

Computational Difficulty of Computing the Density of States

Brielin Brown,^{1,2} Steven T. Flammia,² and Norbert Schuch³

¹University of Virginia, Departments of Physics and Computer Science, Charlottesville, Virginia 22904, USA

²Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada

³California Institute of Technology, Institute for Quantum Information, MC 305-16, Pasadena, California 91125, USA

(Received 26 October 2010; revised manuscript received 28 May 2011; published 20 July 2011)

We study the computational difficulty of computing the ground state degeneracy and the density of states for local Hamiltonians. We show that the difficulty of both problems is exactly captured by a class which we call #BQP, which is the counting version of the quantum complexity class quantum Merlin Arthur. We show that #BQP is not harder than its classical counting counterpart #P, which in turn implies that computing the ground state degeneracy or the density of states for classical Hamiltonians is just as hard as it is for quantum Hamiltonians.

DOI: 10.1103/PhysRevLett.107.040501

PACS numbers: 03.67.Ac, 71.20.-b, 89.70.Eg

Understanding the physical properties of correlated quantum many-body systems is a problem of central importance in condensed matter physics. The density of states, defined as the number of energy eigenstates per energy interval, plays a particularly crucial role in this endeavor. It is a key ingredient when deriving many thermodynamic properties from microscopic models, including specific heat capacity, thermal conductivity, band structure, and (near the Fermi energy) most electronic properties of metals. Computing the density of states can be a daunting task, however, as it in principle involves diagonalizing a Hamiltonian acting on an exponentially large space, though other more efficient approaches which might take advantage of the structure of a given problem are not *a priori* ruled out.

In this Letter, we precisely quantify the difficulty of computing the density of states by using the powerful tools of quantum complexity theory. Quantum complexity aims at generalizing the well-established field of classical complexity theory to assess the difficulty of tasks related to quantum mechanical problems, concerning both the classical difficulty of simulating quantum systems as well as the fundamental limits to the power of quantum computers. In particular, quantum complexity theory has managed to explain the difficulty of computing ground state properties of quantum spin systems in various settings, such as two-dimensional (2D) lattices [1] and even one-dimensional (1D) chains [2], as well as fermionic systems [3].

We will determine the computational difficulty of two problems: first, computing the density of states of a local Hamiltonian, and second, counting the ground state degeneracy of a local gapped Hamiltonian; in both cases, the result holds even if the Hamiltonian is restricted to act on a 2D lattice of qubits or on a 1D chain. To this end, we will introduce the quantum counting class #BQP (sharp BQP), which constitutes the natural counting version of the class QMA (quantum Merlin Arthur) which itself captures the difficulty of computing the ground state energy of a local

Hamiltonian [4,5]. Vaguely speaking, #BQP counts the number of possible “quantum solutions” to a quantum problem that can be verified by using a quantum computer. We show that both problems, computing the density of states and counting the ground state degeneracy, are complete problems for the class #BQP; i.e., they are among the hardest problems in this class.

Having quantified the difficulty of computing the density of states and counting the number of ground states, we proceed to relate #BQP to known classical counting complexity classes and show that #BQP equals #P (under weakly parsimonious reductions). Here, the complexity class #P counts the number of satisfying assignments to any efficiently computable Boolean function. This can be a very hard problem which is believed to take exponential time; in particular, it is at least as hard as deciding whether the function has at least one satisfying input, i.e., the complexity class NP. Examples for #P-complete problems (i.e., the hardest problems in that class) include counting the number of colorings of a graph or computing the permanent of a matrix with binary entries. Phrased in terms of Hamiltonians, what we show is that computing the density of states and counting the ground state degeneracy of a classical spin system is just as hard as solving the same problem for a quantum Hamiltonian.

