

Extremal eigenvalues of local Hamiltonians

Aram W. Harrow¹ and Ashley Montanaro²

¹Center for Theoretical Physics, Department of Physics, MIT

²School of Mathematics, University of Bristol, UK

April 24, 2017

We apply classical algorithms for approximately solving constraint satisfaction problems to find bounds on extremal eigenvalues of local Hamiltonians. We consider spin Hamiltonians for which we have an upper bound on the number of terms in which each spin participates, and find extensive bounds for the operator norm and ground-state energy of such Hamiltonians under this constraint. In each case the bound is achieved by a product state which can be found efficiently using a classical algorithm.

1 Introduction

The eigenvalue statistics of a local Hamiltonian are related to its structure. One example is the level spacings of chaotic vs integrable systems, which can be seen as the small-scale structure of the spectrum. What about the large-scale features, such as the extremal eigenvalues? Do these scale differently for non-interacting or interacting systems? It is generally understood that interacting systems can be frustrated, meaning that all the local terms cannot simultaneously be in their ground state. This situation is generically true for local terms with entangled ground states. But how much of an effect can frustration have on the ground-state energy of a system?

Here we study the extremal eigenvalues of quantum Hamiltonians which are only weakly interacting, in the sense that they can be written as sums of terms where each term depends only on a few qubits, and each qubit is included in only a few terms. With this mild form of locality imposed, how far apart must the largest and smallest eigenvalues be? If the Hamiltonian were non-interacting, the separation should scale with the size of the system. For a more general Hamiltonian, the extremal eigenvectors may be highly entangled and interacting terms may contribute opposite signs. Nevertheless, in this paper we show lower bounds on the norms of local Hamiltonians under very general conditions. An additional argument shows specifically that the ground-state energy is low (or if desired, that the top eigenvalue is high).

Aram W. Harrow: aram@mit.edu

Ashley Montanaro: ashley.montanaro@bristol.ac.uk

Theorem 1. *Let H be a traceless k -local Hamiltonian on n qubits such that $k = O(1)$. Assume that H can be expressed as a weighted sum of m distinct Pauli terms such that each term has operator norm $\Theta(1)$, and each qubit participates in at most ℓ terms. Then $\|H\| \geq \Omega(m/\sqrt{\ell})$ and $\lambda_{\min}(H) \leq -\Omega(m/\ell)$. In each case the bound is achieved by a product state which can be found efficiently using a classical algorithm.*

In the above theorem, $\|H\|$ is the operator norm of H and $\lambda_{\min}(H)$ is the lowest eigenvalue (ground-state energy) of H . (Of course a similar statement could also be made about λ_{\max} . We focus on λ_{\min} because of its relevance to physical systems and to constraint satisfaction problems.) The notation $O(f(x)), \Omega(f(x))$ refers to functions that are $\leq cf(x)$ or $\geq cf(x)$, respectively, for some absolute constant c . We write $\Theta(f(x))$ to mean both a function that is both $O(f(x))$ and $\Omega(f(x))$.

If H were a non-interacting Hamiltonian ($k = \ell = 1$) the largest and smallest eigenvalues would both be $\Theta(m) = \Theta(n)$. Thus Theorem 1 can be viewed as saying that interaction can reduce the norm of H by at most a $O(\sqrt{\ell})$ factor and can reduce the smallest eigenvalue by at most a $O(\ell)$ factor. Observe that the bound on $\lambda_{\min}(H)$ is $-\Omega(n)$ for all lattice Hamiltonians. This proves that for any such system the ground-state energy is smaller than the average energy by an extensive amount. By contrast, using our information about $\text{tr } H^2$ alone would only show that $\lambda_{\min}(H) \leq -\Omega(\sqrt{m})$, which is in general a vanishing fraction of system size.

The restriction to terms of weight $\Theta(1)$ in Theorem 1 is not essential and is only included to simplify the bounds. Further, the hidden constants are not overly large for small k ; more precise statements of our results are given below. For example, for 2-local qubit Hamiltonians, the precise bound on λ_{\min} we obtain is $\lambda_{\min}(H) \leq -\|\hat{H}\|_1/(24\ell)$, where $\|\hat{H}\|_1$ is the sum of the absolute values of the coefficients in the Pauli expansion of H . As a simple instance where this bound can be applied, consider the antiferromagnetic Heisenberg model $H = \sum_{\langle i,j \rangle} X_i X_j + Y_i Y_j + Z_i Z_j$ on a regular lattice with n vertices. Then, for *any* such lattice, we obtain $\lambda_{\min}(H)/n \leq -1/48$.

