

The Complexity of Computing the Density of States

Brielin Brown
Distinguished Major Thesis

Thesis Reader: Professor Gabriel Robins, Department of Computer Science
Second Reader: Professor Olivier Pfister, Department of Physics

This work is an extension of Arxiv:1010.3060 [quant-ph], forthcoming in Physical Review Letters, completed under the supervision of Steve Flammia with assistance from Norbert Schuch at the Perimeter Institute for Theoretical Physics during the Undergraduate Research Program summer 2010. This thesis completed in fulfillment of the University of Virginia Department of Computer Science Bachelor of Arts Distinguished Major.

Abstract

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 QMA . We show that $\#BQP$ is not harder than its classical counting counterpart $\#P$, which in turn implies that computing the ground state degeneracy or density of states for classical Hamiltonians is just as hard as it is for quantum Hamiltonians.

1 Introduction

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 work, 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 lattices [12] and even one-dimensional chains [3], as well as fermionic systems [13].

We will characterize the computational complexity of two problems: the density of states for a local Hamiltonian, and counting the ground state degeneracy of a local gapped Hamiltonian. To do so, we will introduce the quantum counting class $\#BQP$. This class constitutes the natural counting extension to the class QMA , and the natural “quantum version” of the counting function class $\#P$. Vaguely speaking, $\#BQP$ counts the number of possible “quantum solutions” to a problem that can be verified, but not necessarily solved, efficiently using a quantum computer.

Our first major result is 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 equivalent to the hardest problems in this class. After characterizing the complexity of the density of states and ground state solutions problems, we move to relating $\#BQP$ to well-known classical counting classes. Using a weakly parsimonious reduction, we show that $\#BQP$ is equal to its classical counterpart $\#P$. As functions in $\#P$ are thought to be difficult to compute classically, this gives an exact characterization of why computing the density of states and finding the ground state degeneracy of a local Hamiltonian are so difficult.

2 Quantum Complexity Classes

To fix notation, we first review the relevant complexity classes. We begin with QMA, the natural quantum version of the class NP. The *verifier* V verifies an n -qubit quantum state serving as a “quantum solution,” or *proof*, analogous to the n -bit classical proof verified in the class NP. The verifier consists of a length- T ($T \in \text{poly}(n)$) quantum circuit $U = U_T \cdots U_1$ (with local gates U_t) acting on $m \in \text{poly}(n)$ qubits, which takes the n -qubit quantum state $|\psi\rangle_I$ as an input, together with $m - n$ initialized ancillas, $|\mathbf{0}\rangle_A \equiv |0 \cdots 0\rangle_A$, 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”). The acceptance probability for a proof $|\psi\rangle$ is then given by $p_{\text{acc}}(V(\psi)) := \langle \psi | \Omega | \psi \rangle$, with

$$\Omega = (\mathbb{1}_I \otimes \langle \mathbf{0} |_A) U^\dagger (|1\rangle_1 \langle 1|_1 \otimes \mathbb{1}) U (\mathbb{1}_I \otimes |\mathbf{0}\rangle_A). \quad (1)$$

Definition 1 (QMA). *Let a, b satisfy $a - b > \frac{1}{p(n)}$ for some polynomial $p(n)$. A language L is in $\text{QMA}(a, b)$ if there exists a quantum verifier circuit U such that*

$$\begin{aligned} x \in L &\Rightarrow \exists |\psi\rangle \text{ s.t. } \langle \psi | \Omega | \psi \rangle > a \\ x \notin L &\Rightarrow \forall |\psi\rangle \langle \psi | \Omega | \psi \rangle < b. \end{aligned}$$

The decision problem “determine whether a local Hamiltonian has a ground state with energy less than a ” is was the first known complete problem for this class. Thus, this definition of QMA allows us to quantify the difficulty of computing the ground state energy, $E_0(H)$, of a local hamiltonian H up to $1/\text{poly}(n)$ accuracy via binary search and a single query to a black box solving QMA problems [4].

The class NP has a natural counting version, known as #P [14],

Definition 2. *A function problem $f \in \#P$ if and only if it is of the form “given a verifier machine M for a problem in NP, compute the number of strings x of length q such that $M(x) = 1$ ”*

Here, the task is to determine the *number* rather than the existence of satisfying inputs to the verifier (solutions to the problem). In order to define #BQP, we now consider a modified version of QMA. This class by itself holds no intrinsic interest for us and is merely a stepping stone to defining #BQP.

