Quantum Information and Computation, Vol. 11, No. 11&12 (2011) 0901–0912

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, United States of America

Received May 30, 2011
Revised July 29, 2011

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.

Keywords: Quantum complexity, Hamiltonian complexity, commuting Hamiltonians
Communicated by: R Jozsa & B Terhal

1 Introduction

Understanding the ground state properties of spin systems on a lattice is of central importance in many-body physics, but at the same time, it is a highly challenging problem in many scenarios. An important step in understanding its difficulty has been the insight that computing e.g. the ground state energy of a classical spin system is, in general, an NP–complete problem [1]: While the energy of any given spin configuration can be easily computed, finding the configuration with minimal energy is in general a difficult task – it can be as hard as any problem in NP, i.e., any problem whose solution can be efficiently verified. For quantum spin systems, an additional difficulty arises: Generally, we cannot even expect to have an efficient description of the ground state. Thus, it seems that the only statement we can make about the difficulty of the problem is that given a quantum register with the ground state, we will be able to efficiently estimate its energy using a quantum computer. Indeed, it has been shown that this is the best we can say, as the problem of estimating the ground state energy of a quantum system is a complete problem for the class QMA [2, 3], the quantum analogue of NP: It contains all problems which have a quantum solution which can be efficiently checked on a quantum computer, and thus, determining the ground state energy of a quantum spin system is as hard as any of these problems; in fact, the problem retains its hardness even when restricted to two-dimensional lattices of qubits with nearest-neighbor interactions [4] or one-dimensional chains [5].

It is an interesting question to understand the reasons underlying the additional complexity of quantum spin systems as compared to classical systems. To this end, restricted version of

the problem which lie between classical and general quantum spin Hamiltonians have been studied: For instance, it has been shown that so-called *stoquastic* Hamiltonians, where all off-diagonal elements are negative, have a complexity which lies in between NP and QMA, as those systems can be related to classical random processes [6, 7]; in fact, these are exactly the Hamiltonians which allow for Quantum Monte Carlo simulations as they do not exhibit a sign problem.

Another restricted class of Hamiltonians are *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 systems which exhibit rich non-classical physics, in particular models with topological order and even anyonic excitations, such as Kitaev's toric code and quantum double models [8], or Levin and Wen's string net models [9]. Commuting Hamiltonians are also of interest since the fixed points of renormalization flows in gapped phases are expected to be commuting Hamiltonians, and thus understanding their structure might give insight into the structure of gapped quantum phases. Finally, understanding the complexity of commuting Hamiltonians is of interest in quantum complexity, as it might be 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.

What is know about the complexity of finding the ground state energy of commuting Hamiltonians, or rather, of determining whether the ground state minimizes each term in the Hamiltonian individually – the COMMUTING HAMILTONIAN problem? For lattices in two and higher dimensions, COMMUTING HAMILTONIAN is an NP–hard problem, as it e.g. encompasses the Ising model on a planar graph [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 [10] have shown that the problem is in NP by using C^*–algebraic techniques (their result also implies that one-dimensional commuting Hamiltonians are efficiently solvable); subsequently, Aharonov and Eldar [11] 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 both of these works, the classical certificates do not only prove the problem to be in NP, but can in fact be used to construct constant depth quantum circuits which *prepare* the ground state. This, in particular, implies that the corresponding Hamiltonians – including qutrits with three-body interactions – cannot exhibit topological order [11, 12]. 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 paper, 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; it should be noted that the same holds true for the proof in Ref. [10] that COMMUTING HAMILTONIAN with factorizing projectors is in NP.

2 The setup

We will consider a 2D square lattice with spins on the vertices, and either open or periodic boundary conditions. The Hamiltonian

$$H = \sum_p h_p$$

consists of plaquette terms h_p which act on the four spins adjacent to the plaquette p, and we impose that all the terms in the Hamiltonian commute, $[h_p, h_q] = 0 \; \forall p, q$.