Theorem 1 can be applied to qudit Hamiltonians with local dimension $d > 2$ by embedding each subsystem in $\lceil \log_2 d \rceil$ qubits at the expense of increasing

the locality from k to $k\lceil\log_2 d\rceil$.

Proof outline. Both results that make up Theorem 1 are based on the use of a correspondence between local quantum Hamiltonians and low-degree polynomials, which allows us to apply classical approximation algorithms for constraint satisfaction problems. This correspondence uses a qubit 2-design [5, 13] to convert arbitrary qubit Hamiltonians to polynomials over boolean variables.

The operator norm bound in Theorem 1 (stated more precisely as Lemma 3 below) is based on recent work of Barak et al. [2] which gives an efficient randomised algorithm for satisfying a relatively large fraction of a set of linear equations over \mathbb{F}_2 . The bound on λ_{\min} (stated more precisely as Lemma 5 below) is based on analysing a natural greedy algorithm which is similar to a classical algorithm of Håstad [14]. Our results can be seen as generalising these two classical algorithms to the quantum regime.

Other related work. Bansal, Bravyi and Terhal [1] have previously shown that, for 2-local qubit Hamiltonians H on a planar graph with Pauli interactions of weight $\Theta(1)$, $\lambda_{\min}(H) \leq -\Omega(m)$. Similarly to our result, their proof uses a mapping between quantum and classical Hamiltonians and proves the existence of a product state achieving a $-\Omega(m)$ bound. However, the two results are not comparable; ours holds for non-planar graphs and k -local Hamiltonians for $k > 2$, while theirs encompasses two-local Hamiltonians on planar graphs with vertices of arbitrarily high degree. The quantum-classical mapping used is also different. Finally, the constants in our results are somewhat better (for example, they obtain $\lambda_{\min}(H)/n \leq -1/135$ for the antiferromagnetic Heisenberg model on a 2D triangular lattice).

This work was motivated by [2] (whose main result is presented in Section 3). Ref. [2] in turn was inspired by [7, 8], which gives a quantum algorithm for finding low-energy states of classical Hamiltonians. The relative performance of these different algorithms (ours/[2] vs. [7, 8]) is in general unknown, and it is also open to determine the extent to which [7] can be generalised to finding low-energy states of local Hamiltonians.

One other related work is [3], which showed that when $k = 2$ and the degree of the interaction graph is large, then product states can provide a good approximation for any state, with respect to the metric given by averaging the trace distance over the pairs of systems acted on by the Hamiltonian. In particular this means they can approximate the minimum and maximum eigenvalues. Both our result and [3] yield similar error bounds (ours are somewhat tighter), but in this sense apply to incomparable settings: [3] show that product states nearly match the energy of some other state (e.g. the true ground state) with possibly unknown energy while our paper puts explicit bounds on the maximum and/or minimum energy.

Another way to think about our work is as showing that interacting spins must nevertheless behave in some ways like noninteracting spins. In this picture, some vaguely related work is [10, 4], which show that under some conditions lattice systems have a density of states that is approximately Gaussian. These results are incomparable to ours, even aside from the different assumptions, because we put bounds on the extremal eigenvalues while they study the density of states and/or the thermal states at nonzero temperature. Theorem 4 of Ref. [11] also bounded the density of states of k -local Hamiltonians under general conditions, but in the opposite direction: i.e. putting upper bounds on how many eigenvalues could have large absolute value.

Why product states? Ground states of local Hamiltonians may be highly entangled [9]. But our bounds on $\|H\|$ and $\lambda_{\min}(H)$ are achieved only with product states. One reason for this in the case of $\|H\|$ is that we are using *random* states, and product states have much larger fluctuations than generic entangled states. Indeed the variance of $\langle\psi|H|\psi\rangle$ for a random unit vector $|\psi\rangle$ is only $O(m/2^n)$. It is an interesting open question to find a distribution over entangled states that improves the constant factors in Theorem 1 that we achieve with product states.