Definition 3. *Let a, b, c satisfy $a - b > \frac{1}{p(n)}$ and $a - c > \frac{1}{p(n)}$ for some polynomial $p(n)$ and let $|\phi\rangle$ be an eigenvector of Ω . A language L is in *modified-QMA*(a, b, c) if there exists a quantum verifier circuit U such that*

$$\begin{aligned} x \in L &\Rightarrow \exists |\psi\rangle \text{ s.t. } \langle \psi | \Omega | \psi \rangle > a \\ &\quad \nexists |\phi\rangle \text{ s.t. } c < \langle \phi | \Omega | \phi \rangle < a \\ x \notin L &\Rightarrow \forall |\psi\rangle \langle \psi | \Omega | \psi \rangle < b \end{aligned}$$

Of course, this additional promise that we impose on the spectrum of Ω is not present in the original definition of QMA, however similarly restricted definitions of QMA are defined in [2, 8]. This promise gap arises naturally when considering the counting version of QMA: QMA captures the difficulty of determining the existence of an input state with acceptance probability above a , up to a grace interval of $a - b$ in which mistakes are tolerated. Correspondingly, #BQP captures

the difficulty of counting the number of (orthogonal) inputs with an acceptance probability above a , again with a grace interval of $a - b$ in which states are allowed to be miscounted. Similarly, the Hamiltonian formulation of the problem which we will discuss below asks for the existence (QMA) or number ($\#BQP$) of states in a certain energy interval, where states which are in some small $1/\text{poly}(n)$ neighborhood of this interval may be miscounted.

The problem of distinguishing between states exponentially close to the threshold of acceptance and those above the threshold of acceptance is a difficult and open problem that originally motivated this modified definition in [2,8]. In our case, using the original unrestricted definition of QMA would allow us to characterize the number of orthogonal inputs with acceptance probability above the threshold to exponential accuracy.

Using modified-QMA, we define $\#BQP$ as its counting extension,

Definition 4. *Given a verifier map, Ω , for a modified-QMA problem, the class $\#BQP$ consists of all problems of the form “compute the dimension of the space spanned by all eigenvectors of Ω with eigenvalues $\geq a$ ”*

Another definition for $\#BQP$, cf. [2,8] which will become useful is the following:

Definition 5. *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 $a - b > 1/\text{poly}(n)$. Then $\#BQP$ consists of all problems of the form “compute $\dim \mathcal{A}$ ”.*

In fact, these two definitions are equivalent.

Theorem 6. *Equivalence of Definition 4 and Definition 5.*

Proof. First we show Definition 5 is well defined up to the dimension of the subspaces. Consider two decompositions $\mathbb{C}^{2^n} = \mathcal{A} \oplus \mathcal{R}$ and $\mathbb{C}^{2^n} = \mathcal{A}' \oplus \mathcal{R}'$. Without loss of generality, if we assume $\dim \mathcal{A} > \dim \mathcal{A}'$, it follows $\dim \mathcal{A} + \dim \mathcal{R}' > 2^n$, and thus there exists a non-trivial $|\mu\rangle \in \mathcal{A} \cap \mathcal{R}'$, which contradicts the definition.

Now we show this is equivalent to our original definition in terms of the spectrum of Ω : To show Def 4 \Rightarrow Def 5, let \mathcal{A} be spanned by the eigenvectors with eigenvalues $\geq a$. To show Def 5 \Rightarrow Def 4, one can use the minimax-principle for eigenvalues [5] to infer that Ω has a spectral gap between a and b , where the dimension of the eigenspace with eigenvalues $\geq a$ equals $\dim \mathcal{A}$. \square

The class $\#BQP$ inherits the useful property of strong error reduction from QMA: the acceptance probabilities can be amplified from (a, b, c) to $(1 - 2^{-r}, 2^{-r}, 2^{-r})$ by modifying only the verification procedure. That is, without needing additional copies of the proof $|\psi\rangle$. This is established by the following theorem.