Quantum complexity classes.—Let us start by introducing the relevant complexity classes. The central role in the following is taken by the *verifier* V , which verifies quantum solutions (also called *proofs*) to a given problem. More formally, a verifier checking an n -qubit quantum proof (that is, a quantum state $|\psi\rangle$) consists of a $T = \text{poly}(n)$ length quantum circuit $U = U_T \cdots U_1$ (with local gates U_i) acting on $m = \text{poly}(n)$ qubits, which takes the n -qubit quantum state $|\psi\rangle_I$ as an input, together with $m - n$ initialized ancillas, $|0\rangle_A \equiv |0 \dots 0\rangle_A$, applies U , and finally measures the first qubit in the $\{|0\rangle_1, |1\rangle_1\}$ basis to return 1 (“proof accepted”) or 0 (“proof rejected”). Then, the class QMA contains all problems of the form: “Decide whether

there exists a $|\psi\rangle$ such that $p_{\text{acc}}[V(\psi)] > a$, or whether $p_{\text{acc}}[V(\psi)] < b$ for all $|\psi\rangle$, for some chosen $a - b > 1/\text{poly}(n)$, given that one is the case.” Here, the acceptance probability of a state $|\psi\rangle$ is $p_{\text{acc}}[V(\psi)] := \langle \psi | \Omega | \psi \rangle$, with

$$\Omega = (\mathbb{1}_I \otimes \langle \mathbf{0} |_A) U^\dagger (|1\rangle\langle 1|_I \otimes \mathbb{1}) U (\mathbb{1}_I \otimes | \mathbf{0} \rangle_A), \quad (1)$$

which we illustrate in Fig. 1.

The idea behind this definition is that QMA quantifies the difficulty of computing the ground state energy $E_0(H)$ of a local Hamiltonian H up to $1/\text{poly}(n)$ accuracy. Let the verifier be a circuit estimating $\langle \psi | H | \psi \rangle$; then a black box solving QMA problems can be used to compute $E_0(H)$ up to $1/\text{poly}(n)$ accuracy by a binary search using a single QMA query. Note also that QMA is the quantum version of the class NP, where one is given an efficiently computable Boolean function $f(x) \in \{0, 1\}$ and one must figure out if there is an x such that $f(x) = 1$.

The class NP has a natural counting version, known as #P. Here, the task is to determine the *number* rather than the existence of satisfying inputs, i.e., to compute $|\{x: f(x) = 1\}|$. We will now analogously define #BQP, the counting version of QMA. Consider the verifying map Ω of Eq. (1) for a QMA problem, with the additional promise that Ω does not have eigenvalues between a and b , $a - b > 1/\text{poly}(n)$. Then the class #BQP consists of all problems of the form “compute the dimension of the space spanned by all eigenvectors with eigenvalues $\geq a$.”

An equivalent definition for #BQP (cf. also [6,7]) is the following: Consider a verifier Ω with the additional promise that there exist subspaces $\mathcal{A} \oplus \mathcal{R} = \mathbb{C}^{2^n}$ such that $\langle \psi | \Omega | \psi \rangle \geq a$ for all $|\psi\rangle \in \mathcal{A}$, and $\langle \psi | \Omega | \psi \rangle \leq b$ for all $|\psi\rangle \in \mathcal{R}$, where again $a - b > 1/\text{poly}(n)$ —we can think of \mathcal{A} and \mathcal{R} as containing the good and bad witnesses, respectively. (Note that there will always be “mediocre” witnesses—the question is whether there exists a decomposition into a good and a bad witness space.) Then, #BQP consists of all problems of the form “compute $\dim \mathcal{A}$.” This number is well-defined, i.e., independent of the choice of \mathcal{A} and \mathcal{R} , and, moreover, one can easily show that it is equivalent to the definition above; cf. [8].

The gap promise we impose on the spectrum of Ω is not present in the definition of QMA (though similarly




FIG. 1 (color online). A QMA verifier consists of a sequence of T local unitary gates acting on the “quantum proof” $|\psi\rangle$ and an ancillary register initialized to $|\mathbf{0}\rangle$. The final measurement on the first qubit returns $|1\rangle$ or $|0\rangle$ to accept or reject the proof, respectively. Transition probabilities can be computed by doing a “path integral” over all intermediate configurations $(i_k)_k$.

restricted versions of QMA were defined in [6,7]). Nevertheless, this promise emerges naturally when considering the counting version: QMA captures the difficulty of determining the existence of an input state with acceptance probability above a , up to a “grace interval” $[b, a]$ in which mistakes are tolerated (i.e., if the largest eigenvalue of Ω is in $[b, a]$, the oracle can return either outcome). Correspondingly, #BQP captures the difficulty of counting the number of eigenvalues above a , where eigenvalues in the grace interval $[b, a]$ can be miscounted. The reason why we choose to define #BQP with a gap promise rather than with a grace interval is the same as for QMA, namely, to have a unique outcome associated with any input.

