Quantum Information and Computation, Vol. 11, No. 11&12 (2011) 1019–1027

A NOTE ABOUT A PARTIAL NO-GO THEROREM FOR QUANTUM PCP

ITAI ARAD[a]

School of Computer Science and Engineerig The Hebrew University
Jerusalem, Israel

Received March 1, 2011
Revised September 5, 2011

This is not a disproof of the quantum PCP conjecture! In this note we use perturbation on the commuting Hamiltonian problem on a graph, based on results by Bravyi and Vyalyi [1], to provide a very partial no-go theorem for quantum PCP. Specifically, we derive an upper bound on how large the promise gap can be for the quantum PCP still to hold, as a function of the non-commuteness of the system. As the system becomes more and more commuting, the maximal promise gap shrinks. We view these results as possibly a preliminary step towards disproving the quantum PCP conjecture posed in [2]. A different way to view these results is actually as indications that a critical point exists, beyond which quantum PCP indeed holds; in any case, we hope that these results will lead to progress on this important open problem.

Keywords: Quantum PCP, Local Hamiltonians, QSAT

Communicated by: I Cirac & M Mosca

1 Introduction

The PCP (Probabilistically Checkable Proofs) theorem is arguably the most important development in computational complexity over the last two decades. It roughly states that any mathematical proof can be translated efficiently into a proof of a new form with a comparable length, such that a simple test on $\mathcal{O}(1)$ bits chosen randomly from the new proof can decide with high probability whether the proof is correct or not – hence the acronym PCP. Alternatively, it says that given a Constraint Satisfaction Problem (CSP), it can be efficiently replaced by one with a comparable size, such that if the original one was satisfiable, then so is the new one, whereas if for the original CSP every assignment must violate at least one constraint, then any assignment to the new CSP must violate a *constant fraction* of the constraints! An important consequence is what is called "hardness of approximation": the problem of deciding whether a CSP is satisfiable or a constant fraction of its constraints must be violated, is NP-hard.

Is there a quantum analogue to this remarkable theorem? This is perhaps the most important open question in quantum Hamiltonian complexity, and one of the central problems in quantum complexity in general. Both a proof or a disproof of this conjecture would arguably yield deep insights into the basic notions of quantum mechanics, such as entanglement, no-cloning of information, and the quantum to classical transition on large scales.

[a]email: arad.itai@gmail.com

In this note we present a result that can be seen as a weak evidence *against* the existence of a quantum PCP theorem. Hopefully, it may serve as a starting point for a more general framework for disproving this conjecture, or alternatively, as a starting point for better clarifications of the conditions for a quantum PCP theorem to hold. Before stating the result, we first state what is meant by a quantum PCP theorem.

1.1 Background on the quantum PCP conjecture

The quantum PCP conjecture was first stated formally in Ref. [2]. Here we shall roughly follow this presentation. The quantum analog of a classical CSP is the QSAT problem, which is a special instance of the local-Hamiltonian problem. In that problem, we are give a k-local Hamiltonian over a system of n qubits $H = \sum_{i=1}^{M} Q_i$ that is made of $M = \text{poly}(n)$ k-local projections Q_i. We are promised that the ground energy of the system is either 0 (all quantum constrains are satisfied) or it is above some constant $a = 1/\text{poly}(n)$. Just like its classical counter-part, this problem is known to be quantum NP-complete.[b]

A quantum PCP theorem would state that even if $a = rM$, for some constant $0 < r < 1$, the problem is still quantum-NP hard to decide. In other words, it is quantum-NP hard to distinguish between the case when the system is completely satisfiable, or when, roughly, a fraction r of it can not be satisfied.

