2010

Quantum-Merlin-Arthur-complete problems for stoquastic Hamiltonians and Markov matrices

Stephen P. Jordan

Peter John Love
Haverford College

David Gosset

Quantum-Merlin-Arthur–complete problems for stoquastic Hamiltonians and Markov matrices

Stephen P. Jordan

Institute for Quantum Information, California Institute of Technology, Pasadena, California 91125, USA

David Gosset

Center for Theoretical Physics, Massachusetts Institute of Technology, 77 Massachusetts Avenue, 6-304,
Cambridge, Massachusetts 02139, USA

Peter J. Love

Department of Physics, Haverford College, 370 Lancaster Avenue, Haverford, Pennsylvania 19041, USA, and
Institute for Quantum Information, California Institute of Technology, Pasadena, California 91125, USA

We show that finding the lowest eigenvalue of a 3-local symmetric stochastic matrix is Quantum-Merlin-Arthur-complete (QMA-complete). We also show that finding the highest energy of a stoquastic Hamiltonian is QMA-complete and that adiabatic quantum computation using certain excited states of a stoquastic Hamiltonian is universal. We also show that adiabatic evolution in the ground state of a stochastic frustration-free Hamiltonian is universal. Our results give a QMA-complete problem arising in the classical setting of Markov chains and adiabatically universal Hamiltonians that arise in many physical systems.

I. INTRODUCTION

Quantum complexity theory is the study of the capabilities and limitations of computational devices operating according to the principles of quantum mechanics [1]. Because many of the classical constructs of computer science (e.g., circuits and clauses) are replaced by matrices, quantum complexity theory is sometimes referred to as *matrix-valued* complexity theory [2,3]. In addition to its intrinsic interest, this subject has many connections to issues of practical relevance to physical science, such as the difficulty with computing properties of quantum systems using either quantum or classical devices [4–6].

Perhaps the most basic classical complexity classes are P, the class of problems solved by a deterministic Turing machine in polynomial time, and NP (nondeterministic polynomial), the class of problems whose verification lies in P. It is widely believed, but not proven, that NP is strictly larger than P [7].

Because quantum mechanics only predicts probabilities of events, the classical deterministic classes are not the most natural place to start if one seeks their quantum generalizations. The probabilistic generalization of P is bounded-error probabilistic polynomial time (BPP), those problems solvable by a probabilistic Turing machine in polynomial time with bounded error [8]. The quantum generalization of this class is bounded-error quantum polynomial time (BQP), the class of problems solvable in polynomial time with bounded error on a quantum computer [1].

The classical probabilistic generalization of NP is the class MA [9]. This generalizes NP to problems whose verification is in BPP. MA stands for Merlin-Arthur. Merlin, who is computationally unbounded but untrustworthy, provides a proof that Arthur can verify using his BPP machine. The class MA possesses a quantum generalization to quantum-Merlin-Arthur (QMA) [10–12]. QMA may be intuitively understood as the class of decision problems that can be efficiently verified by a quantum computer.

Given a classical description of a decision problem x of length n, the prover, Merlin, provides a witness state $|\psi\rangle$ to the verifier, Arthur. Arthur then peforms a poly(n)-time quantum computation on the witness $|\psi\rangle$ and either accepts or rejects. A problem is contained in QMA if, for all YES instances, there exists a witness causing Arthur to accept with probability greater than 2/3 and, for NO instances, there does not exist any witness that causes Arthur to accept with probability greater than 1/3. A problem X is said to be QMA-complete if it is contained in QMA and every problem in QMA can be converted to an instance of X in classical polynomial time.

Let us consider the following question: What is the ground-state energy of a quantum system? This question lies at the core of many areas of physical science, including electronic-structure theory and condensed-matter physics. In quantum complexity theory, this has been formalized (originally by Kitaev [11]; see also, for example, [13]) as the k-local Hamiltonian problem. For some systems, complexity-theoretic arguments suggest that efficient computation of the ground-state energy is likely to remain beyond reach [11,14].

A Hamiltonian H, acting on n qubits, is said to be k-local if it is of the form

$$H = \sum_s H_s,$$

where each H_s acts on at most k qubits. Thus, for example, 1-local Hamiltonians consist only of external fields acting on individual qubits, and 2-local Hamiltonians consist of 1-local terms and pairwise couplings between qubits. Physically realistic Hamiltonians are usually k-local with small k, often 2 or 1, and each local term has bounded norm. Note that this notion of locality has nothing to do with spatial locality; a 2-local Hamiltonian may have long-range couplings, but they must be pairwise.

Problem: k-local Hamiltonian

Input. We are given a classical description of a k-local Hamiltonian H on n qubits $H = \sum_{j=1}^{r} H_j$ with $r = \text{poly}(n)$. Each H_j acts on at most k qubits and has $O(1)$ operator norm. In addition, we are given two constants a and b such that $0 \leqslant a \leqslant b$, and $b - a = \epsilon > 1/\text{poly}(n)$.

Output. If H has an eigenvalue $\leqslant a$, answer YES. If all eigenvalues of H are $>b$, answer NO.

Promise. The Hamiltonian is such that it will produce either YES or NO.

Perhaps a more obvious formulation of this problem is to ask for an approximate ground-state energy to within $\pm \epsilon$ of the correct answer. However, if one can decide the answer to k-local Hamiltonian in polynomial time, then one can solve the approximation version in polynomial time with a binary search. Thus, the approximation problem is of equivalent difficulty to the "decision" version to within a polynomial factor.

The problem k-local Hamiltonian is QMA-complete for $k \geqslant 2$ [13]. The k-local Hamiltonian problem is specified by the matrix elements of the local terms of H. YES instances possess the ground state as a witness. The verification circuit is the phase estimation algorithm—a suitably formalized version of the notion of energy measurement [11]. If the lowest eigenvalue of H is less than a (a YES instance), then Arthur will accept the ground state as a witness. However, if the lowest eigenvalue of H is greater than b (a NO instance), then Merlin cannot supply any eigenstate or superposition of eigenstates that will result in a measurement of energy less than b.

It is considered unlikely that QMA $\subseteq$ BQP, and therefore it is probably impossible to construct a general quantum (or classical) algorithm that finds ground-state energies in polynomial time. However, many Hamiltonians studied in practice have additional restrictions beyond k-locality. In particular, many physical systems are *stoquastic*, meaning that all of their off-diagonal matrix elements are nonpositive in the standard basis. This includes the ferromagnetic Heisenberg model, the quantum transverse Ising model, and most Hamiltonians achievable with Josephson-junction flux qubits [4]. In [4] it was shown that for any fixed k, stoquastic k-local Hamiltonian is contained in the complexity class AM. Thus, unless QMA $\subseteq$ AM (which is believed to be unlikely), stoquastic k-local Hamiltonian is not QMA-complete.[1] It was also shown in [4] that, for any fixed k, adiabatic quantum computation in the ground state of a k-local stoquastic Hamiltonian can be simulated in BPP$_{\text{path}}$. Thus, unless BQP $\subseteq$ BPP$_{\text{path}}$ (which is also believed to be unlikely), such quantum computation is not universal. The work of [2] also defines a random stoquastic local Hamiltonian problem which is complete for the class MA.