Theorem 7. *Let $\#BQP_m(a, b, c)$ denote $\#BQP$ with an m qubit witness. Then $\#BQP_m(a, b, c) \subset \#BQP_m(1 - 2^{-r}, 2^{-r}, 2^{-r})$ for every $r \in \text{poly}$.*

Here we describe the procedure, but note that the proof is identical to the one presented in [11]: the promise that Ω has a spectral gap implies that the procedure behaves the same way with a bad witness as it does when the considered problem is not in QMA. Assume the witness and a register A of t ancilla qubits are together in a total register R . We have a verification procedure V operating on R with completeness and soundness probabilities given by a and b . Our new procedure V' operates as follows:

1. Set $y_0 = 1$ and $i = 1$
2. Repeat:
 - Apply V to R and measure R by applying $|0\rangle\langle 0| \otimes \mathbb{1}_{m-1} \otimes \mathbb{1}_A$. Let y_i denote the outcome and set $i = i + 1$.
 - Apply V^\dagger to R and measure R by applying $\mathbb{1}_m \otimes |\mathbf{0}\rangle\langle \mathbf{0}|_A$. Let y_i denote the outcome and set $i = i + 1$.

Until $i \geq N$, where $N = 8q^2r$.

3. For each i , set

$$z_i = \begin{cases} 1 & \text{if } y_i = y_{i-1} \\ 0 & \text{if } y_i \neq y_{i-1} \end{cases}$$

Accept if and only if $\sum_{i=1}^N z_i \geq N \frac{a+b}{2}$.

3 The Complexity of the Density of States

We now use the class #BQP to characterize the complexity of the density of states problem and the problem of counting the number of ground states of a local Hamiltonian. First, we define these problems more specifically, and then show that they are both #BQP-complete.

Definition 8. *Given a local Hamiltonian $H = \sum_i H_i$, the problem DOS (density of states) is to 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 states within a small region $\Delta = 1/\text{poly}(n)$ (with $\Delta \ll E_2 - E_1$) around E_1 and E_2 may also be counted.*

Definition 9. *Given a local Hamiltonian $H = \sum_i H_i$ with a spectral gap $\Delta = 1/\text{poly}(n)$ above the ground state subspace, the problem #LH (sharp local Hamiltonian) is to compute the dimension of said ground state subspace.*

Next, we show that #LH and DOS are #BQP-complete. We prove this in two stages, first showing that the problems are contained in #BQP and then showing they are hard for the class.

Theorem 10. *#LH and DOS are contained in #BQP.*

Proof. Clearly #LH can be reduced to solving DOS, so we only need to show that DOS is contained in #BQP. To do this, we must design a quantum circuit which accepts any input state $|\psi\rangle$ with average energy $\langle \psi | H | \psi \rangle \in [E_1, E_2]$ with high probability, while rejecting any state with energy outside $\langle \psi | H | \psi \rangle \in [E_1 - \Delta, E_2 + \Delta]$ with high probability.

We apply the phase estimation algorithm of [6]. Let $\{u_i\}$ be the eigenvectors of H with eigenvalues $\{e_i\}$, U be the unitary associated with H and $\{u_i\}$ be the phases associated with the eigenvalues. We begin with an ancilla register of t qubits in the state $|\mathbf{0}\rangle_A$ and n qubit input $|\psi\rangle = \sum_{i=0}^m \alpha_i |\phi_i\rangle$. We apply a hadamard gate to each qubit in the ancilla register, followed by a series of controlled- U^{2^j} operations controlled by qubit $j = 0 \dots t$. This gives the state,

$$\frac{1}{2^{t/2}} \sum_{k=0}^{2^t-1} \sum_{i=1}^m \alpha_i e^{2\pi i k \phi_i} |k\rangle |u_i\rangle \quad (2)$$

