

A multiprover interactive proof system for the local Hamiltonian problem

Joseph Fitzsimons*

Thomas Vidick[†]

Abstract

We give a quantum interactive proof system for the local Hamiltonian problem on n qubits in which (i) the verifier has a single round of interaction with five entangled provers, (ii) the verifier sends a classical message on $O(\log n)$ bits to each prover, who reply with a constant number of qubits, and (iii) completeness and soundness are separated by an inverse polynomial in n . As the same class of proof systems, without entanglement between the provers, is included in QCMA, our result provides the first indication that quantum multiprover interactive proof systems with entangled provers may be strictly more powerful than unentangled-prover interactive proof systems. A distinguishing feature of our protocol is that the completeness property requires honest provers to share a large entangled state, obtained as the encoding of the ground state of the local Hamiltonian via an error-correcting code. Our result can be interpreted as a first step towards a multiprover variant of the quantum PCP conjecture.

1 Introduction

The PCP theorem [AS98, ALM⁺98] asserts that any language in NP admits proofs of membership that can be efficiently verified using a randomized procedure which makes the correct decision with high probability while only ever reading a constant number of bits of the proof. An equivalent formulation of the PCP theorem, that has been particularly useful in applications to hardness of approximation [Hås01] as well as in devising further improvements to the theorem [Raz98], uses the language of multiplayer games. A two-player game G is specified by question sets Q, Q' , answer sets A, A' , a distribution π on $Q \times Q'$ and a verification criterion $V \subseteq (A \times A') \times (Q \times Q')$. The value ω of G is defined as the maximum, over all assignments $f : Q \rightarrow A, f' : Q' \rightarrow A'$, of the average number of valid answers given by the assignments: $\omega(G) = \sup_{f, f'} \sum_{q, q'} \pi(q, q') V(f(q), f'(q'); q, q')$. The PCP theorem is equivalent to the statement that $\omega(G)$ is NP-hard to approximate to within a constant additive factor, even for the case of answer sets A, A' of constant size. To see the connection, consider the following “consistency game”: the verifier, instead of directly reading bits i_1, \dots, i_k of the proof, asks a first player for the entries at those locations and a second player for the entry corresponding to a single location i_j , where j is chosen uniformly at random in $\{1, \dots, k\}$. The verifier accepts if and only if the first player’s answers correspond to entries that he would have accepted had he read them directly from the proof, *and* the second player’s answer is consistent with that of the first. It is not hard to see that the value of the consistency game is directly related to the fraction of checks satisfied by the optimal PCP proof, so that the respective complexities of deciding whether either is close to 1 (under the appropriate gap promise) are identical.

*Singapore University of Technology and Design and Centre for Quantum Technologies, National University of Singapore, Singapore. Email: joe.fitzsimons@nus.edu.sg.

[†]California Institute of Technology, Pasadena, CA, USA. Email: vidick@cms.caltech.edu.

The quantum analogue of the local proof checking problem was introduced by Kitaev [KSV02]. An instance of the k -local Hamiltonian problem (LH) is specified by m local Hamiltonians H_1, \dots, H_m , where each H_i is a Hermitian matrix of norm at most 1 acting on at most k out of a total of n qubits. The instance is positive if there exists a quantum proof (a quantum state $|\Psi\rangle$ on the n qubits) satisfying a fraction at least $(1 - a)$ of the constraints; precisely, if $H = \sum_i H_i$ (where each H_i is implicitly tensored with the identity on the remaining qubits) has an eigenvalue at most am . If all eigenvalues of H are larger than bm for some $b > a$ the instance is negative. The introduction of the local Hamiltonian problem initiated what is now the burgeoning field of Hamiltonian complexity [Osb12, GH14], expanding well beyond the initial formal connection with classical constraint satisfaction problems to encompass the computational study of a range of problems motivated by condensed-matter physics.

Kitaev proved the “quantum Cook-Levin theorem”: he introduced the class QMA of languages that admit efficiently verifiable quantum proofs, and showed that the local Hamiltonian problem is QMA-complete for some a, b satisfying $b - a = \Theta(\text{poly}^{-1}(n))$. The natural question of whether a quantum analogue of the PCP theorem holds was first posed in [AN02]; it asks whether the local Hamiltonian problem remains QMA-hard for values $b - a = \Omega(1)$. This problem has captured the imagination of many researchers [Aar06, Has13, FH13], but very little is known. If anything recent results [BH13, AAV13] place strong limitations on the parameters, including the locality k or the degree of the constraint graph, for which the conjecture may be valid, showing that it may only hold for ranges of parameters that appear to be much more limited than those for which the classical PCP theorem is known to be true.

In this paper we shed new light on the complexity of the local Hamiltonian problem by recasting it in the language of quantum interactive proofs with entangled provers. In doing so we are motivated by the existing classical connection between local proof verification and multiplayer games, which as already mentioned has been instrumental both in the development of the PCP theorem (and in particular its second proof by Dinur [Din07]) and for applications. Does this connection extend to the quantum setting? While quantum multiprover interactive proof systems have been intensely studied for their own sake [KM03, KKMV09, IV12], prior to our work no nontrivial relation was known between the class QMA_{EXP} , the exponentially scaled-up version of QMA, and the classes QMIP^* or QMIP of languages having quantum interactive proof systems with entangled or unentangled provers respectively. In fact, the latter is known to equal NEXP [KM03], while the former was only recently shown to *contain* NEXP [IV12]. However, no upper bound on QMIP^* is known, so that one may ask — could QMIP^* be a *larger* class than $\text{QMIP} = \text{NEXP}$? The only distinction between the two classes is the presence of entanglement between the provers, which until now (and with some rare exceptions [KKMV09]) has for the most part been understood as a nefarious resource that could be used by the provers in order to break a protocol’s soundness. Giving a positive answer to the question, however, requires finding a *beneficial* use of entanglement, as it entails devising a protocol in which even honest provers are *required* to share an entangled state over a superpolynomial number of qubits in order to succeed on positive instances.¹

A natural target for going beyond $\text{NEXP} \subseteq \text{QMIP}^*$ consists in devising protocols establishing the inclusion of QMA_{EXP} in QMIP^* . Proving such inclusion, however, immediately runs into a number of serious difficulties. To see why, consider the following attempt at designing a quantum interactive proof system for the local Hamiltonian problem that mimics the classical construction of the consistency game (which, as described earlier, easily leads to a proof of $\text{NEXP} \subseteq \text{MIP}$ assuming the PCP theorem). Suppose that the first player is asked to provide a constant-sized subset of the proof qubits, corresponding to a local constraint H_j which the verifier can then check. In the classical case, the second player is asked for just one

¹The class $\text{QMIP}^{(l.e.)}$ of languages having quantum multiprover interactive proof systems in which the provers share an entangled state on at most a polynomial number of qubits is also known to be included in NEXP [KM03].

of the bits asked to the first player; this is used to verify that the first players' answers to any of the bits he was asked about depends on that bit only, and not on the subset of which it is part. In the quantum case this approach is all but ruled out by the no-cloning principle: any given proof qubit can be placed in the hands of one player only, but it cannot be duplicated! Hence the direct quantum analogue of the consistency game does not have *completeness*: even satisfiable instances of the local Hamiltonian problem may not lead to a winning strategy for the players.

Natural workarounds to this difficulty run into different obstacles. For instance, consider splitting the proof (e.g. the ground state of the local Hamiltonian instance) qubits into two (or more) sets S_1 and S_2 , and only asking prover i for qubits coming from set S_i . While this leads to a game which does have perfect completeness, the fact that the sets need to be specified a priori can, at least in some cases, prevent the *soundness* property from holding. To see why, consider the simple example of a one-dimensional nearest-neighbor Hamiltonian in which each term is a projection on the orthogonal complement of an EPR pair split across two adjacent qubits. This Hamiltonian is highly frustrated, as any qubit can only form an EPR pair with its left *or* right neighbor, not both. Nevertheless, the corresponding game in which S_1 (resp. S_2) is the set of all even-numbered (resp. odd-numbered) qubits has a perfect strategy: the players share a single EPR pair and systematically send back their respective half, independently of the question they are asked! Although in this particular case the issue is easily fixed by choosing a different splitting of the proof qubits, in general it seems like any such splitting will be arbitrary and could be taken advantage of by the provers.

1.1 Results

Our main result is the design of an interactive proof system for the local Hamiltonian problem which circumvents the aforementioned difficulties. This is the first time a multiprover interactive proof system is given for a QMA-complete, instead of NP-complete, problem, and it provides strong indication that entangled proof systems may be strictly more powerful than their unentangled counterparts. Formally, we show the following.

