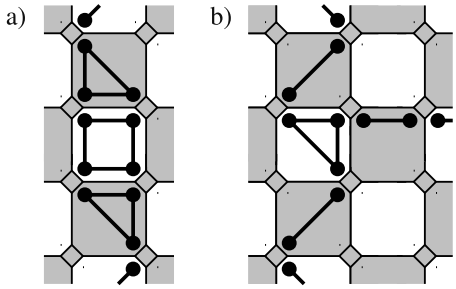

FIG. 1. Computing the overlap Θ for qubits. The diamonds at the vertices of the square lattice denote the qubits. The connected black dots mark qubits on which the ρ_p act non-trivially on plaquette p. If the ρ_p of two plaquettes act non-trivially on the same qubit, we say that they "overlap"; note that this cannot happen for diagonally adjacent plaquettes. Overlapping ρ_p's form structures which we need to contract to compute the overlap Θ. **a)** Patterns forming one-dimensional chains can be contracted efficiently, as the size of the boundary stays constant for any contiguous region. **b)** Branching structures do in general not allow for efficient contraction. However, we prove that such structures cannot occur.

The situation encountered in computing the overlap Θ is depicted in Fig. 1. Here, the dots in each plaquette denote the vertices on which ρ_p acts non-trivially (the lines just connect the vertices involved in ρ_p). If the ρ_p on adjacent plaquettes act non-trivially on the same qubit (we will say they "overlap"), they form connected structures which we need to contract in order to evaluate Θ. For one-dimensional structures as the one on in Fig. 1a, this contraction can be carried out efficiently: One starts from one plaquette and proceeds along one direction of the one-dimensional chain, always tracing out the degrees of freedom on the inside. This way, at every point in the computation only the state at the boundary (which has fixed size) needs to be stored, and thus, the contraction can be carried out efficiently. On the other hand, branching structures like the one in Fig. 1b can in general not be contracted efficiently, since the size of the boundary is a priori not bounded.

However, as we can show, the structures formed by the ρ_p in Θ will always be one-dimensional, and thus Θ can be computed efficiently. To this end, we consider the state ρ_C on a plaquette C (the "central" plaquette), and show that it can overlap non-trivially with the states ρ_p of at most two of the adjacent plaquettes, thus ruling out branching structures as the one in Fig. 1b. The first important ingredient for the proof is that in each layer, at most one plaquette term ρ_p can act non-trivially on any given vertex; in the graphical notation of Fig. 1, we highlight this fact by placing a cross opposite of any dot:

 (5)

This in particular implies that ρ_C can at most overlap non-trivially with the four horizontally and vertically adjacent plaquettes from the other layer, but not with diagonally adjacent plaquettes.

As the second ingredient, we prove the following implication:

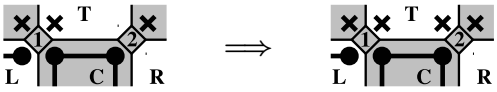

(6)

This is, if in the configuration on the l.h.s., ρ_L acts non-trivially on qubit 1, then ρ_T has to act trivially on both qubits 1 and 2, and thus cannot overlap with ρ_C. (Proof sketch: Since ρ_T cannot act on qubit 1, the corresponding $\Pi_T^{\vec{\beta}}$ has to commute with $\Pi_C^{\vec{\alpha}}$ on qubit 2 only; but since ρ_C acts non-trivially on qubit 2, ρ_T has to act trivially on it.) Using Eq. (6), we then show that the structures in the overlap Θ are all one-dimensional, since every state ρ_C on a "central" plaquette can only overlap with the ρ_p on at most two adjacent plaquettes.

Together, this shows that the overlap Θ can be computed efficiently, which in turn implies that for given $\vec{\alpha}$ and $\vec{\beta}$, $\Omega(\vec{\alpha}, \vec{\beta})$ can be computed efficiently, and thus, the commuting Hamiltonian problem on a square lattice of qubits with plaquette interactions is in NP.

Summary.—We have proven that the COMMUTING HAMILTONIAN problem on a square lattice of qubits with plaquette-wise interaction is NP–complete, i.e., there exists an efficiently checkable classical certificate proving that the ground state of the system minimizes each local term. The idea of the proof is to split the system into two layers, and to argue that the existence of a state minimizing all local terms is equivalent to the existence of a pair of ground states for the two layers with non-zero overlap. The ground state subspace of each layer could be described efficiently, using the methods introduced in [7]. Finally, we showed that the overlap of ground states of two layers can be computed efficiently as it forms one-dimensional structures. A somewhat surprising feature of our approach is that while it certifies the existence of a ground state, it cannot (to our knowledge) be used to devise a way how to prepare the ground state; in fact, due to the possibility of having topological order in such systems, any circuit preparing their ground states would need to have at least logarithmic depth, or linear depth if it was local [9].

Our method does, in principle, also apply beyond qubits, as long as there exists a way to compute $\Omega(\vec{\alpha}, \vec{\beta})$ in NP, i.e., with additional information provided by the prover. In particular, this applies to the case where the decomposition in the direct sum gives one-dimensional spaces, such as in Kitaev's toric code or quantum double models; as well as to cases where the ρ_p are separable states. Our idea also applies to any other graph which can be split into two layers in such a way that the C^*–technique of [7] can be applied to each of them.

[1] F. Barahona, J. Phys. A **15**, 3241 (1982).
[2] D. Aharonov and T. Naveh(2002), quant-ph/0210077.

[3] R. Oliveira and B. M. Terhal, Quant. Inf. Comput. **8**, 900 (2009), quant-ph/0504050.
[4] D. Aharonov, D. Gottesman, S. Irani, and J. Kempe, Commun. Math. Phys. **287**, 41 (2009), arXiv:0705.4077.
[5] A. Kitaev, Ann. Phys. **303**, 2 (2003), quant-ph/9707021.
[6] M. A. Levin and X.-G. Wen, Phys. Rev. B **71**, 045110 (2005), cond-mat/0404617.
[7] S. Bravyi and M. Vyalyi, Quant. Inf. Comput. **5**, 187 (2005), quant-ph/0308021.
[8] D. Aharonov and L. Eldar(2011), arXiv:1102.0770.
[9] S. Bravyi, M. B. Hastings, and F. Verstraete, Phys. Rev. Lett. **97**, 050401 (2006), quant-ph/0603121.