Then by applying the inverse quantum fourier transform, we have the state,

$$\frac{1}{2^t} \sum_{i=1}^m \alpha_i |k_i\rangle |u_i\rangle \quad (3)$$

where k_i can be written $\frac{k}{2^t} = 0.k_i^{(1)} \dots k_i^{(t)}$ as the best t bit approximation to ϕ_i . As the best approximation, $\phi_i = \frac{k_i}{2^t} + \delta$ with $0 < \delta \leq \frac{1}{2^{t+1}}$. This exponentially small error in the phase corresponds to an $e^{1/2^{t+1}} \in 1/\text{poly}(t)$ error in the phase. By using $t = n + \log(2 + \frac{1}{2\epsilon})$ ancilla qubits we achieve $1/\text{poly}(n) = \Delta$ error with probability of success $1 - \epsilon$. Thus by coherently setting the output qubit to $|1\rangle$ if and only if

$$E_1 - \Delta \leq \sum_{i=1}^m |\alpha_i|^2 e^{2\pi i k_i} \leq E_2 + \Delta \quad (4)$$

we determine the energy of the state $|\psi\rangle$ with accuracy Δ . The counting extension of this problem, determining the number of eigenstates in this energy range (DOS), must then lie in $\#\text{BQP}$. That is, $\text{DOS} \in \#\text{BQP}$. \square

Theorem 11. *$\#\text{LH}$ and DOS are $\#\text{BQP}$ -hard.*

Proof. Since $\#\text{LH}$ reduces to DOS , it is sufficient to show the $\#\text{BQP}$ -hardness of $\#\text{LH}$. We show that $\#\text{LH}$ is $\#\text{BQP}$ -hard by starting with an arbitrary QMA verifier circuit $U = U_T \dots U_1$ and construct a Hamiltonian with as many ground states as the circuit has accepting inputs. By amplification, we can assume that the acceptance and rejection thresholds for the verifier are $a = 1 - \epsilon$ and $b = \epsilon$, where we can choose ϵ to be a function of n which decreases like $\exp(-p(n))$ for our choice of polynomial $p(n)$. As before, let \mathcal{A} and \mathcal{R} be the eigenspaces of Ω with eigenvalues $\geq a$ and $\leq b$, respectively. Define

$$U[\mathcal{R}] := \{U|\psi\rangle_I |\mathbf{0}\rangle_A : |\psi\rangle_I \in \mathcal{R}\} \quad (5)$$

and denote the projector onto this space by $\Pi_{U[\mathcal{R}]}$. Notice that for any state $|\chi\rangle \in U[\mathcal{R}]$, due to our rejection threshold $b = \epsilon$, we have

$$\langle \chi | (|1\rangle\langle 1|_1 \otimes \mathbb{1}) | \chi \rangle \leq \epsilon. \quad (6)$$

From here, we mimic Kitaev's original construction to encode a QMA verifier circuit into a Hamiltonian which has a "proof history" as its ground state for any proof $|\phi\rangle_I \in \mathcal{A}$. That is, the ground states of the Hamiltonian are given by

$$|\Phi\rangle = \sum_{t=0}^T U_t \dots U_1 |\phi\rangle_I |\mathbf{0}\rangle_A |t\rangle_T \quad (7)$$

where the third register is used as a "clock". The Hamiltonian has the form

$$H = H_{\text{init}} + \sum_{t=1}^T H_{\text{evol}}(t) + H_{\text{final}} \quad (8)$$

where,

- $H_{\text{init}} = \mathbb{1}_I \otimes (\mathbb{1} - |\mathbf{0}\rangle\langle\mathbf{0}|_A) \otimes |0\rangle\langle 0|_T$ checks that the ancilla is property initialized, penalizing states without properly initialized ancillas;
- $H_{\text{evol}}(t)$ is defined as

$$\begin{aligned} H_{\text{evol}}(t) = & -\frac{1}{2}U_t \otimes |t\rangle\langle t-1|_T - \frac{1}{2}U_t^\dagger \otimes |t-1\rangle\langle t| \\ & + \frac{1}{2}I \otimes |t\rangle\langle t| + \frac{1}{2}I \otimes |t-1\rangle\langle t-1| \end{aligned} \quad (9)$$

and checks that the propagation from time $t-1$ to t is correct, penalizing states with erroneous propagation;

- $H_{\text{final}} = \Pi_{U[\mathcal{R}]} \otimes |T\rangle\langle T|_T$ causes each state $|\phi\rangle$ built from an input $|\psi\rangle_I \in \mathcal{R}$ (but which nonetheless has a correctly initialized ancilla) to receive an energy penalty.