Fourier analysis of boolean functions. We will need some basic facts from classical Fourier analysis of boolean functions [12]. Any function $f : \{\pm 1\}^n \rightarrow \mathbb{R}$ can be written as

$$f(x) = \sum_{S \subseteq [n]} \hat{f}(S)x_S, \quad (1)$$

where $x_S := \prod_{i \in S} x_i$ and $[n] := \{1, \dots, n\}$. This is known as the Fourier expansion of f . Parseval's equality implies that

$$\text{Var}(f) := \mathbb{E}_x[f(x)^2] - \mathbb{E}_x[f(x)]^2 = \sum_{S \neq \emptyset} \hat{f}(S)^2, \quad (2)$$

where the expectation is taken over the uniform distribution on $\{\pm 1\}^n$. In addition, $\hat{f}(\emptyset) = \mathbb{E}_x[f(x)]$. The influence of the j 'th coordinate on f is defined as

$$\text{Inf}_j(f) = \sum_{S \ni j} \hat{f}(S)^2.$$

2 The quantum-classical correspondence

Let H be a k -local Hamiltonian which has Pauli expansion

$$H = \sum_{s \in \{I, X, Y, Z\}^n} \hat{H}_s s_1 \otimes s_2 \otimes \dots \otimes s_n$$

for some weights \hat{H}_s that we can view as a Fourier expansion of H analogous to that in (1). Define the

norms $\|\widehat{H}\|_p := (\sum_s |\widehat{H}_s|^p)^{1/p}$ for $p \geq 1$. In order to apply classical bounds to extremal eigenvalues of H , we observe that the action of a k -local Hamiltonian on product states corresponds to a low-degree polynomial. Define the following set of states [13, 16]:

$$|\psi_{++}\rangle = \frac{1}{\sqrt{6}} \left(\sqrt{3 + \sqrt{3}}|0\rangle + e^{i\pi/4}\sqrt{3 - \sqrt{3}}|1\rangle \right),$$

$$|\psi_{-+}\rangle = Z|\psi_{++}\rangle, |\psi_{+-}\rangle = X|\psi_{++}\rangle, |\psi_{--}\rangle = Y|\psi_{++}\rangle.$$

These four states form a qubit 2-design; equivalently, a symmetric informationally-complete quantum measurement (SIC-POVM) on one qubit [13]. This measurement was studied in detail in [16]. Geometrically, the states describe a tetrahedron within the Bloch sphere [5].

Then define the functions $\chi_s : \{\pm 1\}^2 \rightarrow \mathbb{R}$, for $s \in \{I, X, Y, Z\}$, by

$$\chi_s(x) = \langle \psi_x | s | \psi_x \rangle.$$

These functions are pleasingly simple: one can verify that

$$\chi_I(x) = 1, \chi_X(x) = \frac{x_1}{\sqrt{3}}, \chi_Y(x) = \frac{x_1 x_2}{\sqrt{3}}, \chi_Z(x) = \frac{x_2}{\sqrt{3}}. \quad (3)$$

Split each $x \in \{\pm 1\}^{2n}$ into n consecutive blocks of length 2, written as $x = x^{(1)}x^{(2)} \dots x^{(n)}$, and define the function $f_H : \{\pm 1\}^{2n} \rightarrow \mathbb{R}$ by

$$f_H(x) = \langle \psi_{x^{(1)}} | \dots | \psi_{x^{(n)}} | H | \psi_{x^{(1)}} \rangle \dots | \psi_{x^{(n)}} \rangle$$

$$= \sum_{s \in \{I, X, Y, Z\}^n} \widehat{H}_s \chi_{s_1}(x^{(1)}) \chi_{s_2}(x^{(2)}) \dots \chi_{s_n}(x^{(n)}).$$

As each $x \in \{\pm 1\}^{2n}$ corresponds to a state $|\psi_{x^{(1)}}\rangle \dots |\psi_{x^{(n)}}\rangle$, we immediately have the bounds $\lambda_{\max}(H) \geq \max_{x \in \{\pm 1\}^{2n}} f_H(x)$ and $\lambda_{\min}(H) \leq \min_{x \in \{\pm 1\}^{2n}} f_H(x)$. We will now proceed to show bounds on $\max_{x \in \{\pm 1\}^{2n}} f_H(x)$ and $\min_{x \in \{\pm 1\}^{2n}} f_H(x)$ by viewing $f_H(x)$ as a polynomial.