Formally, we define

Definition 1 (The r-gap k-QSAT problem) *Let $r \in (0,1)$ be some constant. We are given a k-local Hamiltonian $H = \sum_{i=1}^{M} Q_i$ over n qubits, where the Q_i are k-local projections and $M = \text{poly}(n)$. We are promised that the ground energy of H is either 0 (a YES instance) or is greater than rM (a NO instance). We are asked to decide which is which.*

Then the Quantum PCP conjecture is

Conjecture 2 (Quantum PCP) *There exist constants (k, r) for which Problem in Def. 1 is quantum-NP hard.*

To prove that such a problem is quantum-NP hard, we would like to show an efficient reduction of another quantum-NP-hard problem to it. It is natural to start with the k-QSAT problem that was described above. We would like to find an efficient transformation that takes a k-QSAT problem and turns it into a r-gap k-QSAT problem such that if the original system was satisfiable (ground energy is zero), then so is the new system. On the other hand, if it was not satisfiable, with a ground energy above $a = 1/\text{poly}$, then the ground energy of the new system would be above rM.

This type of transformation is called "Gap Amplification", because it amplifies the promise gap of the problem. It is precisely this type of transformation that was used iteratively in Dinur's proof of the classical PCP theorem [3]. Let us describe this transformation in some more details. Dinur achieves gap amplification by an iterative process that amplifies gap by some constant factor > 1 at each round. Each such iteration is made of 3 steps. One step amplifies the promise gap of the system at the expense of making it much less local. The purpose of the two other steps is to fix this, restoring the locality of the system, without

[b] Note that in this paper, when we refer to the complexity class quantum-NP, we formally mean the class QMA_1 (i.e., QMA with one-sided error), and not QMA, which is the more natural "quantum-NP" class. This is mostly done in order to keep this note simple and avoid unnecessary technical difficulties. Nevertheless, we note that our main result, Theorem 4, can be easily generalized to the QMA settings (see the comment below the statement of the theorem). For a technical discussion on the QMA_1 class, we refer the reader to Ref. [1].

compromising too much the gap amplification of the first step.

The entire process is carried over a CSP that is defined on an expander graph, and in addition, the entire proof is very combinatorial in nature. These two facts make it a promising outline for a quantum proof under the natural mapping of classical constraints to projections. Indeed, a first step in that direction was taken in Ref. [2], where it was shown that essentially the amplification step in Dinur's proof can be done also quantumly. Like the classical proof, the quantum proof relies on expander graphs. The quantization of the two other steps, however, remains an open problem.

1.2 Reasons for doubts in a quantum PCP theorem

In the attempts to prove the quantum analogue of Dinur's proof, it seems hard to quantize any classical step that increases the size of the system (and of the witness); such steps seem to conflict with the quantum no-cloning principle. See Ref. [2] for more details; nevertheless, such increase seems unavoidable in the classical case.

Except for these difficulties, there is another reason to believe that there is no quantum PCP theorem, pointed to me by Hastings [4]. As we have seen, such theorem implies the existence of systems in which it is quantum-NP hard to distinguish between a vanishing ground energy and ground energy of the order of the system size. From a physicist point of view, it is equivalent to determining whether or not the free energy of the system becomes negative at a *finite* temperature [5]. It is then argued that at such temperatures, on large scale, the system loses its quantum characteristics; long-range entanglement effects must fade. Consequently, the system can be described (approximately) classically, hence the problem is inside NP. Of course, one should keep in mind that the above intuition is based on physical Hamiltonians that are defined on regular grids, and might be wrong for specially crafted Hamiltonians on general graphs such as expanders.

In this note, we will pursue this direction. The ultimate goal is to show that for any (k, r), the Problem in Def. 1 is inside NP. This, however, seems very difficult. Instead of attacking it directly, it might be beneficial to show first that a more restricted problem is inside NP. This is what we do here.