Note we have chosen H_{final} different from Kitaev's in order to simplify the analysis. With this, H acts independently on the $(T+1)$ -dimensional subspaces spanned by $\{U_t \dots U_1 |\phi\rangle_I |\mathbf{x}\rangle_A |t\rangle_T\}_{t=0, \dots, T}$ for any choice of $|\phi\rangle \in \mathcal{A}, \mathcal{R}$ and basis of ancillas $|\mathbf{x}\rangle_A$ containing $|\mathbf{0}\rangle_A$. By restricting H to any of these subspaces we have a random walk with only two distinct cases: either $|\phi\rangle \in \mathcal{A}$ and $|\mathbf{x}\rangle_A = |\mathbf{0}\rangle_A$ or otherwise. In the second case the ground state energy is $1/\text{poly}(n)$ above the first case, and the ground state degeneracy is $\dim \mathcal{A}$. It follows that the total Hamiltonian H has a $1/\text{poly}(n)$ spectral gap between the ground state space and the first excited state. These claims are made more rigorous in Lemma 12 below.

As we have presented it, H is not a local Hamiltonian, so we must relate it to a new, local one while retaining the relevant properties, namely the ground space dimension (up to small splitting in energies) and the $1/\text{poly}(n)$ spectral gap. The first step is to replace our nonlocal version of H_{final} with a local one and show that this change doesn't significantly affect the spectrum. As we will prove subsequently in Lemma 14, the standard final Hamiltonian term $H_{\text{final}}^{\text{std}} = |0\rangle\langle 0|_1 \otimes \mathbb{1} \otimes |T\rangle\langle T|_T$ obeys the inequality

$$H_{\text{final}}^{\text{std}} \geq H_{\text{final}} - \sqrt{\epsilon} \mathbb{1}. \quad (10)$$

Thus, if instead of projecting onto the space $\mathbb{1} - \Pi_{U[\mathcal{R}]}$ to check proper evolution, we simply check that the application of the operators $U_T \dots U_1$ resulted in a $|1\rangle$ on the first qubit, then the additional error we can incur is at most $\sqrt{\epsilon}$. Put differently, the first excited state energy decreases by no more than $\sqrt{\epsilon}$. Hence, replacing our nonlocal H_{final} with the local $H_{\text{final}}^{\text{std}}$ does not affect the excited states significantly.

Now we consider the ground state space. The error incurred when the input state is composed of a proper proof causes the energy of the ground states to increase by no more than ϵ , for the following reason. Similar to Eq. 6, if we have a state $|\chi\rangle \in U[\mathcal{A}]$, then from strong amplification we have the bound

$$\langle \chi | (|0\rangle\langle 0|_1 \otimes \mathbb{1}) | \chi \rangle = \langle \chi | H_{\text{final}}^{\text{std}} | \chi \rangle \leq \epsilon. \quad (11)$$

Now we show that H' (and thus H) has $\dim(\mathcal{A})$ small eigenvalues. Consider $H'|_{\mathcal{S}_1}$, then create $|\Phi'\rangle = |\phi\rangle_I |0\rangle_A |\psi\rangle$ with $\psi = \frac{1}{\sqrt{T+1}} \sum_{t=0}^T |t\rangle$. Now compute,

$$\begin{aligned} \langle \Phi' | H' |_{\mathcal{S}_1} | \Phi' \rangle &= \langle \Phi' | H'_{\text{init}} |_{\mathcal{S}_1} | \Phi' \rangle + \langle \Phi' | \sum H'_{\text{evol}} |_{\mathcal{S}_1} | \Phi' \rangle + \langle \Phi' | H'_{\text{out}} |_{\mathcal{S}_1} | \Phi' \rangle \\ &= \langle \Phi' | H'_{\text{init}} |_{\mathcal{S}_1} | \Phi' \rangle + \langle \Phi' | \mathbb{1} \otimes E | \Phi' \rangle + \langle \Phi' | (\Pi_{\mathcal{R}} \otimes |\mathbf{0}\rangle\langle \mathbf{0}|_A \otimes |T\rangle\langle T|) | \Phi' \rangle \\ &= 0 + 0 + \langle \phi | \Pi_{\mathcal{R}} | \phi \rangle \frac{1}{T+1} \\ &\leq \frac{\epsilon}{T+1} \end{aligned}$$

where, since we have restricted ourselves to $|\phi\rangle \in \mathcal{A}$, $\langle \phi | \Pi_{\mathcal{R}} | \phi \rangle \leq \epsilon$.