As the Hamiltonian terms commute, there is a basis of eigenstates of H which are simultaneously eigenstates of all the h_p. We would like to know the computational difficulty 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? In the following, we will show that 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 on the other hand, it is clear that the problem is NP–hard, as it e.g. encompasses classical Ising spin glasses in a field which are known to be NP–hard even for two-level systems [1].

For the following, it will be useful to reformulate COMMUTING HAMILTONIAN as follows: Define the *local ground state projectors* Π_p as the projectors onto the ground state subspace of h_p; the Π_p commute again, $[\Pi_p, \Pi_q] = 0$. Then,

$$\Pi_{\mathrm{GS}} = \prod_{p \in \mathcal{P}} \Pi_p$$

is the projector onto the subspace spanned by the states which are ground states of all h_p. Since COMMUTING HAMILTONIAN asks whether such states exist, it is equivalent to asking whether $\Pi_{\mathrm{GS}} \neq 0$.

3 Commuting Hamiltonian in NP

3.1 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

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

be the projectors onto the joint ground state space of the black and white h_p, respectively; then, COMMUTING HAMILTONIAN corresponds to determining whether $\Pi_B \Pi_W \neq 0$. This is equivalent to asking whether

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

(this can be seen using eigendecompositions of Π_W and Π_B), and we will consider this formulation of the problem from now on; 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 in the first part of our discussion is Kitaev's toric code [8]: 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 $p \in \mathcal{P}_W$, with X and Z the Pauli matrices.

3.2 The structure of one layer

In the following, let us study the structure of each layer individually (we will w.l.o.g. choose black). To this end, we will use a result of Bravyi and Vyalyi based on C^*-algebraic techniques [10]; a detailed explanation of those techniques can also be found in [11]. The basic insight from Ref. [10] is the structure of two commuting projectors. Consider two projectors $L \equiv L_{AB} \otimes \mathbb{1}_C$ and $R \equiv \mathbb{1}_A \otimes R_{BC}$ acting on a space $\mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_C$; the two operators overlap on $\mathcal{H}_B$. Now consider the Schmidt decompositions

$$L = \sum A_L^i \otimes B_L^i \quad \text{and} \quad R = \sum B_R^i \otimes C_R^i \ ,$$

i.e. $\mathrm{tr}[A_L^i (A_L^j)^\dagger] = 0$ for $i \neq j$, and similarly for B_L^i, B_R^i, and C_R^i. Then, $[L, R] = 0$ implies that $[B_L^i, B_R^j] = 0$ for all i, j, and thus, B_L^i and B_R^i span commuting C^*-algebras, cf. [10]. Using the standard form of finite C^*-algebras, it follows that the space $\mathcal{H}_B$ has a canonical decomposition

$$\mathcal{H}_B = \bigoplus_\alpha \underbrace{\mathcal{H}_L^\alpha \otimes \mathcal{H}_R^\alpha \otimes \mathcal{H}_Z^\alpha}_{=:\mathcal{H}_B^\alpha} \tag{2}$$

where the B_L^i span the full matrix algebra on $\mathcal{H}_L^\alpha$ while acting trivially on the rest, and correspondingly for B_R^i and $\mathcal{H}_R^\alpha$.

This shows that the space $\mathcal{H}_B$ can be cut into direct sum "slices" (the α) such that in each slice, L and R act on independent subsystems. More formally, there exists a decomposition $\mathbb{1} = \sum_\alpha \pi_\alpha$ of $\mathcal{H}_B$, with π_α the projectors onto $\mathcal{H}_B^\alpha$, such that

$$[\pi_\alpha, L] = 0 \text{ and } [\pi_\alpha, R] = 0$$

and thus

$$L = \sum_{\alpha,\beta} \pi_\alpha L \pi_\beta = \sum_\alpha \underbrace{\pi_\alpha L \pi_\alpha}_{=:L^\alpha} \ , \quad R = \sum_\alpha \underbrace{\pi_\alpha R \pi_\alpha}_{=:R^\alpha} \ ,$$

where L^α and R^α act on different subsystems $\mathcal{H}_L^\alpha$ and H_R^α, i.e., factorize. Note that the above decomposition allows to compute the π_α and thus the L_α and R_α efficiently.