1.3 Results: A partial No-Go theorem for quantum PCP

We are interested in a version of the problem in Def. 1 in which the projections are two-local, sitting on the edges of a D-regular graph:

Definition 3 (The (d, D, r)-gap Hamiltonian problem on a graph) *We consider a QSAT system $H = \sum_{i=1}^{M} Q_i$ that is defined on a D-regular graph, using d-dimensional qudits that sit on its vertices, and projections $\{Q_i\}$ that sit on its edges. We are promised that the ground energy of H is either 0 or is greater than rM for some constant $0 < r < 1$, and we are asked to decide which is which.*

The advantage of working with this restricted set of problems is three-fold. First, its classical analog is the outcome of Dinur's classical PCP proof [3]. It is therefore a natural candidate for a quantum PCP construction. Second, it has a classical, yet non-trivial limit, which was discovered by Bravyi and Vyalyi [6]: When the projections commute, the problem becomes classical in the sense that the ground state of the system can be described by a shallow tensor-network that can be contracted efficiently on a classical computer. This tensor-network can be given as a witness to the prover, and hence the problem is inside NP. Finally, by itself, this class of systems seems rich enough to capture non-trivial quantum effects, if exist. For example, with a 1/poly promise gap, these systems become quantum-NP hard, as can be seen by the results of Ref. [7][c]

In this note we will go slightly beyond that classical limit. We will show that for a system which is only slightly non-commuting, the problem in Def. 3 is inside NP for sufficiently large r's. In other words, there cannot be a quantum PCP construction that yields such slightly non-commuting systems with such large r's.

This is the main theorem we wish to prove:

Theorem 4 *For the set of QSAT systems that are defined on a D-regular graph using d-dimensional qudits, the following holds: if for every two projections,*

$$\|[Q_i, Q_j]\| \le \delta \ , \tag{1}$$

and if the system is satisfiable (has a ground energy 0), then there exists an efficiently contractable tensor network with energy $\le \epsilon M$*, where* $0 < \epsilon < 1$ *depends only on* d, D, δ*, and* $\epsilon \to 0$ *as* $\delta \to 0$*. Consequently, the problem in Def. 3 is inside* NP *when its projections satisfy Eq. (1) and* $r > 2\epsilon$*.*

The idea of the proof is very simple. Using the assumption that the projections in H nearly commute, we will find an auxiliary *commuting* system, $\hat{H} = \sum_{i=1}^{M} \hat{Q}_i$ such that $\|Q_i - \hat{Q}_i\| \le \epsilon/2$ for every i. By Ref. [6], this system has a ground state with an efficient description, that can thus be provided as a classical witness to the NP verifier. The point is that this state is an $M\epsilon$ approximation of the ground state, and thus can provide an NP witness for a $rM \ge 2\epsilon M$ approximation of the ground energy.

It is worth noting that Theorem 4 can be trivially generalized to any eigenvalue of the original system, not just the frustration-free case where all the local energies are zero. This implies that our result can also be stated in terms of the QMA class with 2-sided errors (instead of the QMA_1 class that is used here as quantum-NP), but we will not pursue this direction here.

1.4 Further research

It is interesting to see if the result of this note can be strengthened. For a start, functional dependence of ϵ on d, D, δ is not given here, but is probably not too difficult to find by generalizing the results of Ref. [8].

To show that the problem in Def. 3 is inside NP for every (d, D, r) we would like to show that for $\delta = 1$, one can always find a tensor-network that yields an energy $\le \epsilon M$, for an arbitrarily small, yet constant ϵ. Such a result would be a very strong indication against quantum PCP.

There are two natural directions that may help to prove such a result. First, we note that the tensor-network that results from the construction of Bravyi and Vyalyi is a very simple one. It is essentially a depth-4 local quantum circuit. Using a classical computer, we can in fact, contract similar networks with a logarithmic depth. Can we find a perturbation theory

[c] In that paper, it is shown that the problem of a nearest-neighbor Hamiltonian on a line with qudits of size 12 is QMA-hard, yet a closer inspection reveals that the same construction can also be used to show QMA_1 hardness.

for which the depth-4 network is just the first order approximation? Then by going to higher orders, we may systematically lower ϵ for a given d, D, δ.