Now that we have an upper bound on the eigenvalues of $H|_{\mathcal{S}_1}$, we proceed to determine a lower bound for the eigenvalues of $H|_{\mathcal{S}_2}$. Let $H' = A_1 + A_2$ with $A_1 = H'_{\text{init}} + H'_{\text{out}}$ and $A_2 = H'_{\text{evol}}$. Since A_1 is the sum of commuting projectors it follows that its eigenvalues are all non-negative integers, and thus its minimal non-zero eigenvalue is no less than 1. The smallest positive eigenvalue of A_2 is the smallest positive eigenvalue of E , which is bounded in Ref. [10] to be greater than $1 - \cos(\pi/(T+1)) \geq c/T^2$ with c a positive constant. Now let \mathcal{L}_1 and \mathcal{L}_2 be the null spaces of A_1 and A_2 respectively, and note that $\mathcal{L}_1 \cap \mathcal{L}_2 = \{0\}$. By bounding the angle between these subspaces, $\sin^2(\theta) \geq \frac{1-\sqrt{\epsilon}}{L+1}$, and using the following lemma (both given in Ref. [10])

Lemma 13. *Let A_1 and A_2 be nonnegative operators, and \mathcal{L}_1 and \mathcal{L}_2 their null subspaces, where $\mathcal{L}_1 \cap \mathcal{L}_2 = \{0\}$. Suppose no eigenvalue of A_1 or A_2 is smaller than v . Then,*

$$A_1 + A_2 \geq 2v \sin^2 \frac{\theta}{2} \mathbb{1} \quad (17)$$

we can lower bound the smallest positive eigenvalue of $H'|_{\mathcal{S}_2}$,

$$\langle \Phi' | H' |_{\mathcal{S}_2} | \Phi' \rangle \leq \frac{c(1-\sqrt{\epsilon})}{L^2(L+1)} \quad (18)$$

As these cases give an upper bound to the largest eigenvalue of the ground state and a lower bound to the smallest eigenvalue of the first excited state, it follows that there is a $1/\text{poly}(n)$ spectral gap between the ground state and the first excited state of our Hamiltonian H .

Since we can create exactly one good witness for each $\phi \in \mathcal{A}$, we know that the ground state degeneracy of H is at least $\dim \mathcal{A}$. However, by the definition of $\#\text{BQP}$ there are no more than $\dim \mathcal{A}$ witnesses accepted with high probability, thus the ground state degeneracy of H is equal to $\dim \mathcal{A}$. \square

It remains to prove Eq. 10, which is a corollary to the following lemma. (It follows by choosing $P = |0\rangle\langle 0|_1 \otimes \mathbb{1}$, $Q = \Pi_{U[\mathcal{R}]}$, and noting that ϵ is the bound on the rejection probability for a state in $U[\mathcal{R}]$.)

Lemma 14. *Let P and Q be projectors such that $\|Q(\mathbb{1} - P)Q\| \leq \epsilon$. Then*

$$P - Q \geq -\sqrt{\epsilon} \mathbb{1}. \quad (19)$$

Proof. We begin by recalling the result due to Jordan [9] (see Ref. [7] for a more modern treatment) for the simultaneous canonical form of two projectors. In the subspace where P and Q commute, both operators are diagonal in a common basis and the spectrum is either $(0, 0)$, $(0, 1)$, $(1, 0)$, or $(1, 1)$, and direct sums of those terms. In the subspace where they don't commute, the problem decomposes into a direct sum of two-by-two blocks given by

$$P_j = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad Q_j = \begin{pmatrix} c^2 & cs \\ cs & s^2 \end{pmatrix}, \quad (20)$$

where $s = \sin(\theta_j)$ for some angle θ_j , $c^2 + s^2 = 1$, and the subscript j just labels a generic block.

In fact, this two-by-two block form is completely general if we allow embedding our projectors into a larger space while preserving their rank. The rank-preserving condition guarantees that our bound is unchanged, since we are only appending blocks of zeros, and so we will consider this two-by-two form without loss of generality.