Each vertex in the black sublattice is acted upon by exactly two commuting projectors Π_p; thus, we can apply the preceding argument to all vertices to find decompositions $\pi_{\alpha_v}^v$, $\sum_{\alpha_v} \pi_{\alpha_v}^v = \mathbb{1}$, of the Hilbert space at each vertex v, such that Π_B projected onto the slice $\vec{\alpha} = (\alpha_v)_{v \in V}$ factorizes,

$$\bigotimes_{p \in \mathcal{P}_B} \Pi_p^{\vec{\alpha}} = \prod_{v \in V} \pi_{\alpha_v}^v \, \Pi_B \, \pi_{\alpha_v}^v$$

(this implicitly defines the $\Pi_p^{\vec{\alpha}}$), and Π_B can be written as

$$\Pi_B = \bigoplus_{\vec{\alpha}} \bigotimes_{p \in \mathcal{P}_B} \Pi_p^{\vec{\alpha}} \equiv \sum_{\vec{\alpha}} \bigotimes_{p \in \mathcal{P}_B} \Pi_p^{\vec{\alpha}} \ . \tag{3}$$

Note that in the sum on the right hand side, we implicitly regard the tensor products as being canonically embedded into the full Hilbert space, and we will use this convention in the following.

The analogous decomposition can be performed for the white sublattice, yielding a (in general different!) decomposition

$$\Pi_W = \bigoplus_{\vec{\beta}} \bigotimes_{p\in\mathcal{P}_W} \Pi_p^{\vec{\beta}} \equiv \sum_{\vec{\beta}} \bigotimes_{p\in\mathcal{P}_W} \Pi_p^{\vec{\beta}} \ . \tag{4}$$

Note that in order to distinguish the decomposition for the the black and the white layer, we strictly use labels $\vec{\alpha}$ and $\vec{\beta}$, respectively; moreover, we denote the projectors decomposing the white layer by $\bar{\pi}^v_{\beta_v}$.

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.

3.3 Combining the layers

Using the stucture of Π_B and Π_W, Eqs. (3) and (4), we can rewrite the COMMUTING HAMILTONIAN problem, Eq. (1), as

$$0 \neq \sum_{\vec{\alpha},\vec{\beta}} \mathrm{tr}\left[\Big(\bigotimes_{p\in\mathcal{P}_B} \Pi_p^{\vec{\alpha}}\Big)\Big(\bigotimes_{p\in\mathcal{P}_W} \Pi_p^{\vec{\beta}}\Big)\right] ; \tag{5}$$

recall that we regard the tensor products as being canonically embedded into the full Hilbert space. Since each of the traces is positive (as it is the trace of the product of two positive operators), the above is equivalent to the existence of $\vec{\alpha}$ and $\vec{\beta}$ such that $\Omega(\vec{\alpha},\vec{\beta}) \neq 0$, where

$$\Omega(\vec{\alpha},\vec{\beta}) := \mathrm{tr}\left[\Big(\bigotimes_{p\in\mathcal{P}_B} \Pi_p^{\vec{\alpha}}\Big)\Big(\bigotimes_{p\in\mathcal{P}_W} \Pi_p^{\vec{\beta}}\Big)\right] . \tag{6}$$

Thus, we can ask the prover to provide us as a certificate with some $\vec{\alpha}$ and $\vec{\beta}$ for which $\Omega(\vec{\alpha},\vec{\beta}) \neq 0$; if we can moreover show that $\Omega(\vec{\alpha},\vec{\beta})$ can be computed efficiently (or rather in NP), this will prove that COMMUTING HAMILTONIAN is in NP.

