Constant-Soundness Interactive Proofs for Local Hamiltonians

Anand Natarajan* Thomas Vidick†

December 8, 2015

Abstract

We give a quantum multiprover interactive proof system for the local Hamiltonian problem in which there is a constant number of provers, questions are classical of length polynomial in the number of qubits, and answers are of constant length. The main novelty of our protocol is that the gap between completeness and soundness is directly proportional to the promise gap on the (normalized) ground state energy of the Hamiltonian. This result can be interpreted as a concrete step towards a quantum PCP theorem giving entangled-prover interactive proof systems for QMA-complete problems.

The key ingredient is a quantum version of the classical linearity test of Blum, Luby, and Rubinfeld, where the function $f : \{0,1\}^n \to \{0,1\}$ is replaced by a pair of functions $\mathcal{X}, \mathcal{Z} : \{0,1\}^n \to \mathrm{Obs}_d(\mathbb{C})$, the set of d-dimensional Hermitian matrices that square to identity. The test enforces that (i) each function is exactly linear, $\mathcal{X}(a)\mathcal{X}(b) = \mathcal{X}(a+b)$ and $\mathcal{Z}(a)\mathcal{Z}(b) = \mathcal{Z}(a+b)$, and (ii) the two functions are approximately complementary, $\mathcal{X}(a)\mathcal{Z}(b) \approx (-1)^{a\cdot b}\mathcal{Z}(b)\mathcal{X}(a)$.

1 Introduction

The theory of NP-completeness is a central part of complexity theory, and an important area of research in quantum complexity theory over the last two decades has been to characterize the quantum analog of NP-completeness. The foundations for this were laid by Kitaev, who established a quantum version of the Cook-Levin theorem [KSV02]. His result shows that the *local Hamiltonian problem* is complete for the complexity class QMA, the quantum analog of NP. The local Hamiltonian problem can be cast as a quantum analog of Boolean constraint satisfaction problems (CSPs): instead of the satisfiability of a formula consisting of a conjunction of clauses each acting on a few Boolean variables, one considers the problem of finding the minimum eigenvalue of a Hermitian operator (the Hamiltonian) consisting of the sum of *local terms*, each acting on a constant number of quantum bits (qubits). Boolean CSPs are a special case of the local Hamiltonian problem, obtained by restricting all terms in the Hamiltonian to be matrices diagonal in the computational basis. The Hamiltonian operator plays a fundamental role in quantum mechanics, and the constraint of locality is motivated by problems considered in many-body physics where physical interactions typically only involve small groups of neighboring particles. Eigenvalues of

*Center for Theoretical Physics, MIT, Cambridge, USA. email:anandn@mit.edu.

†Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, USA. email: vidick@cms.caltech.edu.

the Hamiltonian are called energy levels, and the study of the minimal energy (smallest eigenvalue) and associated eigenvector (the ground state) is the focus of condensed-matter physics, as they correspond to the energy and equilibrium state of the system at zero temperature respectively. The local Hamiltonian problem therefore provides a compelling abstraction within which to frame many of the computational problems that arise in fundamental areas of physics.

The NP-hardness of deciding exact satisfiability of Boolean CSPs stated in the Cook-Levin theorem has been greatly strengthened by the PCP theorem [ALM+98, AS98], which extends NP-hardness to the problem of *approximating* the maximum satisfiable fraction of clauses in a formula, even up to *constant factors*. This result forms a cornerstone of modern complexity theory, and in particular implies many optimal (under P$\neq$NP) hardness-of-approximation results for NP-complete problems of independent interest [FGL+96]. The existence of a quantum analog of the PCP theorem, stating QMA-hardness of constant-factor approximations to the minimal energy of a local Hamiltonian H (normalized so that $\|H\| = O(1)$), is a major open problem in quantum complexity theory called the quantum PCP conjecture [AN02, AAV13]. For reasons that will soon be clear we refer to this conjecture as the "constraint satisfaction" variant of QPCP.

The PCP theorem has several different formulations and proofs; arguably the simplest is a combinatorial proof developed by Dinur [Din07]. This proof seems challenging to quantize [AALV09]; for instance, it relies heavily on copying bits, but copying qubits is forbidden by the no-cloning theorem (see [AAV13, Section 3] for a discussion of the many more difficulties that arise). The original proof of the theorem was quite different, drawing very strongly on the new connections that were being established between the theory of error-correcting codes, testing, and the power of interactive proof systems [Aro94]. A major milestone along this route is the inclusion NEXP $\subseteq$ MIP [BFL91], a result that can be interpreted as a "scaled up" precursor to the PCP theorem. A quantum analog of this result was first suggested as an alternative formulation of (and as a step towards a proof of) the quantum PCP conjecture in [FV15]: does the inclusion $\mathrm{QMA}_{\mathrm{EXP}} \subseteq$ QMIP* hold?[1] We henceforth refer to this inclusion as the "multiplayer games" variant of QPCP. Although both formulations, "multiplayer games" and "constraint satisfaction", of the classical PCP theorem are easily seen to be equivalent, whether a similar equivalence holds in the quantum world is an interesting open problem; we refer to [AAV13, Section 5.4] for a more in-depth discussion.