Related to that, one may try to find a reduction of the system to a system that is more commuting, perhaps by some sort of a coarse-graining process. This again, might lead to a larger ϵ for a given d, D, δ.

We now proceed to the proof of theorem 4.

2 Proof of theorem 4

Notation:
We use the natural inner product on the space of matrices

$$\langle A|B\rangle \stackrel{\text{def}}{=} \mathrm{Tr}(A^\dagger B) \ ,$$

which leads to the Frobenius norm

$$\|A\|_F \stackrel{\text{def}}{=} \sqrt{\mathrm{Tr}(A^\dagger A)} \ .$$

Proof of Theorem 4: Using the C^*-algebra machinery of Ref. [6], we start with the following decomposition of every 2-local projection in H.

Lemma 1 *Let Q be a 2-local projection on ${}^d \otimes {}^d$. Then one can write*

$$Q = \sum_{\alpha=1}^{d^2} A_\alpha \otimes B_\alpha \ , \tag{2}$$

with A_α and B_α working locally on one qudit, with the following properties:

1. *$\{A_\alpha\}$ are orthogonal and bounded $\|A_\alpha\|_F \le \|Q\|_F \le d$.*
2. *$\{B_\alpha\}$ are orthonormal.*
3. *The algebra generated by $\{A_\alpha\}$ is close under conjugation, and the same applies for $\{B_\alpha\}$.*

Proof: We treat Q as a vector in a bipartite Hilbert space of operators. Using the Schmidt decomposition in that space, $Q = \sum_\alpha \lambda_\alpha \cdot A'_\alpha \otimes B_\alpha$, with $\{A'_\alpha\}$ and $\{B_\alpha\}$ orthonormal, and $\sum_{\alpha=1}^{d^2} |\lambda_\alpha|^2 = \|Q\|_F^2 \le d^2$. Defining $A_\alpha \stackrel{\text{def}}{=} \lambda_\alpha A'_\alpha$ then proves 1) and 2). The third property follows from the Hermiticity of Q. $\square$

The advantage of working in this representation is that a Frobenius distance between the Q's easily translates into a Frobenius distance between the A_α. Specifically, assume the above decomposition for adjacent projections Q_i, Q_j:

$$Q_i = \sum_\alpha A_\alpha^{(i)} \otimes B_\alpha^{(i)} \qquad\qquad Q_j = \sum_\beta A_\beta^{(j)} \otimes B_\beta^{(j)} \ , \tag{3}$$

where $A^{(i)}_\alpha$ and $A^{(j)}_\beta$ operate on the same qudit, and $B^{(i)}_\alpha$ and $B^{(j)}_\beta$ on two other qudits. Then by the orthonormality of $\{B^{(i)}_\alpha\}$ and $\{B^{(j)}_\beta\}$,

$$\begin{aligned}\|[Q_i,Q_j]\|_F^2 &\stackrel{\text{def}}{=} \mathrm{Tr}\Big([Q_i,Q_j]([Q_i,Q_j])^\dagger\Big) \\ &= \sum_{\alpha,\beta}\mathrm{Tr}\Big([A^{(i)}_\alpha,A^{(j)}_\beta]([A^{(i)}_\alpha,A^{(j)}_\beta])^\dagger \\ &= \sum_{\alpha,\beta}\|[A^{(i)}_\alpha,A^{(j)}_\beta]\|_F^2 .\end{aligned} \tag{4}$$

Note that in the passage of the second equality we have traced out the qudits of the $\{B^{(i)}_\alpha\}$ and the $\{B^{(j)}_\beta\}$ support, and so the Frobenius norm in the third line is over the single qudit space – not to be confused with the Frobenius norm of the first line, which is defined on the space of 3 qudits.

We conclude that for every α,β, we get $\|[A^{(i)}_\alpha,A^{(j)}_\beta]\|_F \le \|[Q_i,Q_j]\|_F$.