Note that $\Omega(\vec{\alpha},\vec{\beta})$ can be interpreted as the overlap of the (unnormalized) states $\bigotimes \Pi_p^{\vec{\alpha}}$ and $\bigotimes \Pi_p^{\vec{\beta}}$, both of which are tensor products of states supported on individual plaquettes, but with different partitions in the two layers. Computing such an overlap can in general be as hard as contracting Projected Entangled Pair States (PEPS) [13], i.e., PP-hard [14]: any PEPS can be expressed this way by using one layer for the bonds and the other for the PEPS projections. Of course, the fact that these states arise from two commuting layers Π_B and Π_W yields additional constraints, and we will show in the following that those constraints allow for the efficient evaluation of $\Omega(\vec{\alpha},\vec{\beta})$ in the case of qubits.

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

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

proving the existence of a zero-energy ground state; note that this certificate does not carry any information on how to prepare the state.

3.4 Computing the overlap

Let us now show that for a lattice of qubits, the overlap $\mathcal{O}(\vec{\alpha},\vec{\beta})$ can be computed efficiently. To this end, we will show that the computation of the overlap can be decomposed into a product of overlaps of one-dimensional structures which can be computed efficiently.

Let us first consider the black layer. For each vertex v, the decomposition (2) of the local Hilbert space can either be trivial (no direct sum) or non-trivial, $\mathbb{1} = \sum \pi^v_{\alpha_v}$. In the former case, this implies that at most one of the adjacent plaquette terms Π_p acts non-trivially of vertex v; in the latter case, the sum consists of exactly two *one-dimensional* projectors $\pi^v_{\alpha_v}$, making use of the fact that the Hilbert space at each site is a qubit, i.e., two-dimensional. We will denote the set of vertices with a non-trivial decomposition in the black layer by $\mathcal{F}_B$, and in the white layer by $\mathcal{F}_W$ (in which the one-dimensional projectors are labelled $\bar{\pi}^v_{\beta_v}$).

In $\Omega(\vec{\alpha},\vec{\beta})$, all vertices in $\mathcal{F}_B \cup \mathcal{F}_W$ contribute only a one-dimensional subspace and thus can be traced out: That is, all vertices in $\mathcal{F}_B \cap \mathcal{F}_W$ can be removed (taking care whether the overlap of the one-dimensional projectors is non-vanishing), while for vertices where only one layer has a one-dimensional decomposition, this yields new effective plaquette terms ρ_p in the other layer by projecting the original plaquette terms Π_p onto that one-dimensional subspace; thus, the problem of checking whether $\Omega(\vec{\alpha},\vec{\beta})$ is non-zero reduces to computing the overlap of the new effective plaquette terms ρ_p. Formally, this reads

$$\begin{aligned}
\Omega(\vec{\alpha},\vec{\beta}) &= \mathrm{tr}\left[\Big(\bigotimes_{p\in\mathcal{P}_B} \Pi_p^{\vec{\alpha}}\Big)\Big(\bigotimes_{p\in\mathcal{P}_W} \Pi_p^{\vec{\beta}}\Big) \right] \\
&\overset{(A)}{=} \mathrm{tr}\Bigg[\underbrace{\Big(\prod_{w\in\mathcal{F}_W} \bar{\pi}^w_{\beta_w} \bigotimes_{p\in\mathcal{P}_B} \Pi_p^{\vec{\alpha}} \prod_{w\in\mathcal{F}_W} \bar{\pi}^w_{\beta_w}\Big)}_{(*)} \Big(\prod_{v\in\mathcal{F}_B\setminus\mathcal{F}_W} \pi^v_{\alpha_v} \bigotimes_{p\in\mathcal{P}_W} \Pi_p^{\vec{\beta}} \prod_{v\in\mathcal{F}_B\setminus\mathcal{F}_W} \pi^v_{\alpha_v}\Big) \Bigg] \\
&\overset{(B)}{=} \prod_{v\in\mathcal{F}_B\cap\mathcal{F}_W} \mathrm{tr}\big[\pi^v_{\alpha_v}\bar{\pi}^v_{\beta_v}\big] \;\times\; \mathrm{tr}\left[\Big(\bigotimes_{p\in\mathcal{P}_B} \rho_p\Big)\Big(\bigotimes_{p\in\mathcal{P}_W} \rho_p\Big) \right] . \qquad (7)
\end{aligned}$$