1.1 Main result

We make progress on the "multiplayer games" variant of QPCP by considering a quantum analog of what is often presented as the first step of the proof of the classical PCP theorem: the (games variant of the) exponentially long PCP for NP, based on the linearity test of Blum, Luby and Rubinfeld [BLR93] (see e.g. Theorem 18.21 in the book [AB09] for a precise formulation). Specifically, we consider the question of proving the inclusion QMA $\subseteq$ MIP*, where the interactive protocol is restricted to a single round with constant answer length (but polynomial question length). The analogous statement with NP and MIP in place of QMA and MIP* is exactly the classical exponentially long PCP.

In fact the inclusion we are seeking *does* hold, and follows from known results in complexity theory. For instance, combining the trivial inclusions QMA $\subseteq$ EXP $\subseteq$ NEXP with NEXP $\subseteq$

[1] Here QMIP* denotes the class of languages that have multi-prover interactive proofs with a quantum polynomial-time verifier and quantum entangled provers; $\mathrm{QMA}_{\mathrm{EXP}}$ is to QMA what NEXP is to NP.

MIP* [IV12] together with the fact that the latter holds for a protocol involving a single round of interaction and a constant answer length [Vid13] suffices to establish the result. A different, perhaps simpler route would be to use QMA $\subseteq$ PSPACE and IP $=$ PSPACE [Sha92], applying arithmetization to obtain an IP protocol and then introducing a second prover and the technique of [KKMV09] to parallelize the protocol to a single round of interaction with two entangled provers; finally one would have to apply some form of parallel amplification [KV11, BVY15] technique and finish with a method that allows a reduction in the answer length [Vid13].

Working through the reductions implied by either of these routes leads to a complex protocol in which the structure of the original instance of the local Hamiltonian problem has all but disappeared, and the "proof" held by the provers bears little relation to the ground state of the local Hamiltonian — even though it still, of course, suffices to certify its ground state energy.

In this work we provide a simpler, more direct and arguably "more quantum" construction of an "exponentially-long quantum PCP". The key ingredient of our protocol is a quantum generalization of the BLR linearity test — a "truly quantum" generalization in the sense that honest provers are *required* to apply quantum operations on a shared entangled state in order to achive completeness. (This is in contrast to the entangled-prover linearity test of [IV12], where the entanglement between the provers is treated as a hurdle against which the protocol has to be "immunized".) Moreover, our protocol uses the structure of the local Hamiltonian problem in a natural way: honest provers are asked to share a distributed encoding of the ground state, on which they perform measurements in order to determine their answers. Theorem 1 gives a precise statement of our result.

Theorem 1. *Let H be a local Hamiltonian on n qubits such that $\|H\| \leq 1$, and $\lambda_{min} \in [0,1]$ its smallest eigenvalue. Let $0 \leq a(n) < b(n) \leq 1$ be such that $b(n) - a(n) \geq 1/\operatorname{poly}(n)$. Then there exists an interactive proof system between a classical polynomial-time verifier and seven entangled provers that decides whether $\lambda_{min} \leq a(n)$ or $\lambda_{min} \geq b(n)$ with completeness $2/3$ and soundness $1/3$. The proof system involves a single round of interaction in which the verifier sends $\operatorname{poly}(n)$ bits to each prover and receives $O(1)$ bits from each.*

Beyond the statement of the theorem itself, we believe our proof technique is an important contribution in the quest for a quantum PCP theorem, either in its "constraint satisfaction" or "multiplayer games" variants. The conjecture is a distant target, and its resolution will undoubtedly require many new ideas. Although the two proofs of the classical PCP theorem provide valuable starting points, it also seems important to find the "right" quantum generalizations of the main ingredients — possibly requiring altogether new ones in order to overcome the specific obstacles posed by quantum information (see e.g. [AAV13, Section 3]). Our results mark a step in this programme. To describe them further we introduce our quantum linearity test next, and then explain its use in devising a protocol for the proof of Theorem 1.