Theorem 1. *Let k be an integer. There exists constants $C, c > 0$ depending on k only such that the following holds. Let $H = \sum_{i=1}^m H_i$ be an instance of the k -local Hamiltonian problem on n qubits, such that the number of constraints is $m = \text{poly}(n)$. There exists a one-round interactive protocol between a quantum polynomial-time verifier and $r = 5$ entangled quantum provers such that:*

- *The verifier sends $O(\log n)$ -bit classical messages to each prover,*
- *The provers respond with at most k qubits each,*
- *If there exists a state $|\Gamma\rangle$ such that $\langle \Gamma | H | \Gamma \rangle \leq am$ then there is a strategy for the provers that is accepted with probability at least $1 - a/2$,*
- *If for every state $|\Psi\rangle$, $\langle \Psi | H | \Psi \rangle \geq bm$ then any strategy of the provers is accepted with probability at most $1 - Cb/n^c$.*

The local Hamiltonian problem is known to be QMA-complete for $k = 2$, a that is exponentially small and b at least an inverse polynomial [KKR06]. The following corollary, which we state using the language of multiplayer games, is thus a direct consequence of Theorem 1:

Corollary 2. *The problem of approximating, to within an additive inverse polynomial, the referee's maximum acceptance probability in a quantum multiplayer game in which questions from the referee are classical*

on $O(\log n)$ -bits and answers from the players are quantum on $O(1)$ qubits is QMA-hard. Furthermore the same holds when restricted to games in which there is a single round of interaction between the referee and at most 5 players.

The same problem but with no entanglement between the players is contained in QCMA: the players' constant-sized quantum answers can be given as a classical description [KM03]. It is also known to be NP-hard, even when restricted to classical answers from the players and for constant additive approximations [Vid13]. However, no *upper bound* is known on the complexity of the problem considered in Corollary 2, which is not even known to be decidable [SW08, JNP⁺11] (and there is no known a priori bound on the amount of entanglement that may be beneficial to the players). Corollary 2 provides the first indication that entanglement indeed *increases* the verifying power of the referee, at least in the range of inverse-polynomial approximations, showing that unless QCMA = QMA the complexity of entangled (quantum) games is strictly *larger* than that of non-entangled (quantum) games.

Consequences for interactive proof systems with entangled provers. We can scale up our result to QMA_{EXP}, the exponential-witness size version of QMA (see Section 2 for the definition) to obtain a formal separation between quantum multiprover interactive proof systems with and without entanglement between the provers. Let QMIP^{*}(r, t, c, s) be the class of languages that have quantum interactive-proof systems with r provers, t rounds of interaction, completeness c and soundness s (see Section 2 for the complete definition).

Corollary 3. *There exists a polynomial q such that*

$$\text{QMA}_{\text{EXP}} \subseteq \text{QMIP}^*(5, 1, 1 - 2^{-(q+1)}, 1 - 2^{-q})$$

and hence

$$\text{QMIP}(5, 1, 1 - 2^{-(q+1)}, 1 - 2^{-q}) \subsetneq \text{QMIP}^*(5, 1, 1 - 2^{-(q+1)}, 1 - 2^{-q})$$

unless NEXP = QMA_{EXP}.

The corollary follows from the fact that $\text{QMIP}(5, 1, 1 - 2^{-(q+1)}, 1 - 2^{-q}) \subseteq \text{NEXP}$ [KM03] and $\text{NEXP} \subseteq \text{MIP}^*(3, 1, 1, 1 - 1/\text{poly})$ [IV12] together with the observation $\text{MIP}^*(3, 1, 1, 1 - 1/\text{poly}) \subseteq \text{QMIP}^*(5, 1, 1 - 2^{-(q+1)}, 1 - 2^{-q})$.

We note that even though it is known that $\text{MIP}^* = \text{QMIP}^*$ [RUV13] the above corollary falls short of proving a separation between $\text{MIP} = \text{NEXP}$ and MIP^* . The reason is that the transformation from a QMIP^{*} to a MIP^{*} protocol in [RUV13] requires the completeness and soundness parameters of the QMIP^{*} protocol to be separated by an inverse polynomial in the input size, whereas our construction only gives an inverse exponential separation.

1.2 Proof idea

Suppose given an instance $H = \sum H_j$ of the local Hamiltonian problem, where each term H_j acts on a subset $S_j = \{i_1, \dots, i_k\}$ of at most k out of the n qubits. Given an explicit description of H , the goal of the verifier is to decide whether there exists a “proof” $|\Psi\rangle$ that satisfies most terms H_j , i.e. such that the total “energy” $\langle \Psi | H | \Psi \rangle$ is below a certain threshold value. As already mentioned, the main challenge in achieving this is that the verifier will only ever receive, at best, a logarithmic number of qubits of the proof from the provers. Although this easily allows him to estimate the energy $\langle \Psi | H_j | \Psi \rangle$ of any local term H_j , the difficulty is to ensure that the qubits received in response to different queries, associated with different local

terms H_j , are *globally consistent* — that they can be “patched together” into an actual proof $|\Psi\rangle$ that has low energy with respect to H . This difficulty is unique to the case of quantum proofs: if we were working with classical assignments, as explained earlier a simple consistency check would be sufficient to enforce that the provers’ answers can be combined into a single assignment satisfying most clauses. But how does one devise a consistency check for quantum proofs, when in general it is not even possible to check whether two quantum states agree locally?²

We suggest the following workaround. Our main goal is to ensure that, when a prover is asked for its share of a certain qubit i_ℓ , or $i_{\ell'}$, of the proof, the actual qubits that it sends back to the verifier in each case do indeed correspond to distinct physical qubits — that they do not “overlap”, or even correspond to the same physical qubit, as was the case in our description of a strategy for the frustrated Hamiltonian projecting on overlapping EPR pairs. To enforce this, instead of asking the (honest) provers to directly split the qubits of the original proof between themselves we ask them to share an *encoding* of the proof: each “logical” qubit of $|\Psi\rangle$ should be individually encoded into five “physical” qubits using a quantum error-correcting code. Each of five provers should then be given one of the five shares associated with each of the original proof’s qubits. (Five is the smallest number of qubits for a quantum error-correcting code satisfying the properties we need; although we did not investigate this further a four-qubit error-detecting code may also suffice.)

Given this (presumed) splitting of the proof, we introduce the following protocol, comprised of two tests each applied with probability $1/2$ by the verifier. Observe first that under our encoding it remains easy for the verifier to estimate the energy of any k -local term H_j : he can ask each of the five provers for its corresponding share of each qubit on which a randomly chosen H_j acts, decode the results, and measure the energy of the resulting qubits with respect to H_j . This only requires each prover to send back k qubits to the verifier, and constitutes the first test in our protocol.

Next consider the following additional test. The verifier chooses a k -element subset $S = \{i_1, \dots, i_k\}$ of $\{1, \dots, n\}$ uniformly at random. He also selects an index $\ell \in \{1, \dots, k\}$ at random and asks four out of the five provers (again chosen at random) for their respective share of qubit i_ℓ only. To the last prover he asks for its respective shares of all qubits in S . (Note that in case S corresponded to the set of k qubits on which a local term H_j acts the last prover cannot distinguish whether it is this test or the first that is being performed, and this will be important for the proof.) The verifier checks that all shares that he received associated with qubit i_ℓ lie in the codespace, and rejects the provers if not.

In this second test the messages sent back by the first four provers only depend on qubit i_ℓ . The key point is that, informally, given their four respective answer qubits there can exist at most one additional qubit that is entangled with them in a way that completes a valid codeword. Indeed (and again informally), if there were two such qubits it would imply that it is possible for the “environment” to entangle itself with a codeword through acting on a single qubit and without being detected by the code — this possibility is excluded as long as the code is required to correct (or even just detect) all single-qubit errors. Thus this additional test enforces that the qubit sent back by the fifth prover in response to query i_ℓ is uniquely specified by the query i_ℓ ; this is achieved by “locking” the qubit with the other four provers’ answers via the codespace.

Although the above provides some intuition, proving soundness of the protocol remains technically challenging. We need to show how a complete proof $|\Psi\rangle$ serving as a witness for the energy of the Hamiltonian H can (at least in principle) be reconstructed from prover strategies that are successful in the protocol. Formally each prover’s strategy is specified by a pair of unitaries, one for each type of query from the verifier.

²Pure quantum states $|\Psi\rangle$ and $|\Phi\rangle$ can be compared using the so-called SWAP test. However, for mixed states this test no longer works, and in fact checking consistency of reduced density matrices, even when specified explicitly, is itself a QMA-complete problem [Liu06]. We refer to [AAV13] for more on the difficulties posed by locally checking consistency of quantum states.

The difficulty in proving that these unitaries are “compatible” and can be composed so as to reconstruct $|\Psi\rangle$ from the provers’ entangled registers — indeed, note a prover may apply an arbitrary transformation to its private space before answering any of the verifier’s queries. Our proof specifies an explicit circuit, based on the provers’ unitaries, for reconstructing $|\Psi\rangle$ from their initial entangled state. The depth of this circuit is linear in the number of qubits, and it is ultimately this which leads to the polynomial dependence of the soundness parameter on the number of qubits in the proof.

1.3 Open questions

Our work gives the first indication that multi-prover interactive proof systems with entangled provers may be strictly more powerful than their purely classical counterparts. Our protocol relies on the use of quantum communication from the provers to the verifier. Although it is known that quantum communication does not increase the power of entangled-prover interactive proof systems, $\text{QMIP}^* = \text{MIP}^*$ [RUV13], the technique used in [RUV13] to replace quantum messages by classical ones introduce a polynomial amount of error that, at least if applied naïvely, would close the completeness/soundness gap of our protocol. We thus leave the possibility of achieving the same results as our ours through a purely classical interaction as an interesting open question.