Similarly, the idea of the Hamiltonian formulation of the problem which we will discuss below is to ask for the number of eigenstates in a certain energy interval, where states which are in some small $1/\text{poly}(n)$ neighborhood of this interval may be miscounted; again, for reasons of rigor we choose to consider only Hamiltonians with no eigenstates in that interval. It should be noted, however, that all of the equivalence proofs we give equally apply if we choose to allow for miscounting of states in those grace intervals instead of requiring them to be empty, as the proofs do not make use of the gap promise itself but rather show that all states outside those grace intervals are mapped (and thus counted) correctly. Thus, while the actual number returned by the grace interval formulation of the counting problems might change under those mappings due to different treatment of states in the grace interval, it will still be in the correct range.

The class #BQP inherits the important property from QMA of being stable under amplification; that is, the definition of #BQP is not sensitive to the choice of a and b . In particular, any $a - b > 1/\text{poly}(n)$ can be amplified (by building a new poly-size Ω' from Ω) such that $a' = 1 - \exp[-\text{poly}(n)]$, $b' = \exp[-\text{poly}(n)]$, and keeping the eigenvalue gap between a' and b' , by using a construction called *strong amplification*; cf. Ref. [9]; as shown there, strong amplification acts on all eigenvalues independently and thus also applies to #BQP. The crucial point is that strong amplification works without changing the proof itself, compared to weak amplification which takes multiple copies of the proof as an input. While this is fine for QMA, it does change the dimension of the accepting subspace in an unpredictable way and is thus not an option for the amplification of #BQP.

Complexity of computing the density of states.—Let us now show why the class #BQP is relevant for physical applications. In particular, we are going to show that computing the density of states of a local n -spin Hamiltonian $H = \sum_i H_i$ with few-body terms H_i , $\|H_i\| \leq 1$, up to accuracy $1/\text{poly}(n)$, is a problem which is complete for #BQP; i.e., it is as hard as any problem in #BQP can be. The same holds true for the (*a priori* weaker) problem of counting the ground state degeneracy of a local Hamiltonian, given a $1/\text{poly}(n)$ spectral gap above (note that Bravyi, Moore, and Russell [10] suggested this as a

definition for a quantum counting class). We can impose additional restrictions on the interaction structure of our Hamiltonian, and, as we will see, the hardness is preserved even for 2D lattices of qubits or 1D systems.

The problem DOS (density of states) is defined as follows: Given a local Hamiltonian $H = \sum_i H_i$, compute the number of orthogonal eigenstates with eigenvalues in an interval $[E_1, E_2]$ with $E_2 - E_1 > 1/\text{poly}(n)$, where we do not allow for eigenvalues within a small grace interval of width $\Delta = (E_2 - E_1)/\text{poly}(n)$ centered around E_1 and E_2 ; alternatively, we can allow for erroneous counts of eigenstates in that interval. Second, the problem #LH (sharp local Hamiltonian) corresponds to counting the number of ground states of a local Hamiltonian which has a spectral gap $\Delta = 1/\text{poly}(n)$ above the ground state subspace, given we are told the ground state energy, and where we allow for a small splitting in the ground state energies; again, we can alternatively allow ourselves to miscount states in the grace interval.

Clearly, #LH is a special instance of DOS, i.e., solving #LH can be reduced to solving DOS. In order to show that DOS is contained in #BQP, we can use a phase estimation circuit [11] to estimate the energy of any given input $|\psi\rangle$ and accept only if its energy $\langle\psi|H|\psi\rangle$ is in the interval $[E_1, E_2]$; as the desired accuracy $\Delta = 1/\text{poly}(n)$, this can be done efficiently. A detailed proof (using a more elementary circuit) is given in [8].