The constraint that $\|Q(\mathbb{1} - P)Q\| \leq \epsilon$ implies constraints on the values that the $\sin(\theta_j)$ can take. In particular, we can directly compute this operator norm in each block separately, and we find that for all j

$$\|Q_j(\mathbb{1} - P_j)Q_j\| = \sin^2(\theta_j) \leq \epsilon. \quad (21)$$

We can also directly compute in each block that

$$P_j - Q_j = \begin{pmatrix} 1 - c^2 & 1 - cs \\ -cs & -s^2 \end{pmatrix}. \quad (22)$$

The spectrum of this operator is easily computed to be $\pm|\sin(\theta_j)|$. Thus, the least eigenvalue of $P - Q$ is bounded from below by $-\sqrt{\epsilon}$. \square

4 Quantum vs. Classical Counting Complexity

We have shown that the quantum counting class $\#\text{BQP}$ exactly captures the complexity of counting the degeneracy of ground states and computing the density of states for local hamiltonians. Now we relate this quantum class to the well known classical class $\#\text{P}$. In particular, we use a *weakly parsimonious* reduction to show that $\#\text{BQP}$ and $\#\text{P}$ are equal. That is, for any function $f \in \#\text{BQP}$ there exist polynomial-time computable functions α and β , and a function $g \in \#\text{P}$, such that $f = \alpha \circ g \circ \beta$. This differs from Karp reductions where no postprocessing is allowed, $\alpha = \text{Id}$, but it still only requires a single call to a $\#\text{P}$ oracle. This implies that counting the number of ground states of a local quantum hamiltonian is no harder than it is for a classical one.

Theorem 15. $\#\text{BQP} =_{w.p.r.} \#\text{P}$

Proof. Clearly, $\#\text{P} \subset \#\text{BQP}$: we simply choose our quantum verifier to compute the value of the classical function encoding the $\#\text{P}$ problem. We now show that any $\#\text{BQP}$ problem can be solved by computing a $\#\text{P}$ function. That is, given the verifier operator Ω we determine the dimension of the satisfying subspace by counting satisfying inputs to an efficiently computable boolean function.

First, we use strong error reduction to let $a = 1 - 2^{-(n-2)}$ and $b = 2^{-(n+2)}$. It follows that

$$|\dim \mathcal{A} - \text{tr } \Omega| \leq 2^n 2^{-(n+2)} = \frac{1}{4} \quad (23)$$

and thus we need to compute $\text{tr } \Omega$ to accuracy $1/4$. This can be done using the “path integral” method previously used to show containments of quantum classes inside PP and $\#\text{P}$. We rewrite $\text{tr } \Omega = \sum_{\zeta} f(\zeta)$, where the sum is over products of transition probabilities along a path, which we label

$$\zeta \equiv (i_0, \dots, i_N, j_1, \dots, j_N), \quad (24)$$

so that

$$f(\zeta) = \langle i_0 |_I \langle \mathbf{0} |_A U_1^\dagger |j_1\rangle \langle j_1| U_1^\dagger \cdots U_T^\dagger |j_T\rangle \times \langle i_T | [|0\rangle \langle 0|_1 \otimes \mathbb{1}] |i_T\rangle \langle i_T| U_T \cdots U_1 |i_0\rangle_I | \mathbf{0} \rangle_A. \quad (25)$$

Since any quantum circuit can be recast in terms of real gates at the cost of doubling the number of qubits, we can simplify the proof by assuming $f(\zeta)$ to be real. Computing $\sum_{\zeta} f(\zeta)$ to the required accuracy can now be mapped to summing over a boolean function

To achieve the desired accuracy it is sufficient to approximate f up to $|\zeta| + 2$ digits, where $|\zeta| = \text{poly}(n)$ is the number of bits in ζ . Now define,

$$g(\zeta) := \text{round}[2^{|\zeta|+2}(f(\zeta) + 1)] \quad (26)$$

and note that $g(\zeta)$ is a positive and integer-valued function satisfying

$$\left| [2^{-|\zeta|-2} \sum_{\zeta} g(\zeta) - 1] - \sum_{\zeta} f(\zeta) \right| \leq \frac{1}{4}. \quad (27)$$

Finally, by defining a boolean indicator function,

$$h(\zeta, \xi) = \begin{cases} 1 & \text{if } 0 \leq \xi < g(\zeta) \\ 0 & \text{otherwise.} \end{cases}$$

we can write $g(\zeta) = \sum_{\xi \geq 0} h(\zeta, \xi)$ which can be solved by a single query to a black box solving $\#\text{P}$ problems.