The main drawback of our protocol is the scaling of the completeness/soundness gap with the size of the local Hamiltonian instance. The most important question that we leave open for future work is to increase this gap from inverse exponential to inverse polynomial, leading to the inclusion $\text{QMA}_{\text{EXP}} \subseteq \text{QMIP}^*$. Together with $\text{QMIP}^* = \text{MIP}^*$ [RUV13] such a result would in particular prove the main result of [IV12], and we expect it to pose a significant challenge. Of importance in itself, research on this question could lead to the development of techniques useful to the study of the quantum PCP conjecture [AAV13]. To stimulate its exploration we propose that the inclusion $\text{QMA}_{\text{EXP}} \subseteq \text{QMIP}^*$ be taken as a second variant of “quantum PCP conjecture” — one we could call the “interactive-proof QPCP”, in contrast to the “proof-checking QPCP” that has so far been the accepted formulation (see e.g. Conjecture 1.4 in [AAV13]). No implication is known between the two conjectures; our work provides a first step towards the former, making it potentially more approachable than the latter.

Acknowledgments. This work was started while both authors were hosted by the Simons Institute in Berkeley, whose financial support we gratefully acknowledge. The second author is grateful to Dorit Aharonov and Umesh Vazirani for pressing him to expose the question investigated in this paper during an open problems session organized at the institute. Joseph Fitzimons’ research is supported by the Singapore National Research Foundation under NRF Award No. NRF-NRFF2013-01. Thomas Vidick’s research was supported in part by the Simons Institute and the Ministry of Education, Singapore under the Tier 3 grant MOE2012-T3-1-009.

2 Preliminaries

Notation. Given a string x we let $|x|$ denote its length. For a set S , $|S|$ is its cardinality. For a positive integer n we abbreviate $\{1, \dots, n\}$ by $[n]$. We use a calligraphic \mathcal{H} to denote finite-dimensional Hilbert spaces, and roman letters Q, R, \dots to denote quantum registers. The Hilbert space associated with register R is \mathcal{H}_R . We will often, though not always, index kets and bras for quantum states by the names of the registers on which the state lies, e.g. $|\Psi\rangle_{QR}$ means that $|\Psi\rangle$ is a bipartite state on $\mathcal{H}_Q \otimes \mathcal{H}_R$. $L(\mathcal{H}_Q, \mathcal{H}_R)$ is the set of all linear maps $\mathcal{H}_Q \rightarrow \mathcal{H}_R$. $\text{Pos}(\mathcal{H})$ is the set of positive operators on \mathcal{H} ; $\text{D}(\mathcal{H})$ the set of density

matrices. Given $\mathcal{F}, \mathcal{G} \in L(\mathcal{H}, \mathcal{H})$ we let $\mathcal{F} \circ \mathcal{G}$ denote their composition. If there are s such maps \mathcal{F}_ℓ , we let $\bigcirc_{\ell=1}^s \mathcal{F}_\ell := \mathcal{F}_s \circ \dots \circ \mathcal{F}_1$.

Given two registers Q and R associated to isomorphic Hilbert spaces $\mathcal{H}_Q, \mathcal{H}_R$ respectively we let SWAP_{QR} be the unitary that swaps their contents: for any two orthonormal bases $|u_i\rangle$ for \mathcal{H}_Q and $|v_j\rangle$ for \mathcal{H}_R , $\text{SWAP}_{QR} = \sum_{i,j} |v_i, u_j\rangle \langle u_j, v_i|$.

Complexity classes. We give relatively informal definitions of the quantum interactive proof classes considered in this paper. For formal definitions we refer the reader to the book [KSV02] and the survey [Wat09].

QMA is the class of all promise problems $L = (L_{yes}, L_{no})$ such that there exists a polynomial p and a quantum polynomial-time verifier V such that:

- (completeness) For every $x \in L_{yes}$, there exists a state $|\Psi\rangle$ on $p(|x|)$ qubits such that $V(x, |\Psi\rangle)$ accepts with probability at least $2/3$,
- (soundness) For every $x \in L_{no}$ and every $|\Psi\rangle$ on $p(|x|)$ qubits, $V(x, |\Psi\rangle)$ accepts with probability at most $1/3$.

We further note that using an amplification technique of Marriott and Watrous [MW05] one can show that for any fixed polynomial q the completeness and soundness parameters can be replaced by $1 - 2^{-q(|x|)}$ and $2^{-q(|x|)}$ respectively without changing the definition of QMA. Furthermore the amplification procedure in [MW05] preserves the witness length, so that the polynomial p does not need to grow if one increases q (only the complexity of the verification procedure increases). We define the exponential-size version of QMA, QMA_{EXP} , by allowing the witness to be on $2^{p(|x|)}$ qubits and the verifier to run in quantum exponential time.

$\text{MIP}(r, t, c, s)$ is the class of all promise problems $L = (L_{yes}, L_{no})$ such that there exists a polynomial p and a classical polynomial-time verifier V , interacting with r non-communicating provers through t rounds of interaction in each of which at most $p(|x|)$ bits of communication are exchanged between the verifier and the provers, such that:

- (completeness) For every $x \in L_{yes}$, there exists a strategy for the provers that is accepted by the verifier with probability at least c ,
- (soundness) For every $x \in L_{no}$ any strategy of the provers is accepted by the verifier with probability at most s .

$\text{QMIP}(r, t, c, s)$ is defined in the same way, except the verifier and communication exchanged are allowed to be quantum. $\text{MIP}^*(r, t, c, s)$ (resp. $\text{QMIP}^*(r, t, c, s)$) is defined as $\text{MIP}(r, t, c, s)$ (resp. $\text{QMIP}(r, t, c, s)$) but the provers are allowed to share an arbitrary entangled state as part of their strategy. (In this paper we only consider protocols for which the number of rounds of interaction is $t = 1$.)

It follows from [BFL91, KM03] that, for any polynomials p_1, p_2 and p_3 ,

$$\begin{aligned} \text{MIP}(p_1, p_2, 2/3, 1/3) &= \text{QMIP}(p_1, p_2, 2/3, 1/3) \\ &= \text{MIP}(2, 1, 1, 2^{-p_3}) = \text{QMIP}(2, 1, 1, 2^{-p_3}) = \text{NEXP}. \end{aligned}$$

In fact, [KM03] even show that the same equalities hold for QMIP^* when the provers are limited to a polynomial number of qubits of entanglement.

The local Hamiltonian problem. Let k be a fixed integer and $a, b : \mathbb{N} \rightarrow [0, 1]$ such that $a(n) < b(n)$ for all integers n . The k -local Hamiltonian problem (LH) is defined as follows. The input is a classical description of a local Hamiltonian $H = \sum_{i=1}^m H_i \in L(\mathbb{C}^{d^n}, \mathbb{C}^{d^n})$ acting on n qudits of dimension d each. Here each H_i is a positive semidefinite matrix of norm at most 1 acting on at most k out of the n qudits, and can thus be represented by a matrix of dimension $d^k \times d^k$; when we write $H = \sum_i H_i$ we implicitly mean that each H_i should be tensored with the identity acting on the remaining $(n - k)$ qudits. We label the qudits from 1 to n , and denote by S_j the set of k qudits on which H_j acts. The problem is to determine which of the following two cases holds:

1. (YES) There exists a n -qudit state $|\Gamma\rangle$ such that $\langle \Gamma | H | \Gamma \rangle \leq am$,
2. (NO) For all states $|\Psi\rangle$, $\langle \Psi | H | \Psi \rangle \geq bm$.

Kempe, Kitaev and Regev showed the following:

Theorem 4 ([KKR06]). *For any fixed polynomial q , there is a polynomial p such that the k -local Hamiltonian problem, where the number of qubits n is specified in unary, is QMA-complete for $k = 2$, $d = 2$, $a = 2^{-q(n)}$ and $b = 1/p(n)$.*

For the case of QMA_{EXP} essentially the same construction yields the following (see also [GI09]):

Theorem 5 ([KKR06]). *For any fixed polynomial q , there is a polynomial p such that the k -local Hamiltonian problem, where the number of qubits N is specified in binary (hence can be exponential in the input size), is QMA_{EXP} -complete for $k = 2$, $d = 2$, $a = 2^{-q(N)}$ and $b = 1/p(N)$.*

Error-correcting codes. Our protocol relies on the use of a quantum error-correcting code C that has the following properties:

- C encodes 1 logical qubit into r physical qubits.
- C detects and corrects all single-qubit Pauli errors on a single qubit.
- The reduced density matrix of any codewords in C on a single qubit is the totally mixed state $\text{Id} / 2$.

An example of a code satisfying all three conditions for $r = 5$ and $e = 1$ is the 5-qubit stabilizer code [BDSW96, LMPZ96]. Given r single-qubit registers R_1, \dots, R_r we let $\text{DEC}_{R_1 \dots R_r} : D((\mathbb{C}^2)^{\otimes r}) \rightarrow D(\mathbb{C}^2)$ be the completely positive trace-preserving (CPTP) map corresponding to the decoding operation. We also let $\text{CHECK}_{R_1 \dots R_r} \in \text{Pos}((\mathbb{C}^2)^{\otimes r})$ be the projection onto the code space.

3 Proof of Theorem 1

In this section we prove Theorem 1. The protocol is described in Figure 1. The first two properties claimed in the theorem, on the structure of the protocol, are clear: there is a single round of interaction, and using the 5-qubit stabilizer code for C the protocol can be executed with $r = 5$ provers. Messages from the verifier to the provers are either the label of a qubit or the description of a set of size k , which require $O(\log n)$ bits to specify. Messages from any prover to the verifier are either 1 or k qubits. In Section 3.1 we establish the completeness property of the protocol; soundness is proved in Section 3.2.