The following lemma uses this decomposition to prove that if the operators $\{Q_i\}$ are slightly non-commuting, we can find a close set of operators which fully commute.

Lemma 2 *There exists a function $\epsilon(\delta)$ with the limit $\lim_{\delta\to 0}\epsilon(\delta)\to 0$ such that the following holds. Every set of operators $\{Q_i\}$, $i=1,\ldots,D$ that work on a given qudit, and for which $\|[Q_i,Q_j]\|_F \le \delta$, can be replaced by operators $\{\hat{Q}_i\}$ with $\|Q_i-\hat{Q}_i\|_F\le\epsilon(\delta)$, that in addition satisfy the following properties*

1. $[\hat{Q}_i,\hat{Q}_j]=0$
2. $\hat{Q}_i$ *are Hermitian*
3. *For any other term in the system, $\|[\hat{Q}_i,P]\|_F \le \|[Q_i,P]\|_F$, so that the system does not become less commuting at other places.*

Notice that $\delta(\epsilon)$ may depend on d,D.

Proof: We use Lemma 1 to decompose

$$Q_i=\sum_\alpha A^{(i)}_\alpha\otimes B^{(i)}_\alpha \ , \tag{5}$$

where $i\in[1,D]$ and $\alpha\in[1,d^2]$.

By Eq. (4), if $\|[Q_i,Q_j]\|_F\le\delta$, then for every $i\ne j$ and every α,β, we have $\|[A^{(i)}_\alpha,A^{(j)}_\beta]\|_F\le\delta$.

We now define the function $\eta(\delta)$ as follows:

Definition 5 (The function $\eta(\delta)$) *For every $\delta>0$, $\eta(\delta)$ is the minimal η that satisfies the following condition: For every set of d^2D operators $\{A^{(i)}_\alpha\}$ as in the decomposition Eq. (5), for which $\|[A^{(i)}_\alpha,A^{(j)}_\beta]\|_F\le\delta$ (for every $i\ne j$ and every α,β), there exists a set of operators $\{\hat{A}^{(i)}_\alpha\}$ such that $\|A^{(i)}_\alpha-\hat{A}^{(i)}_\alpha\|_F\le\eta$, and*

1. $[\hat{A}^{(i)}_\alpha,\hat{A}^{(j)}_\beta]=0$ *(for every $i\ne j$ and every α,β).*

2. *For fixed* i, $\{\hat{A}_\alpha^{(i)}\}$ *are orthogonal and* $\|\hat{A}_\alpha^{(i)}\|_F \le \|A_\alpha^{(i)}\|_F$.

3. *For fixed* i, *the algebra that is generated by* $\{\hat{A}_\alpha^{(i)}\}$ *is close under conjugation.*

We first note that $\eta(\delta)$ is well-defined because for every $\delta > 0$ there exists an $\eta > 0$ that satisfies the above requirements – simply take η large enough, and use the fact that the $\{A_\alpha^{(i)}\}$ operators are bounded (property 1 in Lemma 1), hence it is easy to pick a set of commuting $\{\hat{A}_\alpha^{(i)}\}$ that will do the job.

We now prove:

Claim 6 $\lim_{\delta\to 0}\eta(\delta) = 0$.

Proof: To prove this claim, we use a neat argument from Ref. [9] (page 76).[d] Assume by contradiction that this is not true. Then there is a $\eta_0 > 0$ and a series $\delta_n \to 0$ such that $\eta(\delta_n) \ge \eta_0$. Consequently, there is a corresponding series of operators $\{A_\alpha^{(i)}(n)\}$ such that

$$\text{for every } i \ne j \text{ and } \alpha,\beta \qquad \|[A_\alpha^{(i)}(n), A_\beta^{(j)}(n)]\|_F \le \delta_n \to 0 \ , \tag{6}$$