These results were tightened further for stoquastic frustration-free (SFF) Hamiltonians in [3]. A local Hamiltonian is frustration free if it can be written as a sum of terms

$$H = \sum_{s}^{m} H_s, \qquad (1)$$

such that

(i) each local operator H_s is positive semidefinite;

(ii) the ground state $|\psi\rangle$ of H satisfies $H_s|\psi\rangle = 0$ for each $s \in \{1, \ldots, m\}$.

The work of [3] showed that an adiabatic evolution along a path composed entirely of SFF Hamiltonians may be simulated by a sequence of classical random walks—that is, the adiabatic evolution may be simulated in the complexity class BPP.

These results were extended to the quantum k-satisfiability (k-SAT) problem in [2,3]. The quantum k-SAT problem was defined in [16] and we reproduce the definition here.

Problem: Quantum k-SAT

Input. A set of k-local projectors $\{\Pi_q\}$ for $q \in \{1, \ldots, m\}$, where $m = \text{poly}(n)$ and a parameter $\tilde{\epsilon} > 1/\text{poly}(n)$.

Output. If there is a state $|\phi\rangle$ such that $\Pi_1|\phi\rangle = 0$ for each $q \in \{1, \ldots, M\}$, then this is a YES instance. If every state $|\phi\rangle$ satisfies

$$\sum_{q=1}^{M} \langle\phi|\Pi_q|\phi\rangle \geqslant \tilde{\epsilon},$$

then it is a NO instance.

Promise. The instance is either YES or NO.

In [2] the stoquastic restriction of quantum k-SAT was shown to be contained in MA for any constant k, and MA-complete for $k = 6$—the first nontrivial example of an MA-complete problem. In [3] these results were extended to a simplified form of stoquastic quantum k-SAT in which projectors Π_a all have matrix elements taken from the set $\{0, 1/2, 1\}$, and the stoquastic constraints which appear as terms in the Hamiltonian are of the form $H_a = \mathbb{1} - \Pi_a$.

The main intuition behind these results is that, by the Perron-Frobenius theorem, the ground state of a stoquastic Hamiltonian consists entirely of real positive amplitudes (given the appropriate choice of global phase). Thus, the ground state is proportional to a classical probability distribution. For this reason, ground-state properties are amenable to classical random-walk algorithms and certain problems such as stoquastic k-local Hamiltonian fall into classical probabilistic complexity classes such as AM. Diffusion quantum Monte Carlo calculations for stoquastic Hamiltonians do not suffer from the sign problem because the negativity of the nonzero off-diagonal matrix elements guarantees that the transition probabilities in the associated random walk are all positive.

In this article we first demonstrate that stoquastic Hamiltonians may be constructed which allow universal adiabatic quantum computation in a subspace. Then we show that the 3-local Hamiltonian problem is QMA-complete when restricted to stochastic Hamiltonians. These are Hamiltonians in which all matrix elements are real and non-negative, and the sum of matrix elements in any row or column is 1. Hence determining the lowest eigenstate of a symmetric stochastic matrix is QMA-hard. If H is a stochastic Hamiltonian, then $-H$ is stoquastic. Thus, our result also shows that determination of the highest-lying eigenstate of a stoquastic matrix is QMA-hard, sharpening the intuition that it is the positivity of the ground state which causes its local Hamiltonian problem to fall into a classical class. We then show that universal

---

[1]Like MA, the class AM is a probabilistic generalization of NP, see for example [15].

adiabatic quantum computation is possible in the ground state of a stochastic frustration-free Hamiltonian. Defining the computational problem stochastic k-SAT in analogy to the definition of stoquastic k-SAT given in [3], we show that this problem is QMA$_1$-complete for $k = 6$. (QMA$_1$ is a slight variant of QMA such that in YES instances, Arthur can be made to accept with probability one [16].)

II. QMA-COMPLETENESS AND ADIABATIC UNIVERSALITY OF STOQUASTIC HAMILTONIANS

We start with the result of [17], which shows that for a Hamiltonian of the form

$$H_{XZ} = \sum_i d_i X_i + \sum_i h_i Z_i + \sum_{i,j} K_{ij} X_i X_j + \sum_{i,j} J_{ij} Z_i Z_j,$$

(2)

the 2-local Hamiltonian problem is QMA-complete if the coefficients d_i, h_i, K_{ij}, and J_{ij} are allowed to have both signs. Furthermore, time-dependent Hamilonians that take the form H_{XZ} at all times can perform universal adiabatic quantum computation [17].

Starting with a Hamiltonian of the form H_{XZ} on n qubits, we can eliminate the negative matrix elements in each term using a technique from [18]. Essentially, the idea is that instead of representing the group Z_2 by $\{1, -1\}$ we use its regular representation:

$$\left\{ \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \right\}.$$

H_{XZ} can be rewritten as

$$H_{XZ} = -\sum_k \alpha_k T_k,$$

(3)

where each coefficient α_k is positive and for each k, T_k is one of

$$\pm X, \quad \pm Z, \quad \pm X_i X_j, \quad \pm Z_i Z_j,$$

(4)

with identity acting on the remaining qubits. For any k, T_k is a $2^n \times 2^n$ matrix in which each entry is either $+1$, -1, or 0. From T_k we construct a $2^{n+1} \times 2^{n+1}$ matrix $\widetilde{T}_k$ by making the following replacements:

$$1 \rightarrow \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad -1 \rightarrow \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad 0 \rightarrow \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}.$$

(5)

We can interpret $\widetilde{T}_k$ as acting on $n + 1$ qubits. The 2×2 matrices of (5) act on the ancilla qubit that has been added. Each T_k is 2-local or 1-local; thus, each corresponding $\widetilde{T}_k$ is 3-local or 2-local. Furthermore, each $\widetilde{T}_k$ is a permutation matrix. Let

$$\widetilde{H}_{XZ} = -\sum_k \alpha_k \widetilde{T}_k.$$

(6)

This is a linear combination of permutation matrices with negative coefficients. By construction, $\widetilde{H}_{XZ}$ is therefore a 3-local stoquastic Hamiltonian. We can rewrite $\widetilde{H}_{XZ}$ as

$$\widetilde{H}_{XZ} = H_{XZ} \otimes |-\rangle\langle-| - \bar{H}_{XZ} \otimes |+\rangle\langle+|,$$

(7)

where

$$\bar{H}_{XZ} = \sum_k \alpha_k |T_k|,$$

(8)

$|T_k|$ is the entrywise absolute value of T_k, and

$$|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$$

$$|-\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle).$$

The projectors $|-\rangle\langle-|$ and $|+\rangle\langle+|$ act on the ancilla qubit.

Equation (7) makes the relationship between the spectra of H_{XZ} and $\widetilde{H}_{XZ}$ clear. Let $|\psi_0\rangle, |\psi_1\rangle, \ldots, |\psi_{N-1}\rangle$ denote the eigenstates of H_{XZ} with corresponding eigenvalues $\lambda_0 \leqslant \lambda_1 \leqslant \cdots \leqslant \lambda_{N-1}$, and let $|\bar{\psi}_0\rangle, |\bar{\psi}_1\rangle, \ldots, |\bar{\psi}_{N-1}\rangle$ denote the eigenstates of $\bar{H}_{XZ}$ with corresponding eigenvalues $\bar{\lambda}_0 \leqslant \bar{\lambda}_1 \leqslant \cdots \leqslant \bar{\lambda}_{N-1}$. ($H_{XZ}$ acts on n qubits, so $N = 2^n$.) $\widetilde{H}_{XZ}$, which acts on a $2N$-dimensional Hilbert space, has two N-dimensional invariant subspaces. The first is spanned by $|\psi_j\rangle|-\rangle$ with eigenvalues λ_j. The second is spanned by eigenvectors $|\bar{\psi}_j\rangle|+\rangle$ with eigenvalues $-\bar{\lambda}_j$.