Protocol P

Let $H = \sum_{i=1}^m H_i$ be an instance of the k -local Hamiltonian problem given as input, and n the number of qubits on which H acts. Let C be an error-correcting code which encodes 1 logical qubit into r physical qubits and satisfies the three conditions described at the end of Section 2.

The verifier performs each of the following tests with probability $1/2$ each:

- Test (a) Select a $j \in [m]$ uniformly at random, and let $S_j \subseteq [n]$ be the set of k qubits on which the local term H_j acts. Ask the provers for their respective share of all qubits in S_j . Upon receiving the shares, apply the decoding map independently to each of the k groups of r shares and measure the resulting state using $\{H_j, \text{Id} - H_j\}$. Reject if the outcome is ' H_j '.
 - Test (b) Select a qubit $i \in [n]$ uniformly at random, and a set $S \subseteq [n]$ uniformly at random among all sets of size k that contain i . With probability $1/2$, ask one of the provers at random for his share of all qubits in S , and the remaining $r - 1$ provers for their respective share of the i -th qubit only. With probability $1/2$, ask all provers for their respective share of the i -th qubit. In both cases, verify that all provers' shares of the i -th qubit together lie in the codespace. Reject if not.
-

Figure 1: Protocol for the verification of an instance of the local Hamiltonian problem.

3.1 Completeness analysis

Lemma 6. *Suppose that there exists a state $|\Gamma\rangle$ such that $\langle \Gamma | H | \Gamma \rangle \leq am$. Then there exists a strategy for the provers in Protocol P that is accepted with probability at least $1 - a/2$.*

Proof. We describe a strategy for the provers. Let $|\Gamma\rangle$ be such that $\langle \Gamma | H | \Gamma \rangle \leq am$. Before the protocols start, the provers generate a shared entangled state $|\Psi\rangle$ over rn qubits by independently encoding each qubit of $|\Gamma\rangle$ into r qubits using the code C prescribed by the protocol. Each of the r provers keeps n qubits of $|\Psi\rangle$, corresponding to a share of each of the encoded qubits of $|\Gamma\rangle$. When asked for its share of any set of qubits, the prover complies and sends it to the verifier. It is clear that this strategy is accepted with probability 1 in item (b), and with probability

$$\frac{1}{m} \sum_{i=1}^m \langle \Gamma | (\text{Id} - H_i) | \Gamma \rangle \geq 1 - a$$

in item (a). Using that each test is performed with probability $1/2$, the overall success probability for the strategy is at least $1 - a/2$. \square

3.2 Soundness analysis

In this section we analyze the soundness of protocol P . In section 3.2.1 we introduce the notation used to describe the most general strategy that the provers may employ in the protocol. In section 3.2.2 we show that, provided that all eigenvalues of H are larger than some inverse polynomial, any strategy for the provers is rejected by the verifier with inverse polynomial probability.

	Register	Use
Before application of U_i, V_S .	P^t	Prover t 's register in state $ \Psi\rangle$
After application of U_i, V_S	Q_i^t S^t	Register sent back by prover t if asked for the i -th qubit. Prover t 's remaining private registers.
Auxiliary registers	R_i^t, \bar{R}_i^t	Initialized as an EPR pair.

Figure 2: Notation for the provers' registers.

3.2.1 The provers' strategies

We denote an arbitrary strategy for the r provers in protocol P via a triplet $(U_i^j, V_S^j, |\Psi\rangle)$ (or sometimes (U_i^j, V_S^j, ρ)). Here $|\Psi\rangle$ (or ρ) denotes the initial r -partite entangled state shared by the provers, and U_i, V_S the unitaries that they apply upon receiving questions i, S respectively. More precisely, in the protocol a prover is asked two types of questions. Either it is asked for a single qubit i , in which case we call the unitary U_i^t (where t indexes the prover), or it is asked for a set of k qubits S , in which case we call the unitary V_S^t . We sometimes omit the superscript t , as the labeling of the provers will often be clear from context. We denote the associated completely positive trace-preserving (CPTP) maps by $\mathcal{U}_i^t : \sigma \mapsto U_i^t \sigma (U_i^t)^\dagger$ and $\mathcal{V}_S^t : \sigma \mapsto V_S^t \sigma (V_S^t)^\dagger$.

For $t \in [r]$ we write P^t for the register containing the t -th prover's share of $|\Psi\rangle$. After application of the unitary U_i^t or V_S^t , we relabel registers associated to the prover as S^t, Q_1^t, \dots, Q_n^t . Here the n registers Q_1^t, \dots, Q_n^t are each single-qubit registers such that register Q_i^t (resp. registers $Q_{i_1}^t \dots Q_{i_k}^t$) is sent back to the verifier when the prover is asked for qubit i (resp. set of qubits $S = \{i_1, \dots, i_k\}$). Note that all registers Q_i^t may not exist simultaneously; which ones do depends on the unitary U_i^t or V_S^t that was applied. The remaining register S^t is an auxiliary register of arbitrary dimension. In addition, for each prover $t \in \{1, \dots, r\}$ we introduce $2n$ auxiliary registers R_1^t, \dots, R_n^t and $\bar{R}_1^t, \dots, \bar{R}_n^t$, and define

$$|\tilde{\Psi}\rangle := |\Psi\rangle \bigotimes_{t=1}^r \bigotimes_{i=1}^n \frac{1}{\sqrt{2}} (|00\rangle_{R_i^t \bar{R}_i^t} + |11\rangle_{R_i^t \bar{R}_i^t}), \quad (1)$$

i.e. $|\tilde{\Psi}\rangle$ is $|\Psi\rangle$ adjoined with n EPR pairs for each prover, created in the auxiliary registers. We write $\rho = |\Psi\rangle\langle\Psi|$ and $\tilde{\rho} = |\tilde{\Psi}\rangle\langle\tilde{\Psi}|$. See Figure 2 for a summary of our nomenclature for registers. We will often abbreviate Q_i for the union of the $Q_i^j, j \in [r]$, and write $Q_i^{\neq t}$ for the union of all Q_i^j for $j \in [r] \setminus \{t\}$.

We introduce a new set of unitaries which act on a prover's share of $|\tilde{\Psi}\rangle$ as

$$C_i^t := (U_i^t)^\dagger (\text{SWAP}_{Q_i^t R_i^t} \otimes \text{Id}) U_i^t \quad \text{and} \quad D_{i,S}^t := (V_S^t)^\dagger (\text{SWAP}_{Q_i^t R_i^t} \otimes \text{Id}) V_S^t, \quad (2)$$

where U_i^t and V_S^t are implicitly tensored with the identity on the auxiliary registers. We denote the associated CPTP maps by $\mathcal{C}_i^t : \sigma \mapsto C_i^t \sigma (C_i^t)^\dagger$ and $\mathcal{D}_{i,S}^t : \sigma \mapsto D_{i,S}^t \sigma (D_{i,S}^t)^\dagger$. In order not to overload the notation we often do not specify precisely on which registers the identity acts (sometimes we even omit the symbol Id altogether), as it should always be clear from context. In words, C_i^t corresponds to applying U_i^t , swapping the register Q_i^t containing the output qubit with the i -th ancilla register R_i^t , and applying $(U_i^t)^\dagger$. For $i \in S$, $D_{i,S}^t$ is defined as C_i^t but from the unitary V_S^t instead of U_i^t , while still swapping the output qubit in register Q_i^t only (and not the others). For any subset $T \subseteq S$ we define $D_{T,S}^t$ in the same ways as $D_{i,S}^t$ except all qubits

in the subset T are swapped out; in particular $D_{\{i\},S}^t = D_{i,S}^t$ and $D_{\emptyset,S} = \text{Id}$. The following observation, which follows from $V_S^t(V_S^t)^\dagger = \text{Id}$, will be useful:

$$\forall T \subset S, \forall i \in S \setminus T, \quad D_{i,S}^t D_{T,S}^t = D_{T,S}^t D_{i,S}^t = D_{T \cup \{i\},S}^t. \quad (3)$$

Since $\text{SWAP} = \text{SWAP}^\dagger$ it also holds that $(C_i^t)^\dagger = C_i^t$ and $(D_{T,S}^t)^\dagger = D_{T,S}^t$.

Finally, we define an n -qubit mixed state

$$\sigma := \left(\bigotimes_{i=1}^n \text{DEC}_{R_i^1 \dots R_i^r} \right) \left(\text{Tr}_{\cup_t ((\cup_i R_i^t Q_i^t) S^t)} \left(\left(\bigotimes_{t=1}^r C_n^t \dots C_2^t C_1^t \right) |\tilde{\Psi}\rangle \langle \tilde{\Psi}| \left(\bigotimes_{t=1}^r (C_1^t)^\dagger \dots (C_n^t)^\dagger \right) \right) \right), \quad (4)$$

i.e. σ is the state obtained by, first applying unitaries C_1^t, \dots, C_n^t , for $t = 1, \dots, r$, to the original state $|\Psi\rangle$ and the auxiliary registers (initialized as EPR pairs), then tracing out all but the nr auxiliary registers R_1^t, \dots, R_n^t for $t = 1, \dots, r$, and finally applying the decoding map for code C independently to each group of r auxiliary registers $R_i^1 \dots R_i^r$.