Let us now conversely show that #LH is a hard problem for #BQP; that is, any problem in #BQP can be reduced to counting the ground states of some gapped local Hamiltonian [12]. As in turn #LH can be reduced to DOS, which is contained in #BQP, this proves that both #LH and DOS are complete problems for #BQP; i.e., they capture the full difficulty of this class. To this end, we need to start from an arbitrary verifier circuit $U = U_T \dots U_1$ and construct a Hamiltonian which has as many ground states as the circuit has accepting inputs (corresponding to the outcome $|1\rangle_1$ on the first qubit). Let \mathcal{A} and \mathcal{R} be the eigenspaces of Ω [Eq. (1)] with eigenvalues $\geq a = 1 - 2^{-\text{poly}(n)}$ and $\leq b = 2^{-\text{poly}(n)}$, respectively, and define $U[\mathcal{R}] := \{U|\psi\rangle_I|0\rangle_A : |\psi\rangle_I \in \mathcal{R}\}$.

We will follow Kitaev's original construction for a Hamiltonian encoding a QMA verifier circuit [4,5], which for any proof $|\psi\rangle_I \in \mathcal{A}$ has the "proof history" $|\Phi\rangle = \sum_{t=0}^T U_t \dots U_1 |\psi\rangle_I |0\rangle_A |t\rangle_T$ as its ground state, where the third register is used as a clock. The Hamiltonian $H = H_{\text{init}} + \sum_{t=1}^T H_{\text{evol}}(t) + H_{\text{final}}$ has three types of terms: $H_{\text{init}} = \mathbb{1} \otimes (\mathbb{1} - |0\rangle\langle 0|_A) \otimes |0\rangle\langle 0|_T$ makes sure the ancilla is initialized, $H_{\text{evol}}(t) = -U_t \otimes |t\rangle\langle t-1|_T + \text{H.c.}$ ensures proper evolution from $t-1$ to t , and $H_{\text{final}} = \prod_{U[\mathcal{R}]} \otimes |T\rangle\langle T|_T$ gives an energy penalty to states $|\Phi\rangle$ built from proofs $|\psi\rangle_I \in \mathcal{R}$. Note that our H_{final} differs from the usual choice $|0\rangle\langle 0|_1 \otimes \mathbb{1} \otimes |T\rangle\langle T|_T$ and is in fact nonlocal; as we show in [8], this does not significantly change the relevant spectral properties [in particular, we keep the $1/\text{poly}(n)$ gap, and the ground state subspace is split at most exponentially]. With this choice of H_{final} , H

acts independently on the subspaces spanned by $\{U_t \dots U_1 |\psi\rangle_I |x\rangle_A |t\rangle_T\}_{t=0, \dots, T}$ for any $|\psi\rangle \in \mathcal{A}$ or $|\psi\rangle \in \mathcal{R}$ and $|x\rangle_A$ the computational basis, and the restriction of H to any of these subspaces describes a random walk which is characterized by the choice of $|\psi\rangle_I$ and the number of 1's in $|x\rangle_A$. These cases can be analyzed independently (see [8]), and it follows that H has a $\dim \mathcal{A}$ -fold degenerate ground state space with a $1/\text{poly}(n)$ gap above, proving #BQP-hardness of #LH.

This shows that #LH is #BQP-hard for a Hamiltonian which is a sum of $\log T$ -local terms (i.e., each term acts on $\log T$ sites), as the clock register is of size $\log T$. In order to obtain a k -body Hamiltonian, Kitaev suggested to use a unary encoding of the clock (i.e., $|t\rangle_T$ is encoded as $|1 \dots 10 \dots 0\rangle$, with t 1's), so that each Hamiltonian term acts only on three qubits of the clock. However, this makes the Hilbert space of the clock too big, and terms need to be added to the Hamiltonian to penalize illegal clock configurations. These terms divide the Hilbert space into two parts: $\mathcal{H}_{\text{legal}} \oplus \mathcal{H}_{\text{illegal}}$. Here, $\mathcal{H}_{\text{legal}}$ contains only legal clock states, whereas $\mathcal{H}_{\text{illegal}}$ contains only configurations with illegal clock states [4,5]. Since no Hamiltonian term couples these two subspaces, the Hamiltonian can be analyzed independently on the two subspaces. It turns out that its restriction to $\mathcal{H}_{\text{illegal}}$ has an at least $1/\text{poly}(n)$ higher energy, while on $\mathcal{H}_{\text{legal}}$ the Hamiltonian is exactly the same as before. Thus, one finds that the new Hamiltonian still has the right number of ground states and a $1/\text{poly}(n)$ spectral gap. The very same argument applies in the case of 1D Hamiltonians, by using the QMA construction of Ref. [2]: Again, the Hamiltonian acts independently on a "legal" and an "illegal" subspace, where the latter has a polynomially larger energy and the former reproduces the (low-energy) spectrum of the original Hamiltonian [13].