We can perform universal adiabatic quantum computation in such an eigenstate of a stoquastic Hamiltonian. To prove this, we make use of the universal adiabatic Hamiltonian $H_{XZ}(t)$ from [17], which at all t takes the form shown in Eq. (2). One can use the construction described previously to obtain a stoquastic Hamiltonian $\widetilde{H}_{XZ}(t)$ corresponding to each instantaneous Hamiltonian $H_{XZ}(t)$. In this way we obtain a time-varying Hamiltonian $\widetilde{H}_{XZ}(t)$ whose spectrum in the $|-\rangle$ subspace exactly matches the spectrum of $H_{XZ}(t)$, the only difference being the addition of an ancilla qubit in the $|-\rangle$ state. Because $\widetilde{H}_{XZ}(t)$ has no coupling between the $|-\rangle$ subspace and the $|+\rangle$ subspace, the adiabatic theorem may be applied within the $|-\rangle$ subspace. The relevant eigenvalue gap is thus the same as that of $H_{XZ}(t)$, and so is the run time.

In standard adiabatic quantum computation, the qubits are in the ground state of the instantaneous Hamiltonian. Thus, any disturbance to the state costs energy. This is thought to offer some protection against thermal noise [19]. When performing universal adiabatic quantum computation with $\widetilde{H}_{XZ}(t)$, the qubits are not in the ground state. Thus, it is possible for the system to thermally relax out of the computational state. However, this can only occur by disturbing the ancilla qubit out of the state $|-\rangle$. By protecting the ancilla qubit, one can to a large degree protect the entire computation. Note that an energy penalty against the ancilla qubit leaving the state $|-\rangle$ would be nonstoquastic. This is why the preceding construction fails to prove QMA-completeness and universal adiabatic quantum computation using the *ground state* of a stoquastic Hamiltonian, as we expect it must, based on the complexity-theoretic results of [2–4,20].

III. QMA-COMPLETE PROBLEMS FOR MARKOV MATRICES

The second main result of our article provides an example of a QMA-complete classical problem: finding the lowest eigenvalue of a symmetric Markov matrix. A matrix with all non-negative entries, such that the entries in any given

column sum to 1 is called a stochastic or Markov matrix. These matrices are named after Markov chains, which are stochastic processes such that, given the present state, the future state is independent of the past states. Suppose a system has d possible states. Then, its probability distribution at time t is described by the d-dimensional vector x_t whose entries are non-negative and sum to 1. If the system is evolving according to a Markov process, then its dynamics are completely specified by the equation $x_{t+1} = Mx_t$, where M is a $d \times d$ stochastic matrix. Note that, like quantum Hamiltonians, Markov matrices often have a tensor product structure. For example, suppose we have two independent simultaneous Markov chains governed by $x_{t+1} = Mx_t$ and $y_{t+1} = Ny_t$. Then their joint probability distribution z is governed by $z_{t+1} = (M \otimes N)z_t$.

Markov processes for which the Markov matrix is symmetric correspond to random walks on undirected weighted graphs. (Self-loops are allowed and correspond to diagonal matrix elements.) These matrices are doubly stochastic: The sum of the entries in any row or column is 1. By the Perron-Frobenius theorem, the highest eigenvalue of a symmetric stochastic matrix is 1, and the corresponding eigenvector is the uniform distribution. The eigenvalue with next largest magnitude controls the rate of convergence of the process to its fixed point. A symmetric stochastic matrix is Hermitian and therefore one can also think of these matrices as Hamiltonians.

To prove that finding the lowest eigenvalue of a 3-local symmetric stochastic matrix is QMA-complete, we again use a reduction from the QMA-complete H_{XZ} Hamiltonian of [17]. We must take the opposite sign convention from Eq. (3),

$$H_{XZ} = \sum_k \alpha_k S_k, \qquad (9)$$

where the coefficients α_k are the same as before (all positive) and $S_k = -T_k$. Now define

$$\hat{H}_{XZ} = \frac{1}{N} \sum_k \alpha_k \widetilde{S}_k, \qquad (10)$$

where

$$N = \sum_k \alpha_k$$

and $\widetilde{S}_k$ is the permutation matrix obtained by applying the replacement rules (5) to S_k. By construction, $\hat{H}_{XZ}$ is a 3-local, symmetric, doubly stochastic matrix. We can rewrite $\hat{H}_{XZ}$ as

$$\hat{H}_{XZ} = \frac{1}{N}(H_{XZ} \otimes |-\rangle\langle-| + \bar{H}_{XZ} \otimes |+\rangle\langle+|), \qquad (11)$$

where $\bar{H}_{XZ} = \sum_k \alpha_k |S_k|$. Thus, to determine an eigenvalue of H_{XZ} to within $\pm\epsilon$ we must find the corresponding eigenvalue of $\hat{H}_{XZ}$ to within $\pm\epsilon/N$. Because H_{XZ} is a 2-local Hamiltonian on n qubits with coupling strengths of order unity, N is at most $O(n^2)$. Thus, the problem of determining the eigenvalue of $\hat{H}_{XZ}$ corresponding to the ground state of H_{XZ} to polynomial precision is QMA-hard.

To obtain a cleaner QMA-hard problem, we would like to construct a stochastic matrix whose *lowest* eigenvalue is QMA-hard to find. To do this, let

$$H_p = (1 - p)\sigma_{n+1}^+ + p\hat{H}_{XZ}.$$

Here $\sigma_{n+1}^+ = |+\rangle\langle+| = \frac{1}{2}(\mathbb{1} + X_{n+1})$ acts on the ancilla qubit, thereby giving it an energy penalty of size $(1 - p)$ against leaving the state $|-\rangle$. For $0 \leqslant p \leqslant 1$, H_p is a stochastic Hamiltonian. For $p < 1/3$, the energy penalty is large enough that the highest eigenvalue in the $|-\rangle$ subspace lies below the lowest eigenvalue in the $|+\rangle$ subspace. In this case, the lower half of the spectrum of H_p is the spectrum of H_{XZ} scaled by p/N, and the upper half of the spectrum of H_p is the spectrum of $\bar{H}_{XZ}$ scaled by p/N and shifted up by $1 - p$.

Thus, we can obtain the ground energy of H_{XZ} to polynomial precision by computing the lowest eigenvalue of H_p to a higher but still polynomial precision. This reduction proves that finding the lowest eigenvalue of H_p to polynomial precision is QMA-hard. Using the quantum algorithm for phase estimation, one easily shows that the problem of estimating the lowest eigenvalue of H_p is contained in QMA (see [11]). Thus, this problem is QMA-complete.