3.2.2 Analysis of the strategy

In this section we prove the following lemma, which establishes soundness of protocol P .

Lemma 7. *There exists a universal constant $c_3 > 0$ (depending on k only) such that the following holds. Suppose a strategy for the provers is accepted with probability at least $1 - \varepsilon$ in each of the tests of protocol P , for some $\varepsilon > 0$. Then the state σ defined in (4) satisfies $\frac{1}{m} \text{Tr}(H\sigma) = O(n^{c_3} \varepsilon)$.*

The proof of the lemma follows from a sequence of claims. The first draws a useful consequence of the condition that the provers succeed in test (b) with high probability.

Claim 8. *Suppose the strategy $(U_i^j, V_S^j, |\Psi\rangle)$ succeeds in test (b) with probability at least $1 - \varepsilon$. For any $t \in [r]$, $i \in [n]$ and $S \subseteq [n]$ of cardinality k such that $i \in S$,*

$$\|(C_i^t - D_{i,S}^t) \otimes \text{Id} |\tilde{\Psi}\rangle\|^2 = O(n^k \varepsilon), \quad (5)$$

where $|\tilde{\Psi}\rangle$ is defined from $|\Psi\rangle$ in (1). Furthermore, for any set $S' \subseteq [n]$ of cardinality k and $T \subseteq S \cap S'$,

$$\|(D_{T,S}^t - D_{T,S'}^t) \otimes \text{Id} |\tilde{\Psi}\rangle\|^2 = O(n^k \varepsilon). \quad (6)$$

Proof. For any $t \in [r]$, $i \in [n]$ and set $S \subseteq [n]$ such that $i \in S$ let

$$|\varphi_i\rangle := \bigotimes_{p=1}^r C_i^p |\tilde{\Psi}\rangle \quad \text{and} \quad |\varphi_{i,S}\rangle := D_{i,S}^t \left(\bigotimes_{p \neq t} C_i^p \right) |\tilde{\Psi}\rangle, \quad (7)$$

where for ease of notation the dependence on t of $|\varphi_i\rangle$ and $|\varphi_{i,S}\rangle$ is left implicit. By definition, this strategy's success probability in test (b) of the protocol is exactly

$$\begin{aligned} & \frac{1}{r} \sum_{t=1}^r \frac{1}{n} \sum_{i=1}^n \frac{1}{\binom{n-1}{k-1}} \sum_{S: i \in S} \frac{1}{2} \left(\langle \Psi | \left(\bigotimes_{p=1}^r U_i^p \right)^\dagger \text{CHECK}_{Q_1^t \dots Q_r^t} \left(\bigotimes_{p=1}^r U_i^p \right) | \Psi \rangle \right. \\ & \quad \left. + \langle \Psi | \left(V_S^t \bigotimes_{p \neq t} U_i^p \right)^\dagger \text{CHECK}_{Q_1^t \dots Q_r^t} \left(V_S^t \bigotimes_{p \neq t} U_i^p \right) | \Psi \rangle \right). \end{aligned}$$

Let $\text{CK}_i := \text{CHECK}_{R_i^1 \dots R_i^{r_i}}$. Given the definition of C_i^t and $D_{i,S}^t$ in (2), success $1 - \varepsilon$ in test (b) of the protocol can be rewritten as

$$\frac{1}{n} \sum_{i=1}^n \frac{1}{\binom{n-1}{k-1}} \sum_{S:i \in S} \frac{1}{2} \left(\langle \varphi_i | \text{CK}_i | \varphi_i \rangle + \langle \varphi_{i,S} | \text{CK}_i | \varphi_{i,S} \rangle \right) \geq 1 - \varepsilon_t, \quad (8)$$

where the ε_t satisfy $(1/r)(\varepsilon_1 + \dots + \varepsilon_r) = \varepsilon$. Decompose the action of the unitary $D_{i,S}^t (C_i^t)^\dagger$ as

$$D_{i,S}^t (C_i^t)^\dagger = \text{Id}_{R_i^t} \otimes W_{i,S}^1 + X_{R_i^t} \otimes W_{i,S}^2 + Y_{R_i^t} \otimes W_{i,S}^3 + Z_{R_i^t} \otimes W_{i,S}^4, \quad (9)$$

where the Pauli operators $\{\text{Id}, X, Y, Z\}$ act on the i -th auxiliary register R_i^t associated with the t -th prover, and the $W_{i,S}^\ell$ are arbitrary operators (not necessarily unitary) of norm at most 1 acting on the remaining registers $Q_1^t \dots Q_n^t$ and S^t . Note that both C_i^t and $D_{i,S}^t$ are such that $\text{Tr}_{R_i^t}(C_i^t) = \text{Tr}_{R_i^t}(D_{i,S}^t) = \text{Id}_{Q_1^t \dots Q_n^t S^t}$, hence $W_{i,S}^1 = \text{Id}$. Let

$$|\varphi_i^s\rangle := (\text{CK}_i \otimes \text{Id}) |\varphi_i\rangle \quad \text{and} \quad |\varphi_i^f\rangle := ((\text{Id} - \text{CK}_i) \otimes \text{Id}) |\varphi_i\rangle,$$

so that $|\varphi_i\rangle = |\varphi_i^s\rangle + |\varphi_i^f\rangle$. By assumption the code C corrects all single-qubit Pauli errors, and since by definition the reduced density of $|\varphi_i^s\rangle$ on registers $R_i^1 \dots R_i^r$ is in the codespace, for any single-qubit Pauli error $E_{R_i^t} \in \{X, Y, Z\}$ acting on register R_i^t ,

$$\text{CHECK}_{R_i^1 \dots R_i^r} (E_{R_i^t} \otimes \text{Id}) |\varphi_i^s\rangle = 0. \quad (10)$$

As a consequence, starting from the definition of $|\varphi_{i,S}\rangle$ and using the decomposition (9) we get

$$\begin{aligned} \text{CK}_i |\varphi_{i,S}\rangle &= \text{CK}_i \cdot \left(D_{i,S}^t (C_i^t)^\dagger \otimes \text{Id} \right) |\varphi_i\rangle \\ &= \text{CK}_i \cdot \left(\text{Id}_{R_i^t} \otimes \text{Id} + X_{R_i^t} \otimes W_{i,S}^2 + Y_{R_i^t} \otimes W_{i,S}^3 + Z_{R_i^t} \otimes W_{i,S}^4 \right) |\varphi_i\rangle \\ &= \text{CK}_i \otimes \text{Id} |\varphi_i\rangle + (\text{CK}_i \cdot X_{R_i^t} \otimes W_{i,S}^2 + \text{CK}_i \cdot Y_{R_i^t} \otimes W_{i,S}^3 + \text{CK}_i \cdot Z_{R_i^t} \otimes W_{i,S}^4) |\varphi_i^s\rangle \\ &\quad + (\text{CK}_i \cdot X_{R_i^t} \otimes W_{i,S}^2 + \text{CK}_i \cdot Y_{R_i^t} \otimes W_{i,S}^3 + \text{CK}_i \cdot Z_{R_i^t} \otimes W_{i,S}^4) |\varphi_i^f\rangle \\ &= \text{CK}_i \otimes \text{Id} |\varphi_i\rangle + (\text{CK}_i \cdot X_{R_i^t} \otimes W_{i,S}^2 + \text{CK}_i \cdot Y_{R_i^t} \otimes W_{i,S}^3 + \text{CK}_i \cdot Z_{R_i^t} \otimes W_{i,S}^4) |\varphi_i^f\rangle, \end{aligned} \quad (11)$$

where the last equality follows from (10) and the fact that the $W_{i,S}^j$ do not act on R_i^t . Eq. (8) implies that both

$$\frac{1}{n} \sum_{i=1}^n \frac{1}{\binom{n-1}{k-1}} \sum_{S:i \in S} \left\| |\varphi_i^f\rangle \right\|^2 \leq 2\varepsilon_t \quad (12)$$

and

$$\frac{1}{n} \sum_{i=1}^n \frac{1}{\binom{n-1}{k-1}} \sum_{S:i \in S} \left\| (\text{Id} - \text{CHECK}_{R_i^1 \dots R_i^r}) |\varphi_{i,S}\rangle \right\|^2 \leq 2\varepsilon_t, \quad (13)$$

where we used that $\text{CHECK}_{R_i^1 \dots R_i^r}$ is a projection. Using the triangle inequality as

$$\left\| |\varphi_{i,S}\rangle - |\varphi_i\rangle \right\| \leq \left\| |\varphi_{i,S}\rangle - \text{CK}_i |\varphi_{i,S}\rangle \right\| + \left\| \text{CK}_i |\varphi_{i,S}\rangle - \text{CK}_i |\varphi_i\rangle \right\| + \left\| \text{CK}_i |\varphi_i\rangle - |\varphi_i\rangle \right\|$$

we get

$$\frac{1}{n} \sum_{i=1}^n \frac{1}{\binom{n-1}{k-1}} \sum_{S: i \in S} \|\varphi_{i,S} - |\varphi_i\rangle\|^2 \leq 3(2\varepsilon_t + 9 \cdot 2\varepsilon_t + 2\varepsilon_t) = O(\varepsilon_t), \quad (14)$$

where the first bound is obtained from (13), the second from (11), (12) and $\|\text{CK}_i\|, \|W_{i,j}^\ell\| \leq 1$, and the third from the definition of $|\varphi_i^f\rangle$ and (12). Recalling the definition of $|\varphi_i\rangle$ and $|\varphi_{i,S}\rangle$ in (7), (5) is proved by noting that the operator $(\text{Id} \otimes_{p \neq t} C_i^p)$ is unitary and hence its application does not modify the Euclidean norm.