As H is k -local and each function χ_s ($s \neq I$) is a monomial of degree at most 2, f_H is a polynomial of degree at most $2k$. Because the Fourier expansion of each function χ_s contains only one term, each term in H corresponds to exactly one term in the Fourier expansion of f_H . Indeed

$$\widehat{f}_H(s) = \widehat{H}_s 3^{-|s|/2}, \quad (4)$$

where $s \in \{I, X, Y, Z\}^n$ and $|s| = |\{i : s_i \neq I\}|$. This corresponds to identifying $\{I, X, Y, Z\}^n$ with subsets of $[2n]$ in the natural way. Thus by eqns. (2), (3) and (4) we have

$$\text{Var}(f_H) = \sum_{s \in \{I, X, Y, Z\}^n, s \neq I^n} \widehat{H}_s^2 3^{-|s|} \quad (5)$$

$$\text{Inf}_j(f_H) = \sum_{s, s_j \neq I} \widehat{H}_s^2 3^{-|s|}.$$

3 Operator norm bounds

We will use the following result of Barak et al. [2], which is a constructive version of a probabilistic bound previously shown by Dinur et al. [6]:

Theorem 2 (Barak et al. [2]). *There is a universal constant C and a randomised algorithm such that the following holds. Let $f : \{\pm 1\}^n \rightarrow \mathbb{R}$ be a polynomial with degree at most k such that $\text{Var}(f) = 1$. Let $t \geq 1$ and suppose that $\text{Inf}_i(f) \leq C^{-k} t^{-2}$ for all $i \in [n]$. Then with high probability the algorithm outputs $x \in \{\pm 1\}^n$ such that $|f(x)| \geq t$. The algorithm runs in time $\text{poly}(m, n, \exp(k))$, where m is the number of nonzero monomials in f .*

Recent independent work of Håstad [15] describes an alternative, randomised algorithm achieving a similar bound.

Given the quantum-classical correspondence discussed in the previous section, we can now apply Theorem 2 to f_H to prove the following result, which is one half of Theorem 1.

Lemma 3. *There is a universal constant D and a randomised classical algorithm such that the following holds. Let H be a traceless k -local Hamiltonian given as a weighted sum of m Pauli terms such that, for all j , $\text{Inf}_j(f_H) \leq I_{\max}$. Then with high probability the algorithm outputs a product state $|\psi\rangle$ such that $|\langle \psi | H | \psi \rangle| \geq D^{-k} \|\widehat{H}\|_2^2 / \sqrt{I_{\max}}$. The running time of the algorithm is $\text{poly}(m, n, \exp(k))$.*

Proof. First observe that if we simply pick $x \in \{\pm 1\}^{2n}$ uniformly at random and consider the corresponding product state $|\psi_x\rangle$,

$$\mathbb{E}_x[\langle \psi_x | H | \psi_x \rangle^2] = \text{Var}(f_H) \geq \frac{\|\widehat{H}\|_2^2}{3^k}$$

by (5). In addition (see e.g. [12, Theorem 9.24]), as f_H is a degree- $2k$ polynomial,

$$\Pr_x[|f_H(x)| \geq \sqrt{\text{Var}(f_H)}] \geq \exp(-O(k)).$$

Therefore, simply picking $\exp(O(k))$ random product states of the form $|\psi_x\rangle$ achieves $|\langle \psi_x | H | \psi_x \rangle| \geq \|\widehat{H}\|_2 / 3^{k/2}$ with high probability. Let E be a universal constant to be chosen later. If $I_{\max} \geq E^{-k} \|\widehat{H}\|_2^2$, then $|\langle \psi_x | H | \psi_x \rangle| \geq (\sqrt{3E})^{-k} \|\widehat{H}\|_2^2 / \sqrt{I_{\max}}$ as desired, taking $D = \sqrt{3E}$. So assume henceforth that $I_{\max} \leq E^{-k} \|\widehat{H}\|_2^2$. Let $f'_H = f_H / \sqrt{\text{Var}(f_H)}$ so $\text{Var}(f'_H) = 1$. Then

$$\text{Inf}_j(f'_H) = \frac{\text{Inf}_j(f_H)}{\text{Var}(f_H)} \leq 3^k \frac{I_{\max}}{\|\widehat{H}\|_2^2}.$$