IV. FRUSTRATION-FREE ADIABATIC COMPUTATION

It was stated in [3] that universal adiabatic quantum computation can be performed in the ground state of a 5-local frustration-free Hamiltonian. Let $U = U_L \cdots U_2 U_1$ be a quantum circuit acting on n qubits with $L = \text{poly}(n)$ gates. Let

$$|\psi_j\rangle = U_j \cdots U_1 |0\rangle^{\otimes n} \qquad (12)$$

be a state of n qubits corresponding to the jth state of the time evolution of a quantum circuit specified by gates U_j and

$$|c_t\rangle = |1^{t+1}0^{L-t}\rangle. \qquad (13)$$

be a state of $L + 1$ clock qubits. Bravyi and Terhal construct a parametrized 5-local Hamiltonian $H(s)$ such that the ground state $|\psi(s)\rangle$ satisfies

$$|\psi(1)\rangle = \frac{1}{\sqrt{L+1}} \sum_{j=0}^{L} |\psi_j\rangle|c_j\rangle,$$

$$|\psi(0)\rangle = \frac{1}{\sqrt{L+1}} \sum_{j=0}^{L} |0^n\rangle|c_j\rangle.$$

We can think of the first register in $|\psi(s)\rangle$ as consisting of "work" qubits on which the computation happens and the second register in $|\psi(s)\rangle$ as being a clock containing a time written in unary.

For $s \in [0, 1]$, the minimal eigenvalue gap between the ground state and first excited state of $H(s)$ is $O(1/L^2)$. By the adiabatic theorem,[2] $1/\text{poly}(L)$ eigenvalue gap ensures that given $|\psi_0\rangle$, one obtains $|\psi_1\rangle$ by applying $H(s)$ and varying s from 0 to 1 over $\text{poly}(L)$ time. By measuring the clock register of $|\psi(1)\rangle$, one obtains the result $|1^{L+1}\rangle$ with probability $1/(L + 1)$. If this result is obtained, one finds the output of the circuit U by measuring the first register of qubits in the computational basis. By repeating this process with $O(L)$ copies of $|\psi(1)\rangle$, one succeeds with high

---

[2]Many versions of the adiabatic theorem have been proven. For one example, see Appendix F of [21].

probability. Alternatively, one can pad the underlying circuit with L identity gates, in which case each trial succeeds with probability $1/2$.

The construction from [3] invokes the fact that the spectrum of $H(1)$ is independent of the form of the gates U_j. By choosing a gate set which is composed of elements of simply connected unitary groups such as SU(2) and SU(4), one may construct a continuous path connecting each gate to the identity and use a single parameter s to transform all gates from the identity to the final circuit at once. The Hamiltonian at $s = 0$ corresponds to the identity circuit, and its ground state is the uniform superposition of the clock states tensored with the initial data on the work qubits. In this ground state, the qubits of the clock register are entangled. It is standard to design adiabatic computations such that the initial Hamiltonian has a product state as its ground state, because such states should be easily produced by cooling or single-qubit measurements. In this section we construct a modified version of the construction from [3] that satisfies this condition and is still frustration free.

Let $c(j)$ indicate the jth clock qubit and let $w(j)$ indicate the jth work qubit. Let

$$H_j^{\text{init}} = |1\rangle\langle1|_{w(j)} \otimes |10\rangle\langle10|_{c(1),c(2)},$$
$$H_j^{\text{clock}} = |01\rangle\langle01|_{c(j-1),c(j)}.$$

For $j \in \{1, \ldots, L-1\}$, define

$$\begin{aligned}
H_j^{\text{prop}}(s) = {} & s|100\rangle\langle100|_{c(j),c(j+1),c(j+2)} \\
& + (1-s)|110\rangle\langle110|_{c(j),c(j+1),c(j+2)} \\
& - \sqrt{s(1-s)}[U_j \otimes |110\rangle\langle100|_{c(j),c(j+1),c(j+2)} \\
& + U_j^\dagger \otimes |100\rangle\langle110|_{c(j),c(j+1),c(j+2)}]
\end{aligned}$$

and let

$$\begin{aligned}
H_L^{\text{prop}}(s) = {} & s|10\rangle\langle10|_{c(L),c(L+1)} + (1-s)|11\rangle\langle11|_{c(L),c(L+1)} \\
& - \sqrt{s(1-s)}[U_L \otimes |11\rangle\langle10|_{c(L),c(L+1)} \\
& + U_L^\dagger \otimes |10\rangle\langle11|_{c(L),c(L+1)}].
\end{aligned}$$

It can be directly verified that each H_j^{clock}, H_j^{init}, and $H_j^{\text{prop}}(s)$ is a projector. Here, for convenience, we define s so that it varies from 0 to $1/2$ rather than from 0 to 1 as is done in [3]. Our frustration-free Hamiltonian is the following sum of projectors:

$$H^{\text{clock}} = |0\rangle\langle0|_{c(0)} + \sum_{j=1}^{L} H_j^{\text{clock}},$$

$$H^{\text{init}} = \sum_{j=1}^{n} H_j^{\text{init}},$$

$$H^{\text{prop}}(s) = \sum_{j=1}^{L} H_j^{\text{prop}}(s),$$

$$H^{\text{FF}}(s) = H^{\text{clock}} + H^{\text{init}} + H^{\text{prop}}(s).$$

If $U_1 \cdots U_L$ are chosen from a universal set of two-qubit gates, then $H^{\text{FF}}(s)$ is an efficient 5-local frustration-free adiabatic quantum computer. To see how this Hamiltonian achieves universal adiabatic computation, we examine the various terms one by one. The ground state of $H^{\text{FF}}(s)$ is

the simultaneous zero eigenspace of H^{clock}, H^{init}, and H^{prop}. H^{clock} commutes with $H^{\text{prop}}(s) + H^{\text{init}}$ and provides an energy penalty of at least unit size if the clock register is not in one of the unary states $|c_t\rangle = |1^{t+1}0^{L-t}\rangle$. Thus, the low-lying spectrum of $H^{\text{FF}}(s)$ is strictly contained in the ground space of H^{clock}.

For any bit string $x \in \{0, 1\}^n$ and integer $j \in \{1, 2, \ldots, L\}$, let

$$\left|\chi_x^j\right\rangle = (U_j U_{j-1} \cdots U_1|x\rangle) \otimes |c_j\rangle,$$

where $|c_j\rangle$ is as defined in Eq. (13). (We also define $|\chi_x^0\rangle = |x\rangle \otimes |c_0\rangle$.) There are $2^n(L + 1)$ such states and they form an orthonormal basis for the ground space of H^{clock}. In this basis, $H^{\text{prop}}(s) + H^{\text{init}}$ takes the block-diagonal form

$$H^{\text{prop}}(s) + H^{\text{init}} = \bigoplus_{x\in\{0,1\}^n} M_x,$$

where

$$M_x = \begin{bmatrix} s + |x| & -b & & & & \\ -b & 1 & -b & & & \\ & -b & 1 & -b & & \\ & & \ddots & \ddots & \ddots & \\ & & & -b & 1 & -b \\ & & & & -b & 1-s \end{bmatrix}$$

is an $L + 1$ by $L + 1$ matrix and $b = \sqrt{s(1-s)}$. Here $|x|$ denotes the Hamming weight of the bit string x. The appearance of $|x|$ is the sole manifestation of H^{init}. The rest of the matrix elements all come from the "hopping" action of $H^{\text{prop}}(s)$.