The proof of (6) follows the same steps. Defining vectors $|\varphi_{T,S}\rangle$ and $|\varphi_{T,S'}\rangle$ and using that (8) is satisfied for every $i \in T$ we can decompose $D_{T,S}^t (D_{T,S'}^t)^\dagger$ as in (9), except now the decomposition involves all $|T|$ -qubit Pauli operators on registers R_i^t for $i \in T$. The different qubits are checked independently, and we can define $|\varphi_{T,S'}^s\rangle := (\otimes_{i \in T} \text{CK}_i) |\varphi_{T,S'}\rangle$. The remainder of the derivation follows the same steps, leading to (6) (where factors polynomial in k are hidden in the $O(\cdot)$ notation, using that k is a constant independent of n). \square

For any $i \in [n]$ let \mathcal{F}_i be the completely positive trace non-increasing map, acting on all provers' registers, defined by

$$\mathcal{F}_i : \sigma \mapsto \left(\left(\bigotimes_{j=1}^r X_i^j \right)^\dagger \text{CK}_{Q_i^1 \dots Q_i^r} \left(\bigotimes_{j=1}^r X_i^j \right) \right) \sigma \left(\left(\bigotimes_{j=1}^r X_i^j \right)^\dagger \text{CK}_{Q_i^1 \dots Q_i^r} \left(\bigotimes_{j=1}^r X_i^j \right) \right)^\dagger. \quad (15)$$

Here we use the symbol X_i^j to represent any of U_i^j or V_S^j for any S containing i ; we leave the dependence of \mathcal{F}_i on the choice of X_i^j implicit as all bounds proved will hold irrespective of that choice. We also write $\mathcal{X}_i^j : \sigma \rightarrow X_i^j \sigma (X_i^j)^\dagger$ for the CPTP map associated with X_i^j . Note that, in addition to the presence of the CK operator, the difference between the maps \mathcal{F}_i and e.g. $\otimes_j C_i^j$ is that in the former the t registers Q_i and R_i are not swapped; in particular \mathcal{F}_i acts as identity on R_i .

Our second claim shows that the property that the qubits extracted from the provers' strategies through the maps X_i^j are in the codespace remains preserved even after many layers of application of the \mathcal{F}_i .

Claim 9. *Suppose the strategy $(U_i^j, V_S^j, |\Psi\rangle)$ succeeds in test (b) with probability at least $1 - \varepsilon$. Let s be an integer and $i_1, \dots, i_s \in [n]$. Then*

$$\text{Tr} \left(\left(\bigcirc_{\ell=1}^s \mathcal{F}_{i_\ell} \right) (\tilde{\rho}) \right) = 1 - O(sn^k \varepsilon). \quad (16)$$

Proof. We prove (16) by induction on s . For $s = 1$ it follows immediately by first applying Claim 8 (at most r times to replace each $X_{i_1}^j$ in the definition of \mathcal{F}_{i_1} by $U_{i_1}^j$, and then using the assumption of success in the test, which ensures that the extracted qubits are close to the code space. Suppose (16) verified for some s , and let K be the constant implicit in the $O(\cdot)$ notation; we show it for $s + 1$. Writing $\text{CK}_{i_1} = \text{Id} - (\text{Id} - \text{CK}_{i_1})$

and using that \mathcal{F}_{i_1} reduces to the identity once the operator $\text{CK}_{Q_{i_1}^1 \dots Q_{i_1}^r}$ is removed,

$$\begin{aligned} \text{Tr}\left(\left(\bigcirc_{\ell=1}^{s+1} \mathcal{F}_{i_\ell}\right)(\tilde{\rho})\right) &= \text{Tr}\left(\left(\bigcirc_{\ell=2}^{s+1} \mathcal{F}_{i_\ell}\right)(\tilde{\rho})\right) \\ &\quad - \text{Tr}\left(\left(\bigcirc_{\ell=2}^{s+1} \mathcal{F}_{i_\ell}\right) \circ \left(\bigotimes_{j=1}^r \mathcal{X}_{i_1}^j\right)^\dagger \left(\left(\text{Id} - \text{CK}_{Q_{i_1}^1 \dots Q_{i_1}^r}\right) \left(\bigotimes_{j=1}^r \mathcal{X}_{i_1}^j\right) (\tilde{\rho}) \text{CK}_{Q_{i_1}^1 \dots Q_{i_1}^r}\right)\right) \\ &\geq 1 - Ksn^k \varepsilon - \text{Tr}\left(\left(\text{Id} - \text{CK}_{Q_{i_1}^1 \dots Q_{i_1}^r}\right) \left(\bigotimes_{j=1}^r \mathcal{X}_{i_1}^j\right) (\tilde{\rho})\right) \\ &\geq 1 - Ksn^k \varepsilon - O(n^k \varepsilon), \end{aligned}$$

where the first inequality uses the induction hypothesis for the first term, and that the \mathcal{F}_{i_1} are trace non-increasing for the second, and the last follows from the case $s = 1$ of (16). Provided K is chosen large enough this establishes the induction step and proves the claim. \square

The next claim has a similar flavor as the previous one, that the qubits extracted from the provers' strategies lie in the codespace is preserved even after application of a sequence of maps \mathcal{C}_i^t or $\mathcal{D}_{i,S}^t$ on one of the provers' registers.

Claim 10. *There exists a constant $c_1 > 0$ depending on k only such that the following holds. Suppose the strategy $(U_i^j, V_S^j, |\Psi\rangle)$ succeeds in test (b) with probability at least $1 - \varepsilon$. Let s be an integer and $i_1, \dots, i_s \in [n]$. Then for any $t \in [r]$ and choice of $\mathcal{Y}_{i_\ell}^j \in \{\mathcal{C}_{i_\ell}^j, \mathcal{D}_{i_\ell, S_\ell}^j \mid i_\ell \in S_\ell\}$ for $j \in [r]$ and $\ell \in [s]$,*

$$\text{Tr}\left(\text{CK}_{i_s} \left(\left(\bigcirc_{\ell=1}^s \mathcal{Y}_{i_\ell}^t\right) \bigotimes_{j \neq t} \mathcal{Y}_{i_s}^j\right) (\tilde{\rho})\right) = 1 - O(s^2 n^{c_1} \varepsilon). \quad (17)$$

Proof. For the proof we show that the following holds by downwards induction on s' from s to 1:

$$\text{Tr}\left(\text{CK}_{i_s} \left(\left(\bigcirc_{\ell=s'}^s \mathcal{Y}_{i_\ell}^t\right) \bigotimes_{j \neq t} \mathcal{Y}_{i_s}^j\right) \circ \left(\bigcirc_{\ell=1}^{s'-1} \mathcal{F}_{i_\ell}\right) (\tilde{\rho})\right) = 1 - O(ss' n^{c_1} \varepsilon). \quad (18)$$

Eq. (18) for $s' = s$ is equivalent to (16), so Claim 9 proves the base case for the induction. If $s' = 1$ it reduces to (17), which is what we need to prove. Assume thus (18) verified for $s' + 1$, and prove it for s' . By (16) applied with $s = s'$ we see that the strategy $(U_i^j, V_S^j, \tilde{\rho}_{s'})$, where $\tilde{\rho}_{s'} = \bigcirc_{\ell=1}^{s'} \mathcal{F}_{i_\ell}(\rho)$ (it will be more convenient to leave $\tilde{\rho}_{s'}$ unnormalized) succeeds in test (b) with probability at least $1 - O(sn^k \varepsilon)$. Here in the definition of $\mathcal{F}_{i_{s'}}$ we define $\mathcal{X}_{i_{s'}}^t$ as $\mathcal{U}_{i_{s'}}^t$ if $\mathcal{Y}_{i_{s'}}^t = \mathcal{C}_{i_{s'}}^t$ and $\mathcal{V}_{i_{s'}}^t$ if $\mathcal{Y}_{i_{s'}}^t = \mathcal{D}_{i_{s'}, S}^t$; the remaining $\mathcal{X}_{i_\ell}^j$ can be chosen arbitrarily. Applying Claim 8 we get

$$\left\| \left(\bigcirc_{\ell=s'+1}^s \mathcal{Y}_{i_\ell}^t\right) \bigotimes_{j \neq t} \mathcal{Y}_{i_s}^j (\tilde{\rho}_{s'}) - \left(\bigcirc_{\ell=s'+1}^s \mathcal{Y}_{i_\ell}^t\right) \bigotimes_{j \neq t} \mathcal{D}_{i_s, S}^j (\tilde{\rho}_{s'}) \right\|_1 = O(sn^{2k} \varepsilon), \quad (19)$$

where for S we choose any set containing both $i_{s'}$ and i_s . Applying the claim once more, this time starting from the strategy $(U_i^j, V_S^j, \tilde{\rho}_{s'-1})$, we have

$$\left\| \mathcal{F}_{i_{s'}}(\tilde{\rho}_{s'-1}) - \left(\mathcal{X}_{i_{s'}}^t \bigotimes_{j \neq t} \mathcal{V}_S^j\right)^\dagger \left(\text{CK}_{Q_{i_{s'}}^1 \dots Q_{i_{s'}}^r} \left(\mathcal{X}_{i_{s'}}^t \bigotimes_{j \neq t} \mathcal{V}_S^j\right) (\tilde{\rho}_{s'-1}) \text{CK}_{Q_{i_{s'}}^1 \dots Q_{i_{s'}}^r}\right) \right\|_1 = O(sn^{2k} \varepsilon). \quad (20)$$