Here, we have used in step (A) that for all $v \in \mathcal{F}_W$,

$$\bigotimes_{p\in\mathcal{P}_W} \Pi_p^{\vec{\beta}} = \bar{\pi}^v_{\beta_v} \Big(\bigotimes_{p\in\mathcal{P}_W} \Pi_p^{\vec{\beta}}\Big) \bar{\pi}^v_{\beta_v}$$

and anologously for the black layer; in step (B), we have defined new effective plaquette terms ρ_p by virtue of

$$\bigotimes_{p\in\mathcal{P}_B} \rho_p = \mathrm{tr}_{\mathcal{F}_B\cup\mathcal{F}_W}\left[\Big(\prod_{w\in\mathcal{F}_W\setminus\mathcal{F}_B} \bar{\pi}^w_{\beta_w} \bigotimes_{p\in\mathcal{P}_B} \Pi_p^{\vec{\alpha}} \prod_{w\in\mathcal{F}_W\setminus\mathcal{F}_B} \bar{\pi}^w_{\beta_w}\Big) \right] ,$$

and correspondingly for the white plaquettes; the first factor in (7) takes care of the terms in $\mathcal{F}_B \cap \mathcal{F}_W$ in $(*)$. Note that the ρ_p are now only supported on those vertices not in $\mathcal{F}_B \cup \mathcal{F}_W$, as those have been traced out.

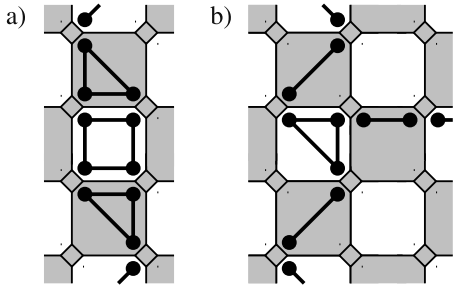

Fig. 1. Computing the overlap Eq. (8) for qubits. The diamonds at the vertices of the square lattice denote the qubits. The connected black dots mark qubits on which the ρ_p act non-trivially on plaquette p. If the ρ_p of two plaquettes act non-trivially on the same qubit, we say that they "overlap"; note that this cannot happen for diagonally adjacent plaquettes. Overlapping ρ_p's form structures which we need to contract to compute the overlap Eq. (8). **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 show that such structures cannot occur, by proving that for any plaquette p, ρ_p can overlap non-trivially at most with two adjacent $\rho_{p'}$'s.

The task of checking whether $\Omega(\vec{\alpha},\vec{\beta}) \neq 0$ has thus been reduced to checking this for (7): For the first term, this can be clearly done efficiently, and so, it remains to prove that the overlap

$$\Theta := \mathrm{tr}\left[\Big(\bigotimes_{p\in\mathcal{P}_B} \rho_p\Big)\Big(\bigotimes_{p\in\mathcal{P}_W} \rho_p\Big) \right] \tag{8}$$

can be computed efficiently. Note that the ρ_p are now supported on plaquettes of a square lattice with vertices missing. Moreover, while the ρ_p do no longer commute, in each layer at most one ρ_p acts non-trivially on each vertex; we will make use of this fact repeatedly in the following.

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; in fact, e.g. quantum circuits, and even postselected quantum circuits, can be encoded this way, making such contractions in general a computationally hard task.

However, as we will show in the following, the structures formed by the ρ_p in Θ, Eq. (8), will always be one-dimensional, and thus Θ can be computed efficiently. To this end, we will 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 on the right of Fig. 1. We will make intensive use of the fact 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 will highlight this fact by placing a cross opposite of any dot:

(9)

Here, the dot indicates that the corresponding ρ_p acts non-trivially on a vertex qubit (the diamond in the center), while the cross indicates that the corresponding ρ_p does act trivially. Note that 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.