$M_{00\ldots0}$ has the unique ground state

$$N \sum_{j=0}^{L} r^j |\psi_j\rangle|c_j\rangle, \tag{14}$$

where $|\psi_j\rangle$ and $|c_j\rangle$ are as defined in Eqs. (12) and (13), $r = \sqrt{\frac{s}{1-s}}$, and N is a normalization factor. This constitutes the ground state of $H^{\text{FF}}(s)$. The first excited state of $M_{00\ldots0}$ has energy $1 - 2\sqrt{s(1-s)}\cos(\frac{\pi}{L+1})$. Because of the direct sum structure of $H^{\text{FF}}(s)$, we can apply the adiabatic theorem directly to $M_{00\ldots0}$. The run time of the adiabatic algorithm is thus determined by the gap between the ground and first excited states of $M_{00\ldots0}$. This takes its minumum at $s = 1/2$, where it is equal to $1 - \cos[\frac{\pi}{(L+1)}] = O(1/L^2)$. For questions of fault tolerance, it is also useful to know the eigenvalue gap between the ground and first excited states of the full Hamiltonian $H^{\text{FF}}(s)$. The first excited energy of $H^{\text{FF}}(s)$ is equal to the ground energy of $M_{10\ldots0}$, which is $1 - 2\sqrt{s(1-s)}\cos[\frac{\pi}{2(L+1)}]$. Thus the minumum eigenvalue gap of $H^{\text{FF}}(s)$ occurs at $s = 1/2$ and is equal to $1 - \cos[\frac{\pi}{2(L+1)}] = O(1/L^2)$.

By Eq. (14), the ground state of $H^{\text{FF}}(0)$ is $|000\cdots\rangle \otimes |1000\cdots\rangle$, and the ground state of $H^{\text{FF}}(1/2)$ is the same state $\frac{1}{\sqrt{L+1}}\sum_{j=0}^{L}|\psi_j\rangle|c_j\rangle$ produced by the scheme of [3].

V. STOCHASTIC FRUSTRATION-FREE COMPUTATION

In [3], Bravyi and Terhal showed that adiabatic quantum computation in the ground state of a SFF Hamiltonian can be efficiently simulated by a classical computer. In this section we show that, in contrast, one can perform universal adiabatic quantum computation in the ground state of a *stochastic* frustration-free (StochFF) Hamiltonian $H^{\text{StochFF}}(s)$. [Alternatively, we can view this as computation in the highest energy state of the stoquastic Hamiltonian $-H^{\text{StochFF}}(s)$.]

It has been shown that the two-qubit controlled-NOT (CNOT) gate, together with any one-qubit rotation whose square is not basis preserving, are sufficient to perform universal quantum computation [22]. All matrix elements in these gates are real numbers. If we choose $U_1, \ldots, U_L$ from this gate set, then $H^{\text{FF}}(s)$ is a 5-local real frustration-free Hamiltonian. Examining the construction of Sec. III, one sees that it can be applied to any Hamiltonian with real matrix elements, and it increases the locality by one. This construction also preserves frustration freeness, as we will show in the next paragraph. We can thus use this construction on $H^{\text{FF}}(s)$ to obtain a 6-local stochastic frustration-free Hamiltonian whose ground state is universal for adiabatic quantum computation.

To show that the mapping of Sec. III preserves frustration freeness, consider applying this mapping to a frustration-free local Hamiltonian $H = \sum_{j=1}^{m} H_j$, where $H_j = \sum_k \alpha_k^j S_k^j$ (where each S_k^j is, up to an overall sign, a tensor product of Pauli operators and each α_k^j is positive). We obtain the Hamiltonian

$$
\begin{aligned}
H_p &= p\hat{H} + (1-p)\left(\frac{\mathbb{1} + X_{n+1}}{2}\right) \\
&= \sum_j \frac{N_j}{N}\left[p\hat{H}_j + (1-p)\left(\frac{\mathbb{1} + X_{n+1}}{2}\right)\right], \quad (15)
\end{aligned}
$$

where $N_j = \sum_k \alpha_k^j$ and $N = \sum_j N_j$. When $p < \frac{1}{3}$, H_p is stochastic and has a zero-energy ground state with an eigenvalue gap which is $\frac{p}{N}$ times the gap of H. Furthermore, we see from (15) (and the fact that each H_j is positive semidefinite) that H_p is a sum of positive semidefinite operators. Hence, the Hamiltonian H_p is frustration free.

VI. GENERALIZATIONS

The constructions of Secs. III and II replace Hamiltonians with real matrix elements of both signs by computationally equivalent Hamiltonians with real positive matrix elements. In this section we show that this technique can be generalized to directly replace Hamiltonians with complex matrix elements by computationally equivalent Hamiltonians with only real positive matrix elements. However, in the process we necessarily introduce a twofold degeneracy of the ground state.

Let H be an arbitrary k-local Hamiltonian. We may expand H as

$$
H = \sum_j \alpha_j O_j, \quad (16)
$$

where each O_j is a tensor product of k or fewer Pauli matrices and each α_j is positive. Each entry in each O_j is ± 1 or $\pm i$.

We can replace the group $\{1, i, -1, -i\}$ with its left-regular representation

$$
\begin{aligned}
i &\mapsto F \\
-1 &\mapsto F^2 \\
-i &\mapsto F^3 \\
1 &\mapsto F^4,
\end{aligned} \quad (17)
$$

where

$$
F = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{pmatrix}. \quad (18)
$$

The eigenvectors of F are $|v_0\rangle, |v_1\rangle, |v_2\rangle, |v_3\rangle$, where

$$
|v_j\rangle = \frac{1}{2}\sum_{l=0}^{3} i^{lj}|l\rangle.
$$

The corresponding eigenvalues are

$$
F|v_j\rangle = i^j|v_j\rangle. \quad (19)
$$

Let S_j and A_j be the real and imaginary parts of $\alpha_j O_j$. That is, S_j and A_j are the unique real symmetric and antisymmetric matrices such that

$$
\alpha_j O_j = S_j + i A_j.
$$

Further, let $S_j^+ = (|S_j| + S_j)/2$ and $S_j^- = (|S_j| - S_j)/2$, and similarly for $A_j^\pm$, where $|\cdot|$ denotes the entrywise absolute value. Applying the replacement (17) to H and dividing by $N = \sum_j \alpha_j$ yields the stochastic Hamiltonian $\tilde{H}$ with the decomposition

$$
\begin{aligned}
\tilde{H} = \frac{1}{N}(&H^{(0)} \otimes |v_0\rangle\langle v_0| + H^{(1)} \otimes |v_1\rangle\langle v_1| \\
&+ H^{(2)} \otimes |v_2\rangle\langle v_2| + H^{(3)} \otimes |v_3\rangle\langle v_3|), \quad (20)
\end{aligned}
$$

where

$$
H^{(0)} = \sum_j S_j^+ + S_j^- + A_j^+ + A_j^-,
$$

$$
H^{(1)} = \sum_j S_j^+ - S_j^- + i A_j^+ - i A_j^-,
$$

$$
H^{(2)} = \sum_j S_j^+ + S_j^- - A_j^+ - A_j^-,
$$

$$
H^{(3)} = \sum_j S_j^+ - S_j^- - i A_j^+ + i A_j^-.
$$

$H^{(1)} = H$; thus, the spectrum of $\tilde{H}$ in the $|v_1\rangle$ subspace matches that of H up to a normalization factor of N and a pair of extra ancilla qubits. If we write each projector $|v_j\rangle\langle v_j|$ in terms of the Pauli basis we obtain

$$
\begin{aligned}
|v_0\rangle\langle v_0| &= \Pi_+^X \Pi_+^X, \\
|v_1\rangle\langle v_1| &= \Pi_-^X \Pi_+^Y, \\
|v_2\rangle\langle v_2| &= \Pi_+^X \Pi_-^X, \\
|v_3\rangle\langle v_3| &= \Pi_-^X \Pi_-^Y,
\end{aligned}
$$