Set

$$t = \frac{C^{-k/2}}{\sqrt{\max_i \text{Inf}_i(f'_H)}} \geq E^{-k/2} \frac{\|\widehat{H}\|_2}{\sqrt{I_{\max}}} \geq 1,$$

where C is the constant in Theorem 2 and we choose E large enough for the first inequality to hold. Then the algorithm of Theorem 2 outputs $x \in \{\pm 1\}^{2^n}$ such that $|f'_H(x)| \geq E^{-k/2} \|\widehat{H}\|_2 / \sqrt{I_{\max}}$. Renormalising again by multiplying by $\sqrt{\text{Var}(f_H)} \geq 3^{-k/2} \|\widehat{H}\|_2$, $|\langle \psi_x | H | \psi_x \rangle| = |f_H(x)| \geq (\sqrt{3E})^{-k} \|\widehat{H}\|_2^2 / \sqrt{I_{\max}}$, which completes the proof.

Note that the algorithm does not need to know whether I_{\max} is large or not, since it can simply try both strategies and see which one results in the larger value of $|\langle \psi | H | \psi \rangle|$. \square

The first part of Theorem 1 is now immediate from Lemma 3. Let H be a k -local Hamiltonian with $k = O(1)$ such that H is a sum of m distinct Pauli terms, each of weight $\Theta(1)$, with each qubit participating in ℓ terms. Then $\|\widehat{H}\|_2^2 = \Theta(m)$, $\text{Inf}_j(f_H) = O(\ell)$.

4 Bounds on extremal eigenvalues

We now describe an algorithm for bounding extremal eigenvalues which is weaker, but holds for both the largest and smallest eigenvalues. Once again, the algorithm is based on applying the quantum-classical correspondence in Section 2 to a classical algorithm. We first describe the classical algorithm, which is a simple greedy approach to find large values taken by a low-degree polynomial on the boolean cube.

Given $f : \{\pm 1\}^n \rightarrow \mathbb{R}$ define $W := W(f) = \sum_{S \neq \emptyset} |\widehat{f}(S)|$. Assume that, for all $i \in [n]$,

$$|\{T \subseteq [n] : \widehat{f}(T) \neq 0 \text{ and } i \in T\}| \leq \ell.$$

Consider the following algorithm, based on ideas of [14] but somewhat simpler:

1. Find S such that $|\widehat{f}(S)|$ is maximal.
2. Substitute values for x_i , $i \in S$, such that $\widehat{f}(\emptyset)$ increases by at least $|\widehat{f}(S)|$.
3. Repeat until f is constant; call this constant f_{end} .

It is not obvious that step (2) can be achieved, because there might exist $T \subsetneq S$ such that $\widehat{f}(T) \neq 0$. Define a function f_S by $f_S(x) = \sum_{T \subsetneq S} \widehat{f}(T) x_T$. For each $T \subsetneq S$ such that $T \neq \emptyset$ and each $a \in \{\pm 1\}$, $\mathbb{E}_{x, x_S=a} [x_T] = 0$. So $\mathbb{E}_{x, x_S=\text{sgn}(\widehat{f}(S))} [f_S(x)] = \widehat{f}(\emptyset) + |\widehat{f}(S)|$, and there must exist some y achieving $f_S(y) \geq \widehat{f}(\emptyset) + |\widehat{f}(S)|$. Searching over at most 2^k different values x is sufficient to find y .