The simplest case is when the state ρ_C on the central plaquette C involves only two vertices non-trivially, for instance

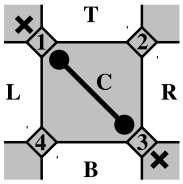

Now, both qubits 1 and 3 can be acted upon non-trivially by at most one white plaquette – the other has to be empty, following the rule (9) that opposite of any dot there has to be a cross; this way, only one-dimensional structures can arise:

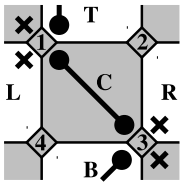

This clearly holds for any possible ρ_C which acts non-trivially on only two qubits, and for any configuration of the adjacent plaquettes; it follows that only one-dimensional structures can emerge this way.

In order to understand the cases where ρ_C acts non-trivially on three or four qubits, let us first analyze the following situation:

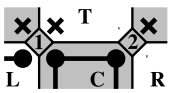

Here, ρ_C acts non-trivially at least on qubits 1 and 2, and ρ_L acts non-trivially on qubit 1 which implies that ρ_T acts trivially on it. In the following, we will show that this implies that ρ_T also has to act trivially on qubit 2. We will prove this by contradiction, so assume that ρ_T acts non-trivially on qubit 2. Since ρ_T is obtained from the original projector Π_T on that plaquette by a partial projection on some of the other vertices, this implies that Π_T acts

non-trivially on qubit 2 (where it spans the whole C^*–algebra, since we have traced out all vertices where this was not the case). On the other hand, ρ_L and thus Π_L acts non-trivially on qubit 1, and thus spans the whole C^*–algebra on it. Since Π_L and Π_T commute, this means that Π_T acts trivially on qubit 1; that is, Π_T and Π_C need to commute on qubit 2 alone: However, since Π_T spans the whole C^*–algebra on qubit 2, this would imply that Π_C and thus ρ_C has to act trivially on qubit 2, giving a contradiction. Thus, we have the following "Lemma":

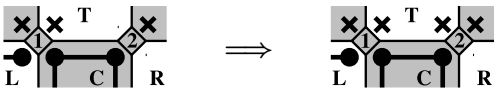

(10)

Let us now consider the case where the state on the central plaquette involves all four qubits non-trivially, and let us start by assuming w.l.o.g. that ρ_L acts non-trivially on qubit 1:

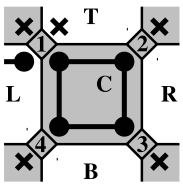

This implies that qubit 1 is not acted upon by ρ_T. Using Eq. (10), we infer that ρ_T cannot involve qubit 2 either,

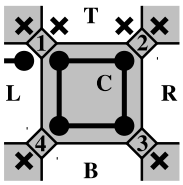

Eq. (10) also shows that ρ_B has to act trivially on qubit 4 – otherwise, ρ_L would act trivially on qubit 4 and thus qubit 1, which it doesn't:

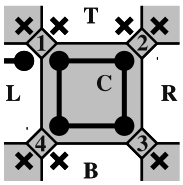

In order to obtain a branching structure, both ρ_B and ρ_R need to act non-trivially on some of the qubits. However, if ρ_B acts non-trivially on qubit 3, Eq. (10) implies that ρ_R has to act trivially on qubits 2 and 3:

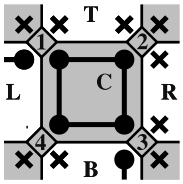

This shows that for ρ_C acting non-trivially on all four qubits, the central plaquette can only couple to at most two adjacent plaquettes, forming one-dimensional structures.