An alternative way to prove QMA-hardness on restricted interaction graphs is to use so-called *perturbation gadgets*, which yield the Hamiltonian of the Kitaev construction above from a perturbative expansion; in particular, this way one can show QMA-hardness of Pauli-type nearest-neighbor Hamiltonians on a 2D square lattice of qubits [1]. As shown in Ref. [14], such gadgets do in fact approximately preserve the whole low-energy part of the spectrum, and, thus, our #BQP-hardness proof for #LH still applies to these classes of Hamiltonians. It should be noted, however, that since the eigenvalues are preserved only up to an error $1/\text{poly}(n)$, the splitting of the ground state space will now be of the order of $1/\text{poly}(n)$; however, it can still be chosen to be polynomially smaller than the spectral gap.

Quantum vs classical counting complexity.—As we have seen, the quantum counting class #BQP exactly captures the difficulty of counting the degeneracy of ground states and computing the density of states of local quantum Hamiltonians. In the following, we will relate #BQP to classical counting classes and prove that #BQP is equal to #P, counting the number of satisfying inputs to a Boolean function [15]. In physical terms, this shows that counting

the number of ground states or determining the density of states for a quantum Hamiltonian is not harder than either problem is for a classical Hamiltonian.

Clearly, #P is contained in #BQP, as the latter includes classical verifiers. It remains to be shown that any #BQP problem can be solved by computing a #P function. We start from a verifier operator Ω , Eq. (1), and wish to show that the dimension of its accepting subspace, i.e., the subspace \mathcal{A} with eigenvalues $\geq a$, can be computed by counting satisfying inputs to some efficiently computable Boolean function. Using amplification, we can ensure that $|\dim \mathcal{A} - \text{tr} \Omega| \leq 1/4$; i.e., we need to compute $\text{tr} \Omega$ to accuracy $1/4$. This can be done by using a “path integral” method, which has been used previously to show containments of quantum classes in the classical classes PP and #P (see, e.g., [16]). We rewrite $\text{tr} \Omega = \sum_{\zeta} f(\zeta)$ as a sum over products of transition probabilities along a path $\zeta \equiv (i_0, \dots, i_N, j_1, \dots, j_N)$, where

$$f(\zeta) = \langle i_0 |_I \langle \mathbf{0} |_A U_1^\dagger | j_1 \rangle \langle j_1 | U_1^\dagger \cdots U_T^\dagger | j_T \rangle \\ \times \langle i_T | [|0\rangle\langle 0| \otimes \mathbb{1}] | i_T \rangle \langle i_T | U_T \cdots U_1 | i_0 \rangle | \mathbf{0} \rangle_A \quad (2)$$

(cf. Figure 1). Since any such sum over an efficiently computable $f(\zeta)$ can be determined by counting the satisfying inputs to some Boolean formula (see [8] for details), it follows that Ω can be computed by using a single query to a black box solving #P problems.

Summary and discussion.—In this work, we considered two problems: computing the density of states and computing the ground state degeneracy of a local Hamiltonian of a spin system. In order to capture the computational difficulty of these problems, we introduced the quantum complexity class #BQP, the counting version of the class QMA. We proved that this complexity class exactly captures the difficulty of our two problems, even when restricting to local Hamiltonians on 2D lattices of qubits or to 1D chains, since all these problems are #BQP-complete [17].

We have further shown that #BQP is no harder than its classical counterpart #P. In particular, this implies that computing the density of states is no harder for quantum Hamiltonians than it is for classical ones. While this quantum-classical equivalence might seem surprising at the Hamiltonian level, it should be noted that the classes #P and PP quite often form natural “upper bounds” for many quantum *and* classical problems.