where $\Pi_{\pm}^a$ is the projector onto the eigenvalue ± 1 eigenstate of the Pauli matrix a. Thus, an X penalty on the first ancilla qubit will separate the $|v_1\rangle, |v_3\rangle$ subspace from the $|v_0\rangle, |v_2\rangle$ subspace. So, taking, $0 < p < \frac{1}{3}$, the stochastic Hamiltonian

$$H'_p = (1 - p)\left(\frac{\mathbb{1} + X_{n+1}}{2}\right) + p\tilde{H}$$

has ground space spanned by $|\psi^{(1)}\rangle|v_1\rangle$ and $|\psi^{(3)}\rangle|v_3\rangle$, where $|\psi^{(1)}\rangle$ is the ground state of $H^{(1)}$ and $|\psi^{(3)}\rangle$ is the ground state of $H^{(3)}$. $H^{(1)} = H$; thus, $|\psi^{(1)}\rangle$ is the ground state of H. $H^{(3)} = H^*$; thus, $|\psi^{(3)}\rangle$ is the complex conjugate of the ground state of H.

A simple argument shows that the doubling in the spectrum of H'_p is a necessary property for any construction which maps an arbitrary Hamiltonian onto a real Hamiltonian H_R, where H_R is equal to H within a fixed 1D subspace of the ancillas. Suppose that we have such a map which sends an arbitrary Hamiltonian H which acts on a Hilbert space $\mathcal{H}_1$ to a real Hamiltonian H_R on a larger Hilbert space $\mathcal{H}_1 \otimes \mathcal{H}_2$ with the property that

$$H_R = H \otimes |\phi\rangle\langle\phi| + H^{\text{other}} \otimes (\mathbb{1} - |\phi\rangle\langle\phi|), \qquad (21)$$

where the state $|\phi\rangle \in \mathcal{H}_2$ does not depend on the particular Hamiltonian H but the operator H^{other} may depend on H. Then for any eigenvector $|\psi\rangle$ of H with energy E, we have

$$H_R|\psi\rangle|\phi\rangle = E|\psi\rangle|\phi\rangle. \qquad (22)$$

Since H_R is real, complex conjugating this equation gives

$$H_R|\psi^\star\rangle|\phi^\star\rangle = E|\psi^\star\rangle|\phi^\star\rangle. \qquad (23)$$

To show that doubling exists in the spectrum, it is sufficient to show that $\langle\phi|\phi^\star\rangle = 0$. To prove this, first use Eq. (23) to obtain

$$(\mathbb{1} \otimes |\phi\rangle\langle\phi|)H_R|\psi^\star\rangle|\phi^\star\rangle = E|\psi^\star\rangle|\phi\rangle\langle\phi|\phi^\star\rangle. \qquad (24)$$

Then use Eq. (21) to obtain

$$(\mathbb{1} \otimes |\phi\rangle\langle\phi|)H_R|\psi^\star\rangle|\phi^\star\rangle = (H|\psi^\star\rangle)|\phi\rangle\langle\phi|\phi^\star\rangle. \qquad (25)$$

Equating these expressions gives

$$H|\psi^\star\rangle\langle\phi|\phi^\star\rangle = E|\psi^\star\rangle\langle\phi|\phi^\star\rangle. \qquad (26)$$

This must hold for all Hamiltonians H and eigenstates $|\psi\rangle$, and therefore it must be the case that $\langle\phi|\phi^\star\rangle = 0$. So we have shown that the doubling in the spectrum of H'_p is a necessary feature of the type of maps we consider. For constructing universal adiabatic quantum computers, the degeneracy induced by this construction may be problematic. However, for proving complexity-theoretic completeness results, it is often irrelevant, as we see in the next section.

VII. STOCHASTIC k-SAT

The methods of the previous section can be used to show, roughly speaking, that deciding whether or not a Hamiltonian which is a sum of positive semidefinite stochastic operators is frustration free is as difficult as the general problem of deciding whether a Hamiltonian is frustration free. In this section we formalize this by defining a problem called stochastic k-SAT, which we show to be QMA$_1$-complete for $k = 6$.

We first recall the definition of stoquastic k-SAT which is given in [3].

Problem: Stoquastic k-SAT

Input. A set of k-local Hermitian operators $\{H_j\}$ for $j \in \{1, \ldots, m\}$, where $m = \text{poly}(n)$ and a parameter $\epsilon > 1/\text{poly}(n)$:

(i) each H_j is positive semidefinite;

(ii) each H_j has norm which is bounded by a polynomial in n;

(iii) every H_j is stoquastic.

Output. If $H = \sum_j H_j$ has a zero-energy ground state, then this is a YES instance. Otherwise, if every eigenstate of H has energy $> \epsilon$, then it is a NO instance.

Promise. Either the ground state of H has energy 0, or else it has energy $> \epsilon$.

The stoquastic k-SAT problem is therefore the problem of deciding if a given stoquastic Hamiltonian that is a sum of positive definite operators is frustration free, given that either this is the case or else its ground energy exceeds ϵ [3]. Note that this definition of stoquastic k-SAT looks somewhat different from the definition of quantum k-SAT which was given in Sec. I, which was stated entirely in terms of projectors. Given an instance of stoquastic k-SAT, we can define operators Π_j which project onto the zero eigenspaces of the H_j. When the Hamiltonians H_j are stoquastic, these projectors are guaranteed to have non-negative matrix elements in the computational basis [3]. So given an instance of stoquastic k-SAT with Hermitian positive semidefinite operators H_j, it is possible to construct another instance of stoquastic k-SAT with operators $\tilde{H}_j = \{1 - \Pi_j\}$ that are all projectors.

We now define a problem called stochastic k-SAT, which is identical to stoquastic k-SAT except that condition (iii) is replaced by

(iii$'$) Every H_j is a stochastic matrix.

We note that there does not appear to be an equivalence between this definition of stochastic k-SAT and the corresponding definition where all the H_j are (in addition) required to be projectors.

Given these two definitions and the foregoing map from an arbitrary Hamiltonian to a stochastic Hamiltonian, we now show how to reduce any instance of quantum 4-SAT to an instance of stochastic 6-SAT. Starting with an instance of quantum 4-SAT specified by a set of projectors $\{\Pi_j\}$ (for $j \in 1, \ldots, m$), we use the map of the previous section (with $p = \frac{1}{3}$ for concreteness) on each projector to obtain a set of 6-local positive semidefinite stochastic Hamiltonians $\{H_j\}$, where

$$H_j = \frac{2}{3}\left(\frac{\mathbb{1} + X_{n+1}}{2}\right) + \frac{1}{3}\tilde{\Pi}_j. \qquad (27)$$

[Note that $\tilde{\Pi}_j$ refers to the operator obtained by applying the mapping from Eq. (20).] If the 4-SAT instance is satisfiable, then the stochastic 6-SAT instance will also be satisfiable. Define $N_{\max}$ to be the maximum value of N obtained for one of the terms Π_j when using the mapping of Eq. (20). If the 4-SAT instance is not satisfiable, then for any state $|\phi\rangle$ there is some projector Π_k such that $\langle\phi|\Pi_k|\phi\rangle \geqslant \frac{\epsilon}{m}$. If we take the parameter $\tilde{\epsilon}$ of the stochastic 6-SAT instance to be related to the parameter ϵ of the quantum 4-SAT instance

by $\tilde{\epsilon} = \frac{\epsilon}{3mN_{\max}}$, then the stochastic 6-SAT instance will also be unsatisfiable. Therefore, stochastic 6-SAT is QMA$_1$-hard. Stochastic 6-SAT is contained in QMA$_1$ since every instance of stochastic 6-SAT can be mapped to an instance of quantum 6-SAT by taking projectors Π_j which project onto everything but the zero eigenspaces of the H_j.