This shows that $\text{tr } \Omega$ can be approximated to accuracy $\frac{1}{4}$, and thus $\dim \mathcal{A}$ can be determined by counting the number of satisfying assignments of a single boolean function $h(\zeta, \xi)$ that can be efficiently constructed from Ω , i.e., by a single query to a black box solving $\#\text{P}$ problems. \square

5 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 $\#\text{BQP}$, the natural 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 two-dimensional

lattices of qubits or to one-dimensional chains, since all these problems are complete problems for the class $\#\text{BQP}$.

We have further shown that $\#\text{BQP}$ is no harder than its classical counterpart $\#\text{P}$. 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 $\#\text{P}$ and PP quite often form natural “upper bounds” for many quantum *and* classical problems. As an example, consider the class PQP , the unbounded version of the class BQP . It follows from arguments in [1] used to show $\text{BQP} \subset \text{PP}$ that $\text{PQP} = \text{PP}$. That is, PP forms a natural upper bound for what is computable using both quantum circuits and classical probabilistic ones. The oracle equivalence of PP and $\#\text{P}$ (that is, $\text{P}^{\text{PP}} = \text{P}^{\#\text{P}}$) then motivates the notion that $\#\text{P}$ forms an upper bound for quantum counting.

This work leaves open a number of important questions. In the future, we wish to investigate the feasibility of amplifying the “intermediate” witnesses of the traditional QMA to a low probability of acceptance as explored previously in Ref. [2,8]. We are also interested in the question of whether $\#\text{P}$ and $\#\text{BQP}$ are *set* equivalent. That is, whether they contain exactly the same constituent problems.

6 Acknowledgements

We gladly 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 the Perimeter Institute for Theoretical Physics is supported by the Government of Canada through Industry Canada and by the Province of Ontario through the Ministry of Research and Innovation. Norbert Schuch is supported by the Gordon and Betty Moore Foundation through Caltech’s Center for the Physics of Information and NSF Grant No. PHY-0803371. Brielin Brown would like to thank Professor’s Gabriel Robins and Olivier Pfister for evaluating this document in completion of the Computer Science distinguished majors program.

References

- [1] L. Adleman, J. DeMarras, , and M.-D. Huang. Quantum computability. *SIAM J. Comput.*, 26:1524, 1997.
- [2] Dorit Aharonov, Michael Ben-Or, Fernando G.S.L. Brandao, and Or Sattath. The pursuit for uniqueness: Extending valiant-vazirani theorem to the probabilistic and quantum settings. 2008.
- [3] Dorit Aharonov, Daniel Gottesman, Sandy Irani, and Julia Kempe. The power of quantum systems on a line. *Commun. Math. Phys.*, 287:41, 2009.
- [4] Dorit Aharonov and Tomer Naveh. Quantum np - a survey. 2002.
- [5] R. Bhatia. *Matrix Analysis*. Springer, 1997.
- [6] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca. Quantum algorithms revisited. *Proc. R. Soc. Lond.*, 454:339–354, 1998.

- [7] P.R. Halmos. Two subspaces. *Trans. Amer. Math. Soc.*, 144:381, 1969.
- [8] Rahul Jain, Iordanis Kerenidis, Miklos Santha, and Shengyu Zhang. On the power of a unique quantum witness. 2009.
- [9] C. Jordan. Essai sur la géométrie à n dimensions. *Bulletin de la Société Mathématique de France*, 3:103–174, 1875.
- [10] A. Yu. Kitaev, A. H. Shen, and M. N. Vyalyi. *Classical and quantum computation*. American Mathematical Society, Providence, Rhode Island, 2002.
- [11] C. Marriott and J. Watrous. Quantum arthur-merlin games. *Comput. Complex.*, 14:122, 2005.
- [12] Roberto Oliveira and Barbara M. Terhal. The complexity of quantum spin systems on a two-dimensional square lattice. *Quant. Inf. Comput.*, 8:900, 2009.
- [13] Norbert Schuch and Frank Verstraete. Interacting electrons, density functional theory, and quantum merlin arthur. *Nature Physics*, 5:732, 2009.
- [14] Leslie G. Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8:189–201, 1979.