Combining (19) and (20) by the triangle inequality and evaluating the overlap with CK_{i_s} we get

$$\begin{aligned} & \left| \text{Tr} \left(\text{CK}_{i_s} \left(\left(\bigcirc_{\ell=s'+1}^s \mathcal{Y}_{i_\ell}^t \right) \bigotimes_{j \neq t} \mathcal{Y}_{i_s}^j \right) (\tilde{\rho}_{s'}) \right) \right. \\ & \quad \left. - \text{Tr} \left(\text{CK}_{i_s} \left(\bigcirc_{\ell=s'+1}^s \mathcal{Y}_{i_\ell}^t \circ (\mathcal{X}_{i_{s'}}^t)^\dagger \right) \left(\text{CK}_{\mathcal{Q}_{i_{s'}}^1 \dots \mathcal{Q}_{i_{s'}}^r} \left(\mathcal{X}_{i_{s'}}^t \bigotimes_{j \neq t}^r \mathcal{V}_S^j \right) (\tilde{\rho}_{s'-1}) \text{CK}_{\mathcal{Q}_{i_{s'}}^1 \dots \mathcal{Q}_{i_{s'}}^r} \right) \right) \right| = O(sn^{2k}\varepsilon), \end{aligned} \quad (21)$$

where we also used the definition of $\mathcal{D}_{i_{s'}, S}^j$ to simplify the successive application of unitaries V_S^j and $(V_S^j)^\dagger$ on provers $j \neq t$ in the second term above. The fact that every codeword of the code \mathcal{C} has a reduced density on any single qubit that is totally mixed implies that if we trace out registers $\mathcal{Q}_{i_{s'}}^{\neq t}$ and $\bar{\mathcal{R}}_{i_{s'}}^t$ in the (unnormalized) density $\tilde{\sigma}_{s'} = \text{CK}_{i_{s'}}(\mathcal{X}_{i_{s'}}^t \bigotimes_{j \neq t} \mathcal{V}_S^j)(\tilde{\rho}_{s'-1})\text{CK}_{i_{s'}}$, the registers $\mathcal{R}_{i_{s'}}^t \mathcal{Q}_{i_{s'}}^t$ are jointly in the totally mixed state. Swapping the two registers thus leaves the state invariant, and from (21) we get

$$\begin{aligned} & \left| \text{Tr} \left(\text{CK}_{i_s} \left(\left(\bigcirc_{\ell=s'+1}^s \mathcal{Y}_{i_\ell}^t \right) \bigotimes_{j \neq t} \mathcal{Y}_{i_s}^j \right) (\tilde{\rho}_{s'}) \right) \right. \\ & \quad \left. - \text{Tr} \left(\text{CK}_{i_s} \left(\bigcirc_{\ell=s'+1}^s \mathcal{Y}_{i_\ell}^t \right) \left(\text{CK}_{\mathcal{R}_{i_{s'}}^t \mathcal{Q}_{i_{s'}}^{\neq t}} \left(\mathcal{Y}_{i_{s'}}^t \bigotimes_{j \neq t}^r \mathcal{V}_S^j \right) (\tilde{\rho}_{s'-1}) \text{CK}_{\mathcal{R}_{i_{s'}}^1 \mathcal{Q}_{i_{s'}}^{\neq t}} \right) \right) \right| = O(sn^{2k}\varepsilon). \end{aligned} \quad (22)$$

Finally, using Claim 9 once more, we can remove the application of $\text{CK}_{\mathcal{R}_{i_{s'}}^1 \mathcal{Q}_{i_{s'}}^{\neq t}}$ in the second term in (22) to obtain

$$\left| \text{Tr} \left(\text{CK}_{i_s} \left(\left(\bigcirc_{\ell=s'+1}^s \mathcal{Y}_{i_\ell}^t \right) \bigotimes_{j \neq t} \mathcal{Y}_{i_s}^j \right) (\tilde{\rho}_{s'}) \right) - \text{Tr} \left(\text{CK}_{i_s} \left(\left(\bigcirc_{\ell=s'+1}^s \mathcal{Y}_{i_\ell}^t \right) \bigotimes_{j \neq t} \mathcal{V}_S^j \right) (\tilde{\rho}_{s'-1}) \right) \right| = O(sn^{2k}\varepsilon). \quad (23)$$

To conclude, since the registers $\mathcal{Q}_{i_{s'}}^{\neq t}$ are being traced out, and using Claim 8 for the strategy $(U_i^j, V_S^j, \tilde{\rho}_{s'-1})$ the application of $\bigotimes_{j \neq t} \mathcal{V}_S^j$ in the second term in (23) can be replaced by $\bigotimes_{j \neq t} \mathcal{Y}_{i_s}^j$ for \mathcal{Y} of our choice. Using the induction hypothesis (18) to bound the first term in (23), this proves (18) for s' and establishes the induction step, proving the claim. \square

The following corollary is a simple consequence of Claim 8.

Corollary 11. *Suppose the strategy $(U_i^j, V_S^j, |\Psi\rangle)$ succeeds in test (b) with probability at least $1 - \varepsilon$. Let s be an integer and $i_1, \dots, i_s \in [n]$. Then for any $t \in [r]$, set S containing i_s , and choice of $\mathcal{Y}_{i_\ell}^t \in \{\mathcal{C}_{i_\ell}^t, \mathcal{D}_{i_\ell, S_\ell}^t \mid i_\ell \in S_\ell\}$ for $\ell \in [s-1]$,*

$$\left\| \left(\left(\mathcal{C}_{i_s}^t \bigcirc_{\ell=1}^{s-1} \mathcal{Y}_{i_\ell}^t \right) \otimes \text{Id} \right) (\tilde{\rho}) - \left(\left(\mathcal{D}_{i_s, S}^t \bigcirc_{\ell=1}^{s-1} \mathcal{Y}_{i_\ell}^t \right) \otimes \text{Id} \right) (\tilde{\rho}) \right\|_1 = O(s^2 n^{c_1+k} \varepsilon), \quad (24)$$

where c_1 is as in Claim 10.

Proof. Using the freedom in the choice of the operators \mathcal{Y} , (17) from Claim 10 shows that the strategy $(U_i^j, V_S^j, (\bigcirc_{\ell=1}^{s-1} \mathcal{Y}_{i_\ell}^t) \otimes \text{Id})(\tilde{\rho})$ succeeds with probability $1 - O(s^2 n^{c_1} \varepsilon)$ in test (b) of the protocol. The corollary then follows directly from Claim 8. \square

Our final claim shows that if the provers have a high success probability in both tests of protocol P the state σ defined in (4) must have low energy with respect to the local Hamiltonian H .

Claim 12. *There exists a constant $c_2 > 0$ depending on k only such that the following holds. Let $\delta, \varepsilon > 0$ be such that the provers succeed in test (a) of protocol P with probability at least $1 - \delta$, and in test (b) with probability at least $1 - \varepsilon$. Then*

$$\frac{1}{m} \text{Tr}(H\sigma) = O(\delta + n^{c_2}\varepsilon).$$

Proof. For any $i \in [n]$ we abbreviate $\text{DEC}_{R_1^i \dots R_r^i}$ as DEC_i . By definition,

$$\sigma = \left(\bigotimes_{i=1}^n \text{DEC}_{R_1^i \dots R_r^i} \right) \left(\text{Tr}_{U_t(U_i(\overline{R}_i^t Q_i^t) S^t)}(\tau) \right),$$

where

$$\tau := \left(\bigotimes_{j=1}^r (\mathcal{C}_n^j \circ \dots \circ \mathcal{C}_2^j \circ \mathcal{C}_1^j) \right) (\tilde{\rho}).$$

Fix a local term H_j acting on k qubits $S := S_j = \{i_1, \dots, i_k\}$, and let $U = [n] \setminus S$. Let $T_0 = \emptyset$, and for $s = 1, \dots, k$ let $T_s = \{i_1, \dots, i_s\}$. We show the following by induction on $s = 0, \dots, k$:

$$\left\| \text{Tr}_{Q_U \overline{R}_U}(\tau) - \text{Tr}_{Q_U \overline{R}_U} \left(\bigotimes_{j=1}^r \left(\bigcirc_{\substack{i=1 \\ i \notin T_s}}^n \mathcal{C}_i^j \right) \circ \left(\bigotimes_{j=1}^r \mathcal{D}_{T_s, S}^j \right) (\tilde{\rho}) \right) \right\|_1 = O(sn^{c_2}\varepsilon), \quad (25)$$

for some constant $c_2' > 0$. The equation is trivially true for $s = 0$. Suppose it true for some $s < k$. Let $\tilde{\rho}_s = \bigotimes_{j=1}^r \mathcal{D}_{T_s, S}^j(\tilde{\rho})$. We again proceed by induction on $\ell = i_{s+1}, \dots, 1$ to show

$$\begin{aligned} & \left\| \text{Tr}_{Q_U \overline{R}_U} \left(\left(\bigotimes_{j=1}^r \left(\bigcirc_{\substack{i=1 \\ i \notin T_s}}^{i_{s+1}} \mathcal{C}_i^j \right) \right) (\tilde{\rho}_s) \right) - \text{Tr}_{Q_U \overline{R}_U} \left(\left(\bigotimes_{j=1}^r \left(\bigcirc_{\substack{i=\ell \\ i \notin T_s}}^{i_{s+1}-1} \mathcal{C}_i^j \right) \circ \mathcal{D}_{i_{s+1}, S}^j \bigcirc_{\substack{i=1 \\ i \notin T_s}}^{\ell-1} \mathcal{C}_i^j \right) \right) (\tilde{\rho}_s) \right\|_1 \\ & = O((i_{s+1} + 1 - \ell)n^{c_2'-1}\varepsilon). \end{aligned} \quad (26)$$