So QMA$_1$ completeness of stochastic 6-SAT follows from the results of Bravyi [16] on quantum k-SAT. This is in contrast to stoquastic k-SAT, which is contained in MA for every constant k [3].

VIII. QMA-COMPLETENESS FOR EXCITED STATES

The local Hamiltonian problem refers specifically to ground state energies. Similarly, we have formulated a computational problem based on the highest energy of a given Hamiltonian. It is natural to ask about the complexity of estimating the cth excited state. We can formulate this as follows. Let H be a k-local Hamiltonian on the Hilbert space $\mathcal{H}$ of n qubits. Let $\lambda_1 \leqslant \lambda_2 \leqslant \cdots \leqslant \lambda_{2^n}$ denote the eigenvalues of H, with corresponding eigenvectors $|\psi_1\rangle, |\psi_2\rangle, \ldots, |\psi_{2^n}\rangle$. The (k,c,ϵ)-energy problem is as follows.

Problem: (k,c,ϵ)-energy

Input. We are given a classical description of H, an integer $c \geqslant 1$, and a pair of parameters a, b such that $b - a = \epsilon > 1/\text{poly}(n)$.

Output. If $\lambda_c \leqslant a$, answer YES. If $\lambda_c \geqslant b$, output NO.

Promise. H is such that the answer is YES or NO.

In this section we will show that the (k, c, ϵ)-energy problem is QMA-complete for any $c = O(1)$. Showing QMA-hardness is the easier of the two proofs. This can be achieved as follows. Let $H^{(a)}$ and $H^{(b)}$ be a pair of k-local Hamiltonians on n qubits, with spectra $\lambda_1^{(a)}, \ldots, \lambda_{2^n}^{(a)}, |\psi_1^{(a)}\rangle, \ldots, |\psi_{2^n}^{(a)}\rangle$ and $\lambda_1^{(b)}, \ldots, \lambda_{2^n}^{(b)}, |\psi_1^{(b)}\rangle, \ldots, |\psi_{2^n}^{(b)}\rangle$, respectively. Then

$$H^{(ab)} = H^{(a)} \otimes |0\rangle\langle 0| + H^{(b)} \otimes |1\rangle\langle 1|$$

is a $(k+1)$-local Hamiltonian on $n+1$ qubits. Its complete set of eigenvalues is $\lambda_1^{(a)}, \ldots, \lambda_{2^n}^{(b)}$; $\lambda_1^{(b)}, \ldots, \lambda_{2^n}^{(b)}$, with corresponding eigenvectors $|\psi_1^{(a)}\rangle|0\rangle, \ldots, |\psi_{2^n}^{(a)}\rangle|0\rangle; |\psi_1^{(b)}\rangle|1\rangle, \ldots, |\psi_{2^n}^{(b)}\rangle|1\rangle$. To prove QMA-hardness of a low-lying excited state, let H_0 to be a Hamiltonian such that determining whether the ground energy is close to zero is QMA-hard. Given an integer c, let $d = \lceil \log_2 c \rceil$, $P_k = \frac{1}{2}(Z_k + \mathbb{1})$, and

$$H_c = \sum_{k=0}^{d} 2^k P_k + \sum_{k=d+1}^{n} 2^{d+1} P_k - \left(c - \frac{1}{2}\right)\mathbb{1}.$$

H_c has exactly c states with negative energy, and its lowest non-negative eigenvalue is $\frac{1}{2}$. Thus determining the cth excited energy of $H_c \otimes |0\rangle\langle 0| + H_0 \otimes |1\rangle\langle 1|$ is QMA-hard. In particular, it is interesting to note that by choosing $c = 2$ we construct a Hamiltonian whose eigenvalue gap between the ground state and the first excited state is QMA-hard to compute.

Next we show containment in QMA. The naive protocol would be for Merlin to provide Arthur with the state

$|\psi_1\rangle|\psi_2\rangle \ldots |\psi_c\rangle$ and for Arthur to use phase estimation to check that the c registers each contain a state of energy at most a. The problem is that for NO instances there are many ways for Merlin to cheat. For example if $\lambda_1 \leqslant a$ but $\lambda_c \geqslant b$, the answer is NO, but Merlin can provide the state $|\psi_1\rangle^{\otimes c}$ as a supposed witness. To prevent this, Arthur needs to somehow check that he has been given a set of c orthogonal states that each have energy at most a. Thus, we propose the following protocol.

Arthur demands that Merlin give him the state

$$|W\rangle = \frac{1}{\sqrt{c!}} \sum_{\pi \in S_c} \text{sgn}(\pi)|\psi_{\pi(1)}\rangle|\psi_{\pi(2)}\rangle \cdots |\psi_{\pi(c)}\rangle. \quad (28)$$

Arthur performs the projective measurement to see that the state given to him by Merlin lies in the antisymmetric subspace of $\mathcal{H}^{\otimes c}$. If this fails, he rejects. He then throws away all but the first register and performs phase estimation of H to precision better than ϵ. If the state has energy above b, he rejects. Otherwise, he accepts.

It is clear that for YES instances, Arthur will accept the state $|W\rangle$ with high probability. (The only source of error is imprecision in phase estimation.) We will next prove that for NO instances the acceptance probability is at most $1 - \frac{1}{c}$. Using standard methods [11,23,24], we can amplify this protocol to obtain polynomially small acceptance probability for NO instances.

Lemma 1. For any state $|\phi\rangle$ in the antisymmetric subspace of $\mathcal{H}^{\otimes c}$ and any state $|\alpha_1\rangle \in \mathcal{H}$, $\langle\phi|(|\alpha_1\rangle\langle\alpha_1| \otimes \mathbb{1})|\phi\rangle \leqslant \frac{1}{c}$ where $\mathbb{1}$ is the identity operator on $\mathcal{H}^{\otimes(c-1)}$.