and at the same time, for any set operators $\{\hat{A}_\alpha^{(i)}\}$ that satisfies the 3 requirements with respect to $\{A_\alpha^{(i)}(n)\}$ (for a fixed n) there is a at least one pair (i,α) such that

$$\|A_\alpha^{(i)}(n) - \hat{A}_\alpha^{(i)}\|_F \ge \eta_0 \ . \tag{7}$$

But this is wrong because of the following reason. Since we work in a compact space (we work with $d^2 \cdot D$ bounded operators on a finite-dimensional Hilbert space), the series $\{A_\alpha^{(i)}(n)\}$ must have at least one limit point, which we denote by $\{a_\alpha^{(i)}\}$. Then we can pick n_0 for which $\|A_\alpha^{(i)}(n_0) - a_\alpha^{(i)}\| < \eta_0/4$ for every (i,α). Let us now define the set of operators $\{\hat{A}_\alpha^{(i)}\}$ as follows:

$$\hat{A}_\alpha^{(i)} \stackrel{\text{def}}{=} c_\alpha^{(i)} \cdot a_\alpha^{(i)} \ , \tag{8}$$

where $c_\alpha^{(i)}$ is a factor defined by

$$c_\alpha^{(i)} = \begin{cases} 1, & a_\alpha^{(i)} = 0 \\ \frac{\|A_\alpha^{(i)}(n_0)\|_F}{\|a_\alpha^{(i)}\|_F}, & a_\alpha^{(i)} \ne 0 \end{cases} \tag{9}$$

We now use $\{\hat{A}_\alpha^{(i)}\}$ to contradict the assumption about $\{A_\alpha^{(i)}(n_0)\}$. First, it is easy to see that for all (i,α), $\|A_\alpha^{(i)}(n_0) - \hat{A}_\alpha^{(i)}\|_F < \eta_0$. Indeed, if $a_\alpha^{(i)} = 0$ then $\|A_\alpha^{(i)}(n_0) - \hat{A}_\alpha^{(i)}\|_F = \|A_\alpha^{(i)}(n_0) - a_\alpha^{(i)}\|_F \le \eta_0/4$. Otherwise,

$$\begin{aligned}\|A_\alpha^{(i)}(n_0) - \hat{A}_\alpha^{(i)}\|_F &\le \|A_\alpha^{(i)}(n_0) - a_\alpha^{(i)}\|_F + \left|1 - \frac{\|A_\alpha^{(i)}(n_0)\|_F}{\|a_\alpha^{(i)}\|_F}\right| \cdot \|a_\alpha^{(i)}\|_F \\ &= \|A_\alpha^{(i)}(n_0) - a_\alpha^{(i)}\|_F + \left|\|A_\alpha^{(i)}(n_0)\|_F - \|a_\alpha^{(i)}\|_F\right| \\ &\le 2\|A_\alpha^{(i)}(n_0) - a_\alpha^{(i)}\|_F \le \eta_0/2 \ .\end{aligned}$$

[d]This nice trick was first brought to my attention by Matthew Hastings [10]

In addition, $\{\hat{A}^{(i)}_\alpha\}$ satisfy the 3 properties: The $\|\hat{A}^{(i)}_\alpha\|_F \le \|A^{(i)}_\alpha\|_F$ condition of property 2) is trivially satisfied by definition. To prove properties 1), 3) and the orthogonality condition of property 2), we first notice that they hold for $\{a^{(i)}_\alpha\}$: property 1) follows directly from Eq. (6), while property 3) and the orthogonality condition follow from that fact that it holds for every item $\{A^{(i)}_\alpha(n)\}$ in the series. This implies that also $\{\hat{A}^{(i)}_\alpha\}$ satisfy these properties because each of them is obtained from $\{a^{(i)}_\alpha\}$ by a non-zero multiplicative factor. We found a contradiction, and this completes the proof of claim 6. □