It remains to study what happens in the case of tripartite entanglement on the central plaquette. We start from the following configuration:

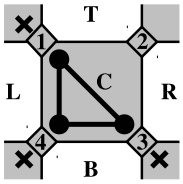

(Note that we don't make any assumptions on how the Hilbert space of qubit 2 decomposes.) Clearly, in order to obtain a branching structure, either ρ_T has to act non-trivially on qubit 1, or ρ_R has to act non-trivially on qubit 3. We consider w.l.o.g. the first possibility and infer from Eq. (10) that ρ_L has to act trivially on both qubits 1 and 4:

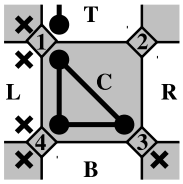

In order to obtain a branching structure, we now have to get both ρ_R and ρ_B involved. However, if ρ_R acts non-trivially on qubit 3, Eq. (10) implies that ρ_B has to act trivially on both qubits 3 and 4, and thus, the structure formed around ρ_C will again be one-dimensional.

Together, this shows that the overlap Θ, Eq. (8), decays into one-dimensional structures for which the overlap can be computed efficiently. In turn, this 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.

3.5 Finite accuracy

In the preceding proof, we have assumed infinite accuracy, but as we will now show, our argument still applies if we compute with finite accuracy. To this end, let N denote the number of qubits in the system; we will need to show that the computation time scales as poly(N). We assume that the Hamiltonian terms are given exactly and can be represented with poly(N) digits. First, note that the trace in Eq. (1), which equals the sum in Eq. (5), evaluates to an integer, and thus, there exists at least one pair $(\vec{\alpha},\vec{\beta})$ such that $\Omega(\vec{\alpha},\vec{\beta}) \geq 2^{-2N}$. If we request this particular $(\vec{\alpha},\vec{\beta})$ as a certificate, it is sufficient if we can determine $\Omega(\vec{\alpha},\vec{\beta})$ to $2N+1 = \text{poly}(N)$ digits. (This is crucial, since there can be $(\vec{\alpha},\vec{\beta})$ for which $\Omega(\vec{\alpha},\vec{\beta})$ is arbitrarily small.) $\Omega(\vec{\alpha},\vec{\beta})$ is obtained from contracting a polynomial number of terms which are either Π_p (which are known exactly) or $\pi^v_{\alpha_v}$ and $\bar{\pi}^v_{\beta_v}$, and the latter can be determined to poly(N) accuracy from the C^*–decomposition (2), which is the solution to a (fixed-size) eigenvalue problem. It follows that $\Omega(\vec{\alpha},\vec{\beta})$ can be computed to the required poly(N) accuracy in poly(N) time, and our proof still applies.

4 Summary

We have studied the COMMUTING HAMILTONIAN problem on a square lattice of qubits with plaquette-wise interaction and shown that the problem is NP–complete. Differently speaking, we have shown that there exists a classical certificate for the fact that the ground state of the system minimizes each term locally which can be checked efficiently on a classical computer. The central idea for our proof has been to split the system into two layers in each of which the commuting terms overlap on individual sites, 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. Each layer could be decomposed using the C*–algebraic techniques introduced to the problem in [10], allowing to find an efficient description of its ground state subspace. Finally, we showed that the overlap of ground states of two layers can be computed efficiently, by showing that it gives rise to of one-dimensional structures only. 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 [12].

Our method does, in principle, also apply beyond qubits: We can still split the system into two layers, decompose both of them into direct sum slices $\vec{\alpha}$ and $\vec{\beta}$, and ask the prover to provide labels $\vec{\alpha}$ and $\vec{\beta}$ with non-zero overlap $\Omega(\vec{\alpha},\vec{\beta})$. While we cannot make sure any more that $\Omega(\vec{\alpha},\vec{\beta})$ can be computed efficiently, we can always ask the prover to also provide us with an instruction on how to efficiently contract the states, in case there is a way to do so, e.g. by providing the optimal contraction order. 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 [10] can be applied to each of them, and in fact to any type of Hamiltonian which is composed of two layers with eigenbases for the zero-energy subspace whose overlap can be computed efficiently, such as for product bases. Note that on the other hand, a decomposition into three layers cannot be used for our purposes, since for three positive operators A, B, and C, $\mathrm{tr}[ABC]$ can have both real and imaginary parts of either sign, so that Eq. (5) is no longer equivalent to Eq. (6) being non-zero (as the $\Pi_p^{\vec{\alpha}}$ and $\Pi_p^{\vec{\beta}}$ do not commute any more).