Proof. Extend $|\alpha_1\rangle$ to an orthonormal basis $|\alpha_1\rangle, |\alpha_2\rangle, \ldots, |\alpha_{2^n}\rangle$ for $\mathcal{H}$. Let F be the set of functions $f : \{1, 2, \ldots, c\} \to \{1, 2, \ldots, 2^n\}$ such that $f(1) < f(2) < \cdots < f(c)$. Thus, $|F| = \binom{2^n}{c}$. For any $f \in F$ we have the corresponding Slater determinant state:

$$|D_f\rangle = \frac{1}{\sqrt{c!}} \sum_{\pi \in S_c} \text{sgn}(\pi)|\alpha_{f[\pi(1)]}\rangle|\alpha_{f[\pi(2)]}\rangle \ldots |\alpha_{f[\pi(c)]}\rangle.$$

It a standard result that these $\binom{2^n}{c}$ states form a complete orthonormal basis for the antisymmetric subspace of $\mathcal{H}^{\otimes c}$. For any $f, g \in F$, we have

$$\langle D_f|(|\alpha_1\rangle\langle\alpha_1| \otimes \mathbb{1})|D_g\rangle = \frac{\delta_{f,g}\delta_{f(1),1}}{c}, \quad (29)$$

where each δ denotes a generalized Kronecker-δ. Because $|\phi\rangle$ is antisymmetric, it can be decomposed in the Slater determinant basis,

$$|\phi\rangle = \sum_{f \in F} \phi_f |D_f\rangle,$$

and $p_f = \phi_f^* \phi_f$ is a corresponding probability distribution on F. Thus,

$$\langle\phi|(|\alpha_1\rangle\langle\alpha_1| \otimes \mathbb{1})|\phi\rangle = \sum_{f,g \in F} \phi_f^*\langle D_f|(|\alpha_1\rangle\langle\alpha_1| \otimes \mathbb{1})|D_g\rangle\phi_g.$$

By Eq. (29), this is

$$= \frac{1}{c} \sum_{f \in F} p_f \delta_{f(1),1} \leqslant \frac{1}{c}. \qquad \blacksquare$$

The quantity

$$p_j(\phi) = \langle\phi|(|\psi_j\rangle\langle\psi_j| \otimes \mathbb{1})|\phi\rangle$$

is the probability of obtaining $|\psi_j\rangle$ if we measure the first register of a state $|\phi\rangle$ in the eigenbasis of H. By Lemma 1 ,

$$\sum_{j=1}^{c-1} p_j(\phi) \leqslant 1 - \frac{1}{c}.$$

Thus, with probability at least $\frac{1}{c}$, such a measurement would yield $|\psi_j\rangle$, with $j > c - 1$. Thus, if $\lambda_c > b$, then with probability at least $\frac{1}{c}$ a measurement of the observable H would yield energy of at least b. The phase estimation algorithm can in poly($1/\epsilon$) time perform such an energy measurement with an exponentially smaller chance of making an error as large as ϵ. Thus, the protocol is sound, which completes the proof that the (c, k, ϵ)-energy problem is QMA-complete for constant c and k and polynomially small ϵ.

IX. DISCUSSION AND CONCLUSIONS

The results presented in this article have several applications. Although calculating the ground-state energy of stoquastic Hamiltonians appears easier than calculating the ground-state energy of generic Hamiltonians, our results suggest that calculating other eigenstates of stoquastic Hamiltonians remains hard. Because the wave functions of these states have amplitudes which are both positive and negative, the hardness of determining their energy supports the intuition that it is the positivity of the amplitudes which makes the ground-state problem for stoquastic Hamiltonians easier. An extreme distinction between stochastic and stoquastic Hamiltonians arises when the Hamiltonians are also frustration free. Although adiabatic evolution with SFF Hamiltonians is simulable in BPP [3], we have shown that adiabatic evolution in the ground state of a stochastic frustration-free Hamiltonian is universal.

Second, these results may be relevant for the physical implementation of quantum computers. The first proof of universality of adiabatic quantum computation used 5-local interactions [25]. Since then, the Hamiltonians have been brought into incrementally more physically feasible form by various techniques while retaining universality [13,26,27]. The universal Hamiltonian H_{XZ} of [17] is one outcome of this chain of reductions. Here we add one more step to this chain, obtaining universal stochastic and stoquastic Hamiltonians which resemble those arising in some systems of superconducting qubits [4]. The constructions given here are at least 3-local, and so would require the use of perturbative gadgets to implement in terms of physical 2-local interactions.

Finally, our results are of interest from a purely complexity-theoretic point of view. Stochastic matrices arise outside the context of quantum mechanics, in Markov chains. Our reduction shows that finding the lowest eigenvalue of a certain class of exponentially large but efficiently describable doubly stochastic matrices is QMA-complete. (These stochastic matrices correspond to Markov chains in which the "update rule" is a probabilistic selection over some set of updates which are local in the tensor product sense.) In general, the problem of finding eigenvalues of stochastic matrices is of interest because the eigenvalue of second-largest magnitude determines the mixing time of the corresponding Markov chain. There exist Markov chains in which the eigenvalue of second-largest magnitude is negative and is the lowest-lying eigenvalue. We hope that the demonstration of a QMA-complete problem arising in a classical setting will help shed further light on the class QMA itself.

[1] E. Bernstein and U. Vazirani, SIAM J. Comput. **26**, 1411 (1997).

[2] S. Bravyi, A. J. Bessen, and B. M. Terhal, e-print arXiv:quant-ph/0611021.

[3] S. Bravyi and B. Terhal, SIAM J. Comput. **39**, 1462 (2009).

[4] S. Bravyi, D. DiVincenzo, R. Oliveira, and B. Terhal, Quant. Inf. Comp. **8**, 0361 (2008).

[5] Y. Liu, M. Christandl, and F. Verstraete, Phys. Rev. Lett. **98**, 110503 (2007).

[6] N. Schuch and F. Verstraete, Nat. Phys. **5**, 732 (2009).

[7] M. Sipser, *Introduction to the Theory of Computation* (PWS Publishing, Boston, 1997).

[8] L. Adleman, in *FOCS '78: Proceedings of the 19th IEEE Symposium on Foundations of Computer Science* (IEEE Computer Society, 1978), pp. 75–83.

[9] L. Babai, in *STOC '85: Proceedings of the 17th Annual ACM Symposium on Theory of Computing* (ACM, 1985), pp. 421–429.

[10] D. Aharonov and T. Naveh (2002), e-print arXiv:quant-ph/0210077v1.

[11] A. Yu Kitaev, A. H. Shen, and M. N. Vyalyi, *Classical and Quantum Computation* (American Mathematical Society, 2002).

[12] J. Watrous, in *FOCS '00: Proceedings of the 41st IEEE Symposium on Foundations of Computer Science* (IEEE Computer Society, 2000).

[13] J. Kempe, A. Kitaev, and O. Regev, SIAM J. Comput. **35**, 1070 (2004).

[14] F. Barahona, J. Phys. A **15**(10), 3241 (1982).

[15] S. Arora and B. Barak, *Computational Complexity: A Modern Approach* (Cambridge University Press, 2009).

[16] S. Bravyi (2006), e-print arXiv:quant-ph/0602108v1.

[17] J. D. Biamonte and P. Love, Phys. Rev. A **78**, 012352 (2008).

[18] D. Janzing and P. Wocjan (2006), e-print arXiv:quant-ph/0610235.

[19] A. M. Childs, E. Farhi, and J. Preskill, Phys. Rev. A **65**, 012322 (2001).

[20] Y.-K. Liu (2007), e-print arXiv:0712.1388.

[21] S. P. Jordan, Ph.D. thesis, Massachusetts Institute of Technology, 2008.

[22] Y. Shi, Quantum Inf. Comput. **3**, 84 (2003).

[23] C. Marriott and J. Watrous, Comput. Complexity **14**, 122 (2005).

[24] D. Nagaj, P. Wocjan, and Y. Zhang, Quantum Inf. Comput. **9**, 1053 (2009).

[25] D. Aharonov, W. van Dam, J. Kempe, Z. Landau, S. Lloyd, and O. Regev, SIAM J. Comput. **37**, 166 2007.

[26] J. Kempe and O. Regev, Quantum Inf. Comput. **3**, 258 (2003).

[27] R. Oliveira and B. M. Terhal, Quantum Inf. Comput. **8**, 900 (2008).