We now proceed to prove the lemma. The idea is to replace $A^{(i)}_\alpha \mapsto \hat{A}^{(i)}_\alpha$ in Eq. (5). We start by assuming that the $\{Q_i\}$ operators satisfy $\|[Q_i, Q_j]\|_F \le \delta$, which, as we have seen, implies that for every $i \ne j$, $\|[A^{(i)}_\alpha, A^{(j)}_\beta]\|_F \le \delta$. Then let $\{\hat{A}^{(i)}_\alpha\}$ be their commuting version from claim 6. We now use the decomposition Eq. (5) to introduce the operators

$$\tilde{Q}_i \stackrel{\text{def}}{=} \sum_\alpha \hat{A}^{(i)}_\alpha \otimes B^{(i)}_\alpha \ , \tag{10}$$

$$\hat{Q}_i \stackrel{\text{def}}{=} \frac{1}{2}\Big(\tilde{Q}_i + \tilde{Q}_i^\dagger\Big) \ , \tag{11}$$

which all work on the *same* pair of qudits.

What are the properties of these operators? First, by definition, the $\{\hat{Q}_i\}$ are Hermitian. Second, since $\{\hat{A}^{(i)}_\alpha\}$ are commuting then $[\tilde{Q}_i, \tilde{Q}_j] = 0$. Moreover, since the generated algebra of $\{\hat{A}^{(i)}_\alpha\}$ is close under conjugation, it follows that also $[\tilde{Q}_i, \tilde{Q}_j^\dagger] = 0$, which implies $[\hat{Q}_i, \hat{Q}_j] = 0$. Third, to check property 3) of the lemma, consider $Q \in \{Q_i\}$ and its replacement $\hat{Q}$, and assume P is an operator that intersects with Q but that does not belong to $\{Q_i\}$. We use Lemma 1 to write $Q = \sum_\alpha A_\alpha \otimes B_\alpha$, $\tilde{Q} = \sum_\alpha \hat{A}_\alpha \otimes B_\alpha$, $\hat{Q} = \frac{1}{2}\big(\tilde{Q} + \tilde{Q}^\dagger\big)$, $P = \sum_\beta C_\beta \otimes D_\beta$, where it is assumed that $\{B_\alpha\}$ and $\{C_\beta\}$ act on the same qudit. Then by the orthogonality of $\{A_\alpha\}, \{\hat{A}_\alpha\}$ and the orthonormality of $\{D_\beta\}$, we conclude that

$$\|[Q, P]\|_F^2 = \sum_{\alpha,\beta} \|[B_\alpha, C_\beta]\|_F^2 \cdot \|A_\alpha\|_F^2 \ ,$$

$$\|[\tilde{Q}, P]\|_F^2 = \sum_{\alpha,\beta} \|[B_\alpha, C_\beta]\|_F^2 \cdot \|\hat{A}_\alpha\|_F^2 \ ,$$

and so the fact that $\|\hat{A}_\alpha\|_F \le \|A_\alpha\|_F$ implies that $\|[\tilde{Q}, P]\|_F \le \|[Q, P]\|_F$, and since P is Hermitian also $\|[\hat{Q}, P]\|_F \le \|[Q, P]\|_F$.

Finally, let us bound the distance $\|Q_i - \hat{Q}_i\|_F$: by the orthonormality of the $\{B^{(i)}_\alpha\}$, we get

$$\|Q_i - \tilde{Q}_i\|_F^2 = \sum_\alpha \|A^{(i)}_\alpha - \hat{A}^{(i)}\|_F^2 \le d^2\eta^2(\delta) \ . \tag{12}$$

So $\|Q_i - \tilde{Q}_i\|_F \le d\eta(\delta)$, and since $Q_i = Q_i^\dagger$, it follows that also $\|Q_i - \hat{Q}_i\|_F \le d\eta(\delta)$. The proof now follows by setting $\epsilon(\delta) \stackrel{\text{def}}{=} d\eta(\delta)$, and noticing that $\lim_{\delta\to 0}\epsilon(\delta) = \lim_{\delta\to 0}\eta(\delta) = 0$. □

We use the lemma to construct a new 2-local Hamiltonian by sequentially going over all qudits and replacing $Q_i \mapsto \hat{Q}_i$. We obtain a new system $\hat{H} = \sum_i \hat{Q}_i$, which is commuting and the distance between two corresponding operators Q_i and $\hat{Q}_i$ is $\|Q_i - \hat{Q}_i\|_F \le 2\epsilon(\delta)$ (This is because Q_i works on 2 qudits, and so it undergoes two replacements).