Lemma 4. *When the above algorithm terminates, $f_{\text{end}} \geq \widehat{f}(\emptyset) + W/(2k\ell)$.*

Proof. Let f_j be the new function produced at the j 'th stage of the algorithm, with $f_0 = f$. Let M_j be the value of $|\widehat{f}_{j-1}(S)|$ corresponding to the set S chosen at stage j . Then $\widehat{f}_j(\emptyset) = \widehat{f}_{j-1}(\emptyset) + M_j$ and

$W(f_j) \geq W(f_{j-1}) - 2k\ell M_j$. The latter inequality is shown as follows. For each $i \in S$, there are at most ℓ subsets T such that $\widehat{f}_{j-1}(T) \neq 0$ and $i \in T$. So there are at most $k\ell$ subsets T such that $T \cap S \neq \emptyset$. For each such T , the substitution of values x_i , $i \in S$, implies that $\widehat{f}_j(T)$ is set to 0, and for some other subset T' , $\widehat{f}_j(T') = \widehat{f}_{j-1}(T') \pm \widehat{f}_{j-1}(T)$. These are the only coefficients modified by the substitution process. Thus $W(f_j)$ can only decrease by at most $2|\widehat{f}_{j-1}(T)| \leq 2M_j$ for each T such that $S \cap T \neq \emptyset$. Rearranging we have $M_j \geq (W(f_{j-1}) - W(f_j))/(2k\ell)$ for $j \geq 1$ and hence

$$\begin{aligned} f_{\text{end}} &= \widehat{f}(\emptyset) + \sum_j M_j \\ &\geq \widehat{f}(\emptyset) + \frac{1}{2k\ell} \sum_j W(f_{j-1}) - W(f_j) \\ &= \widehat{f}(\emptyset) + \frac{W}{2k\ell} \end{aligned}$$

as claimed. \square

The following lemma is now essentially immediate.

Lemma 5. *There is a universal constant E and a deterministic classical algorithm such that the following holds. Let H be a traceless k -local Hamiltonian which can be written as a weighted sum of m distinct Pauli terms such that each qubit participates in at most ℓ terms. Then the algorithm outputs a product state $|\psi\rangle$ such that $\langle \psi | H | \psi \rangle \geq E^{-k} \|\widehat{H}\|_1 / \ell$. The algorithm runs in time $\text{poly}(m, n, \exp(k))$.*

Proof. Apply the algorithm of Lemma 4 to the degree- $2k$ polynomial $f_H : \{\pm 1\}^{2^n} \rightarrow \mathbb{R}$ defined as in Section 2. We have $W(f_H) \geq 3^{-k/2} \|\widehat{H}\|_1$, $\widehat{f}_H(\emptyset) = 0$. Hence the algorithm finds $|\psi\rangle$ such that $\langle \psi | H | \psi \rangle \geq 3^{-k/2} \|\widehat{H}\|_1 / (4k\ell) \geq E^{-k} \|\widehat{H}\|_1 / \ell$ for a large enough universal constant E . \square

Applying the same procedure to $-H$ is of course sufficient to also find $|\psi\rangle$ such that $\langle \psi | H | \psi \rangle \leq -E^{-k} \|\widehat{H}\|_1 / \ell$. Observing that $\|\widehat{H}\|_1 = \Theta(m)$ for a Hamiltonian H which is a sum of m distinct Pauli terms of weight $\Theta(1)$ completes the proof of the second part of Theorem 1.

5 Optimality

Both bounds in Theorem 1 are tight, even for classical Hamiltonians, as demonstrated by the following examples. The first is based on an example in [2]. Consider the 2-local Hamiltonian on n qubits

$$H = \sum_{(i,j) \in E} \widehat{H}_{ij} Z_i Z_j,$$

where E is the edges of an arbitrary r -regular undirected graph on n vertices, and each weight $\widehat{H}_{ij} \in$

$\{\pm 1\}$ is picked uniformly at random. Then $m = rn/2$, $\ell = r$. For each fixed $x \in \{0, 1\}^n$,

$$\langle x|H|x \rangle = \sum_{(i,j) \in E} \widehat{H}_{ij} (-1)^{x_i+x_j}$$

is a sum of $rn/2$ uniformly distributed elements of $\{\pm 1\}$. Via a standard Chernoff bound argument,

$$\Pr_H[|\langle x|H|x \rangle| \geq t] \leq 2e^{-\frac{t^2}{rn}}.$$

Fixing $t = \Theta(n\sqrt{r})$ and taking a union bound over all $x \in \{0, 1\}^n$, $\|H\| = O(n\sqrt{r}) = O(m/\sqrt{\ell})$ with high probability.