An interesting open question relating to the present approach to the problem is whether it can be generalized beyond qubits. For four-level systems and beyond, this is likely not the case, since the local Hilbert space can decompose into two qubits, and thus operators commuting on a single spin can both act non-trivially on it, i.e., Eq. (9) does not hold any more. On the other hand, for qutrits this is not the case once we have fixed a slice in the direct sum; yet, it is not clear how to establish a version of Eq. (10). In particular, the non-trivial projections $\pi_{\alpha_v}^v$ can now have both rank 1 and 2, and in the latter case we cannot simply trace out the corresponding degree of freedom; it is however not clear that this does rule out an analogue to Eq. (10).

Acknowledgements

We acknowledge helpful conversations with Dorit Aharonov, Sergey Bravyi, Lior Eldar, Tobias

Osborne, and Volkher Scholz. This work has been supported by the Gordon and Betty Moore Foundation through Caltech's Center for the Physics of Information, the NSF Grant No. PHY-0803371, and the ARO Grant No. W911NF-09-1-0442.

References

1. F. Barahona, *On the computational complexity of Ising spin glass models*. J. Phys. A **15**, 3241 (1982).
2. A. Y. Kitaev, A. H. Shen, and M. N. Vyalyi, *Classical and quantum computation*, American Mathematical Society, Providence, Rhode Island, 2002.
3. D. Aharonov and T. Naveh, *Quantum NP - A Survey*. quant-ph/0210077 (2002).
4. R. Oliveira and B. M. Terhal, *The complexity of quantum spin systems on a two-dimensional square lattice*. Quant. Inf. Comput. **8**, 900 (2009); quant-ph/0504050.
5. D. Aharonov, D. Gottesman, S. Irani, and J. Kempe, *The power of quantum systems on a line*. Commun. Math. Phys. **287**, 41 (2009); arXiv:0705.4077.
6. S. Bravyi, D. P. DiVincenzo, R. I. Oliveira, and B. M. Terhal, *The Complexity of Stoquastic Local Hamiltonian Problems*. Quant. Inf. Comput. **8**, 361 (2008); quant-ph/0606140.
7. S. Bravyi and B. Terhal, *Complexity of stoquastic frustration-free Hamiltonians*. SIAM J. Comput. **39**, 1462 (2009); arXiv:0806.1746.
8. A. Kitaev, *Fault-tolerant quantum computation by anyons*. Ann. Phys. **303**, 2 (2003); quant-ph/9707021.
9. M. A. Levin and X.-G. Wen, *String-net condensation: A physical mechanism for topological phases*. Phys. Rev. B **71**, 045110 (2005); cond-mat/0404617.
10. S. Bravyi and M. Vyalyi, *Commutative version of the k-local Hamiltonian problem and common eigenspace problem*. Quant. Inf. Comput. **5**, 187 (2005); quant-ph/0308021.
11. D. Aharonov and L. Eldar, *On the complexity of Commuting Local Hamiltonians, and tight conditions for Topological Order in such systems*. arXiv:1102.0770 (2011).
12. S. Bravyi, M. B. Hastings, and F. Verstraete, *Lieb-Robinson bounds and the generation of correlations and topological quantum order*. Phys. Rev. Lett. **97**, 050401 (2006); quant-ph/0603121.
13. F. Verstraete and J. I. Cirac, *Valence Bond Solids for Quantum Computation*. Phys. Rev. A **70**, 060302 (2004); quant-ph/0311130.
14. N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, *The computational complexity of PEPS*. Phys. Rev. Lett. **98**, 140506 (2007); quant-ph/0611050.