To finish the proof of Theorem 4, we use the fact that both the Hilbert space of a single qudit, or the Hilbert space on which the Q_i operators work are of constant dimension (of d and d^2 respectively). Therefore in these spaces all norms are equivalent up to some d-dependent factor. It follows that we can re-define the function $\epsilon(\delta)$ to apply to the *operator* norm, as stated in the theorem. In other words, we re-scale $\epsilon(\delta)$ such that if $\|[Q_i, Q_j]\| \le \delta$, then there exists a commuting system $\{\hat{Q}_i\}$ such that

$$\|Q_i - \hat{Q}_i\| \le \epsilon(\delta)/2 \ . \tag{13}$$

Let $|\psi_0\rangle$ be the ground state of the original system. By assumption, $\langle\psi_0|H|\psi_0\rangle = 0$. Then by Eq. (13), $\langle\psi_0|\hat{H}|\psi_0\rangle \le M\epsilon(\delta)/2$, hence the ground energy of $\hat{H}$ is upper-bounded by $M\epsilon(\delta)/2$. By Ref. [6], $\hat{H}$ has a ground state $|\psi_0'\rangle$ that can be written as an efficient tensor-network. By using Eq. (13) once more, we see that $\langle\psi_0'|H|\psi_0'\rangle \le \epsilon(\delta)M$, as required. □

Acknowledgments

I wish to thank Dorit Aharonov, Matt Hastings, Lior Eldar, , Zeph Landau, Tobias Osborne, Umesh Vazirani, for inspiring discussions about the quantum PCP and to Dorit Aharonov for her great help on the manuscript.

I gratefully acknowledge support by Julia Kempe's ERC Starting Researcher Grant QUCO and ISF 759/07.

References

1. S. Bravyi, "Efficient algorithm for a quantum analogue of 2-SAT," *arXiv:quant-ph/0602108*, 2006.
2. D. Aharonov, I. Arad, Z. Landau, and U. Vazirani, "The detectability lemma and quantum gap amplification," in *STOC '09: Proceedings of the 41st annual ACM symposium on Theory of computing, arXiv:0811.3412*, (New York, NY, USA), pp. 417–426, ACM, 2009.
3. I. Dinur, "The pcp theorem by gap amplification," *J. ACM*, vol. 54, no. 3, p. 12, 2007.
4. M. B. Hastings. personal communication, 2008.
5. D. Poulin and M. B. Hastings, "Markov entropy decomposition: a variational dual for quantum belief propagation," *arXiv:1012.2050*, 2010.
6. S. Bravyi and M. Vyalyi, "Commutative version of the k-local Hamiltonian problem and common eigenspace problem," *arXiv:quant-ph/0308021*, 2003.
7. D. Aharonov, D. Gottesman, S. Irani, and J. u. Kempe, "The power of quantum systems on a line," *Communications in Mathematical Physics*, vol. 287, pp. 41–65, 2009. 10.1007/s00220-008-0710-3.
8. C. Pearcy and A. Shields, "Almost commuting matrices," *Journal of Functional Analysis*, vol. 33, no. 3, pp. 332 – 338, 1979.
9. P. Halmos, "Some unsolved problems of unknown depth about operators on Hilbert space," in *Proc. R. Soc. Edinb., Sect. A*, vol. 76, pp. 67–76, 1976.
10. M. B. Hastings. personal communication, 2009.