For $\ell = i_{s+1}$ the bound follows from Corollary 11 applied to each of the provers, provided c_2' is chosen large enough. For $\ell = 1$, using $\mathcal{D}_{i_{s+1}, S}^j \circ \mathcal{D}_{T_s, S}^j = \mathcal{D}_{T_{s+1}, S}^j$, together with the triangle inequality it establishes the induction step for the proof of (25). Suppose (26) verified for some $\ell > 1$. To prove it for $\ell - 1$, we apply Corollary 11 to the state

$$\left(\bigotimes_{j=1}^r \left(\bigcirc_{\substack{i=\ell \\ i \notin T_s}}^{i_{s+1}-1} \mathcal{C}_i^j \right) \circ \mathcal{D}_{i_{s+1}, S}^j \bigcirc_{\substack{i=1 \\ i \notin T_s}}^{\ell-1} \mathcal{C}_i^j \right) (\tilde{\rho}_s)$$

a number of times: first, the maps $\mathcal{C}_{\ell-1}^j$ are replaced by $\mathcal{D}_{\ell-1, S'}^j$ for some S' containing both $\ell - 1$ and i_{s+1} . Second, the $\mathcal{D}_{i_{s+1}, S}^j$ are replaced by $\mathcal{D}_{i_{s+1}, S'}^j$. Next we use the relation $\mathcal{D}_{i_{s+1}, S'}^j \circ \mathcal{D}_{\ell-1, S'}^j = \mathcal{D}_{\ell-1, S'}^j \circ \mathcal{D}_{i_{s+1}, S'}^j$ and perform the same replacements in reverse. This establishes (26) for $\ell - 1$ (provided c_2' is chosen large enough) and completes the induction step. We have now proven (25).

Using the definition of $\mathcal{D}_{T_k, S}$ and that both H_j and DEC_{i_s} , for $s = 1, \dots, k$, do not act on any of the registers in Q_U or \overline{R}_U , from (25) we get

$$\left| \text{Tr}(H_j \sigma) - \text{Tr} \left(H_j \left(\bigotimes_{\ell=1}^k \text{DEC}_{i_\ell} \right) \left(\bigotimes_{t=1}^r V_{S_j}^t \right) \tilde{\rho} \left(\bigotimes_{t=1}^r V_{S_j}^t \right)^\dagger \right) \right| = O(n^{c_2'+1}\varepsilon). \quad (27)$$

By definition success $1 - \delta$ in test (a) of protocol P implies

$$\frac{1}{m} \sum_{S_j=\{i_1, \dots, i_k\}} \text{Tr} \left(H_j \left(\bigotimes_{s=1}^k \text{DEC}_{i_s} \right) \left(\left(\bigotimes_{t=1}^r V_{S_j}^t \right) \tilde{\rho} \left(\bigotimes_{t=1}^r V_{S_j}^t \right)^\dagger \right) \right) \leq \delta,$$

which together with (27) proves the claim for an appropriate choice of c_2 . \square

Lemma 7 now follows directly from Claim 12 and the fact that any strategy with success $1 - \varepsilon$ in Protocol P must have success probability at least $1 - 2\varepsilon$ in each of the two tests (a) and (b) of the protocol.

References

- [Aar06] Scott Aaronson. The quantum PCP manifesto, 2006. Blog entry available at <http://www.scottaaronson.com/blog/?p=139>.
- [AAV13] Dorit Aharonov, Itai Arad, and Thomas Vidick. The quantum PCP conjecture. Technical report, arXiv:1309.7495, 2013. Appeared as guest column in ACM SIGACT News archive Volume 44 Issue 2, June 2013, Pages 47–79.
- [ALM⁺98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998.
- [AN02] Dorit Aharonov and Tomer Naveh. Quantum NP – a survey. Technical report, arXiv:quant-ph/0210077, 2002.
- [AS98] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of NP. *J. ACM*, 45(1):70–122, 1998.
- [BDSW96] Charles H Bennett, David P DiVincenzo, John A Smolin, and William K Wootters. Mixed-state entanglement and quantum error correction. *Phys. Rev. A*, 54(5):3824, 1996.
- [BFL91] László Babai, Lance Fortnow, and Carsten Lund. Non-deterministic exponential time has two-prover interactive protocols. *Comput. Complexity*, 1:3–40, 1991.
- [BH13] Fernando G.S.L. Brandao and Aram W. Harrow. Product-state approximations to quantum ground states. In *Proc. 45th STOC*, 2013.
- [Din07] Irit Dinur. The PCP theorem by gap amplification. *J. ACM*, 54(3), June 2007.
- [FH13] Michael H. Freedman and Matthew B. Hastings. Quantum systems on non- k -hyperfinite complexes: A generalization of classical statistical mechanics on expander graphs. *arXiv preprint arXiv:1301.1363*, 2013.
- [GHL14] Sevag Gharibian, Yichen Huang, and Zeph Landau. Quantum hamiltonian complexity. Technical report, arXiv:1401.3916, 2014.
- [GI09] Daniel Gottesman and Sandy Irani. The quantum and classical complexity of translationally invariant tiling and hamiltonian problems. In *Proc. 50th FOCS*, pages 95–104, Oct 2009.
- [Hås01] Johan Håstad. Some optimal inapproximability results. *J. ACM*, 48:798–859, 2001.

- [Has13] Matthew B. Hastings. Trivial low energy states for commuting Hamiltonians, and the quantum PCP conjecture. *Quantum Information and Computation*, 13(5 & 6):393–429, 2013.
- [IV12] Tsuyoshi Ito and Thomas Vidick. A multi-prover interactive proof for NEXP sound against entangled provers. *Proc. 53rd FOCS*, pages 243–252, 2012.
- [JNP⁺11] Marius Junge, Miguel Navascues, Carlos Palazuelos, David Perez-Garcia, Volkher B. Scholz, and Reinhard F. Werner. Connes’ embedding problem and tsirelson’s problem. *J. Math. Physics*, 52(1):–, 2011.
- [KKMV09] Julia Kempe, Hirotada Kobayashi, Keiji Matsumoto, and Thomas Vidick. Using entanglement in quantum multi-prover interactive proofs. *Computational Complexity*, 18:273–307, 2009.
- [KKR06] Julia Kempe, Alexei Kitaev, and Oded Regev. The complexity of the local hamiltonian problem. *SIAM J. Comput.*, 35(5):1070–1097, May 2006.
- [KM03] Hirotada Kobayashi and Keiji Matsumoto. Quantum multi-prover interactive proof systems with limited prior entanglement. *Journal of Computer and System Sciences*, 66(3):429–450, 2003.
- [KSV02] Alexei Yu. Kitaev, Alexander H. Shen, and Mikhail N. Vyalyi. *Classical and Quantum Computation*, volume 47 of *Graduate Studies in Mathematics*. American Mathematical Society, 2002.
- [Liu06] Yi-Kai Liu. Consistency of local density matrices is qma-complete. In Josep Daz, Klaus Jansen, JosD.P. Rolim, and Uri Zwick, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, volume 4110 of *Lecture Notes in Computer Science*, pages 438–449. Springer Berlin Heidelberg, 2006.
- [LMPZ96] Raymond Laflamme, Cesar Miquel, Juan Pablo Paz, and Wojciech Hubert Zurek. Perfect quantum error correcting code. *Phys. Rev. Lett.*, 77(1):198, 1996.
- [MW05] Chris Marriott and John Watrous. Quantum Arthur—Merlin games. *Comput. Complexity*, 14(2):122–152, June 2005.
- [Osb12] Tobias J Osborne. Hamiltonian complexity. *Reports on Progress in Physics*, 75(2):022001, 2012.
- [Raz98] Ran Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27:763–803, 1998.
- [RUV13] Ben Reichardt, Falk Unger, and Umesh Vazirani. A classical leash for a quantum system: Command of quantum systems via rigidity of CHSH games. *Nature*, 496(7446):456–460, 2013.
- [SW08] Volkher B. Scholz and Reinhard F. Werner. Tsirelson’s problem. Technical report, arXiv:0812.4305v1 [math-ph], 2008.
- [Vid13] Thomas Vidick. Three-player entangled XOR games are NP-hard to approximate. In *Proc. 54th FOCS*, 2013.
- [Wat09] John Watrous. Quantum computational complexity. In *Encyclopedia of complexity and systems science*, pages 7174–7201. Springer, 2009.