What about the problem of computing the density of states for fermionic systems, such as many-electron systems? On the one hand, this problem will be still in #BQP and thus #P, since any local fermionic Hamiltonian can be mapped via the Jordan-Wigner transform to a (nonlocal) Hamiltonian on a spin system, whose energy can still be estimated efficiently by a quantum circuit [20]. On the other hand, hardness of the problem for #BQP can be shown, e.g., by using the #BQP-hardness of #LH and encoding each spin by using one fermion in two modes, similar to Ref. [20]. Thus, computing the density of states for fermionic systems is #BQP-complete as well.

We thank S. Aaronson, D. Gottesman, D. Gross, Y. Shi, M. van den Nest, J. Watrous, and S. Zhang for helpful discussions and comments. Research at PI is supported by Industry Canada and Ontario MRI. N. S. is supported by the Gordon and Betty Moore Foundation through Caltech’s Center for the Physics of Information and NSF Grant No. PHY-0803371.

Note added.—After completion of this work, we learned that Shi and Zhang [21] have independently defined #BQP and shown its relation to #P by using the same technique.

-
- [1] R. Oliveira and B.M. Terhal, *Quantum Inf. Comput.* **8**, 900 (2009).
 - [2] D. Aharonov, D. Gottesman, S. Irani, and J. Kempe, *Commun. Math. Phys.* **287**, 41 (2009).
 - [3] N. Schuch and F. Verstraete, *Nature Phys.* **5**, 732 (2009).
 - [4] A. Y. Kitaev, A. H. Shen, and M. N. Vyalyi, *Classical and Quantum Computation* (American Mathematical Society, Providence, RI, 2002).
 - [5] D. Aharonov and T. Naveh, [arXiv:quant-ph/0210077](http://arxiv.org/abs/quant-ph/0210077).
 - [6] D. Aharonov, M. Ben-Or, F.G. Brandao, and O. Sattath, [arXiv:0810.4840](http://arxiv.org/abs/0810.4840).
 - [7] R. Jain, I. Kerenidis, M. Santha, and S. Zhang, [arXiv:0906.4425](http://arxiv.org/abs/0906.4425).
 - [8] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.107.040501> for technical details.
 - [9] C. Marriott and J. Watrous, *Comput. Complex.* **14**, 122 (2005).
 - [10] S. Bravyi, C. Moore, and A. Russell, [arXiv:0907.1297](http://arxiv.org/abs/0907.1297).
 - [11] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca, *Proc. R. Soc. A* **454**, 339 (1998).
 - [12] Note that the connection between QMA with a unique “good witness,” such as in our #BQP definition, and local Hamiltonians with a unique ground state and a $1/\text{poly}(n)$ gap has been shown in Ref. [6].
 - [13] N. Schuch, I. Cirac, and F. Verstraete, *Phys. Rev. Lett.* **100**, 250501 (2008).
 - [14] J. Kempe, A. Kitaev, and O. Regev, *SIAM J. Comput.* **35**, 1070 (2006).
 - [15] Formally speaking, the reduction from #BQP to #P is *weakly parsimonious*; i.e., for any function $f \in \#BQP$, there exist polynomial-time computable functions α and β and a function $g \in \#P$, such that $f = \alpha \circ g \circ \beta$. This differs from Karp reductions $\alpha = \text{Id}$, but it still requires only a single call to a #P oracle.
 - [16] L. Adleman, J. DeMarrais, and M.-D. Huang, *SIAM J. Comput.* **26**, 1524 (1997).
 - [17] Note that computing the density of states to multiplicative accuracy is less difficult; see Refs. [18,19].
 - [18] F.G.S.L. Brandao, [arXiv:0810.0026](http://arxiv.org/abs/0810.0026).
 - [19] T.J. Osborne, [arXiv:cond-mat/0605194](http://arxiv.org/abs/cond-mat/0605194).
 - [20] Y.-K. Liu, M. Christandl, and F. Verstraete, *Phys. Rev. Lett.* **98**, 110503 (2007).
 - [21] Y. Shi and S. Zhang (private communication), notes available at <http://www.cse.cuhk.edu.hk/~syzhang/papers/SharpBQP.pdf>.