Second, consider the 2-local Hamiltonian on n qubits

$$H = \sum_{i < j} Z_i Z_j.$$

Then H is a sum of $m = \Theta(n^2)$ Pauli terms of weight 1, where each qubit participates in $\ell = \Theta(n)$ terms. We have

$$\langle x|H|x \rangle = \sum_{i < j} (-1)^{x_i+x_j} = \frac{(n - 2|x|)^2 - n}{2},$$

so $\lambda_{\min}(H) \geq -n/2 = -\Theta(m/\ell)$.

It is an open question whether both the bounds on $\|H\|$ and $\lambda_{\min}(H)$ can be saturated at the same time.

Acknowledgements

We would like to thank Oded Regev for the key idea behind the proof of Lemma 3. This work was initiated at the BIRS workshop 15w5098, ‘‘Hypercontractivity and Log Sobolev Inequalities in Quantum Information Theory’’. We would like to thank BIRS and the Banff Centre for their hospitality. AM was supported by the UK EPSRC under Early Career Fellowship EP/L021005/1. AWH was funded by NSF grants CCF-1111382 and CCF-1452616 and ARO contract W911NF-12-1-0486. No new data were created during this study.

References

[1] N. Bansal, S. Bravyi, and B. Terhal. Classical approximation schemes for the ground-state energy of quantum and classical Ising spin hamiltonians on planar graphs. *Quantum Inf. Comput.*, 9(8):0701–0720, 2009, [arXiv:0705.1115](#).

[2] B. Barak, A. Moitra, R. O’Donnell, P. Raghavendra, O. Regev, D. Steurer, L. Trevisan, A. Vijayaraghavan, D. Witmer, and J. Wright. Beating the random assignment on constraint satisfaction problems of bounded degree, 2015, [arXiv:1505.03424](#).

[3] F. G. S. L. Brandão and A. W. Harrow. Product-state approximations to quantum ground states. *Communications in Mathematical Physics*, 342(1):47–80, 2006. [arXiv:1310.0017](#).

[4] F. G. S. L. Brandao and M. Cramer. Equivalence of statistical mechanical ensembles for non-critical quantum systems, 2015, [arXiv:1502.03263](#).

[5] P. Delsarte, J. Goethals, and J. Seidel. Spherical codes and designs. *Geometriae Dedicata*, 6(3):363–388, 1977.

[6] I. Dinur, E. Friedgut, G. Kindler, and R. O’Donnell. On the Fourier tails of bounded functions over the discrete cube. *Israel Journal of Mathematics*, 160:389–412, 2007.

[7] E. Farhi, J. Goldstone, and S. Gutmann. A quantum approximate optimization algorithm. Technical Report MIT-CTP/4610, MIT, 2014, [arXiv:1411.4028](#).

[8] E. Farhi, J. Goldstone, and S. Gutmann. A quantum approximate optimization algorithm applied to a bounded occurrence constraint problem. Technical Report MIT-CTP/4628, MIT, 2014, [arXiv:1412.6062](#).

[9] M. H. Freedman and M. B. Hastings. Quantum systems on non- k -hyperfinite complexes: A generalization of classical statistical mechanics on expander graphs, 2013, [arXiv:1301.1363](#).

[10] M. Hartmann, G. Mahler, and O. Hess. Gaussian quantum fluctuations in interacting many particle systems. *Letters in Mathematical Physics*, 68(2):103–112, 2004, [arXiv:math-ph/0312045](#).

[11] A. Montanaro. Some applications of hypercontractive inequalities in quantum information theory. *Journal of Mathematical Physics*, 53(12):122206, 2012, [arXiv:1208.0161](#).

[12] R. O’Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014.

[13] J. Renes, R. Blume-Kohout, A. J. Scott, and C. Caves. Symmetric informationally complete quantum measurements. *J. Math. Phys.*, 45(6):2171–2180, 2003. [quant-ph/0310075](#).

[14] J. Håstad. On bounded occurrence constraint satisfaction. *Inf. Proc. Lett.*, 74:1–6, 2000.

[15] J. Håstad. Improved bounds for bounded occurrence constraint satisfaction, 2015. <http://www.csc.kth.se/~johanh/bounded2.pdf>.

[16] J. Řeháček, B.-G. Englert, and D. Kaszlikowski. Minimal qubit tomography. *Phys. Rev. A*, 70:052321, 2004. [quant-ph/0405084](#).