1.2 A quantum linearity test

What is a good quantum analogue of the linearity test? Recall that in our context instead of evaluating *clauses* on *Boolean variables* we are interested in evaluating *local Hamiltonians* on *qubits*. In addition we may without loss of generality assume that each local term is a tensor product of single-qubit Pauli operators $\mathcal{P} = \{I, X, Z\}$ (defined in (1)), as the corresponding local Hamiltonian problem is known to be QMA-complete [CM14].

Thus a natural starting point consists in replacing the Boolean function $f : \{0,1\}^n \to \{0,1\}$ by a function $\hat{F} : \{0,1\}^n \to \mathcal{P}^{\otimes n}$. Linearity would then amount to the requirement that $\hat{F}(a+b) = \hat{F}(a)\hat{F}(b)$ for every $a, b \in \{0,1\}^n$, where the underlying operation is matrix multiplication. Can such a property be tested?

Our quantum linearity test provides a positive answer for a specific instantiation of the question that is the most appropriate to our setting. The test considers *pairs* of functions $\hat{X}, \hat{Z} : \{0,1\}^n \to \mathcal{O}$, where $\mathcal{O}$ denotes the set of all observables (Hermitian matrices that square to identity), *of any dimension*: indeed, we do not impose a priori that the functions take values in the set $\mathcal{P}^{\otimes n}$ of n-qubit Pauli operators. The goal of the test is to enforce, to the extent possible, that $\hat{X}(a) \simeq X^{a_1} \otimes \cdots \otimes X^{a_n}$ and $\hat{Z}(b) \simeq Z^{b_1} \otimes \cdots \otimes Z^{b_n}$ for all $a, b \in \{0,1\}^n$, where $\simeq$ denotes "behaves as", in a sense soon to be made precise.

We design a test for this property. The test can be implemented by a verifier interacting classically with r entangled provers, where r is a parameter of the test related to the use of an r-qubit quantum stabilizer code; we can take $r = 7$. Each of the provers is sent a query of the form (X, a) or (Z, b), where X, Z are treated as formal labels and $a, b \in \{0,1\}^n$, and replies with a single ± 1 bit.[2] We show that the test has the following completeness and soundness properties (see Lemma 8 for a formal statement):

(*Completeness*) If the provers each own one of the r shares of each of the n qubits of an arbitrary n-qubit state $|\psi\rangle$, encoded via an r-qubit stabilizer code specified in the protocol, and measure their respective share using observable $X(a) = X^{a_1} \otimes \cdots \otimes X^{a_n}$ (resp. $Z(b) = Z^{b_1} \otimes \cdots \otimes Z^{b_n}$) when queried with (X, a) (resp. (Z, b)) then they succeed in the test with probability ω^*_{encode} (a universal constant).

(*Soundness*) For any $\epsilon > 0$ and strategy of the provers which succeeds with probability at least $\omega^*_{encode} - \epsilon$ in the test, for each prover there exists $\mathcal{X}, \mathcal{Z} : \{0,1\}^n \to \mathcal{O}$, where $\mathcal{O}$ is the set of observables acting on the prover's Hilbert space, such that

(i) $\mathcal{X}$ and $\mathcal{Z}$ are *exactly linear*, in the sense that $\mathcal{X}(a+b) = \mathcal{X}(a)\mathcal{X}(b)$ for all $a, b \in \{0,1\}^n$ and similarly for $\mathcal{Z}$;

(ii) The action of $\mathcal{X}$ and $\mathcal{Z}$ on the provers' shared entangled state is indistinguishable, up to an additive $O(\epsilon^{1/16})$ in Euclidean norm, from the action of the provers' actual observables $\hat{X}, \hat{Z}$ in their strategy;

(iii) There exists a local isometry acting on each provers' local Hilbert space such that the action of $\mathcal{X}(a)$, $\mathcal{Z}(b)$ on the isometry's input (the provers' shared entangled state in the strategy) is (up to $O(\epsilon^{1/16})$ in Euclidean norm) equivalent to the action of the Pauli operators $X(a), Z(b)$ on the isometry's ouptut.

A key feature of the test, essential to obtaining a gap-preserving reduction, is that n does not appear in the soundness bound — in the language of testing, the test has *constant completeness-soundness gap*. (In addition the test enforces a number of useful consistency properties between the provers' operations that are used to prove our main theorem; we refer to Section 4 for details.) The actual test (see the three *encoding tests* in Figure 1 for a description) and its analysis combine two main ingredients.

The first ingredient is the entangled-prover linearity test of [IV12], adapted into a two-prover test using the oracularization technique from [IKM09]. This test is used to verify that the provers'

[2] In fact the test considers an two additional types of queries, with two-bit answers; see Section 3 for details.

observables associated with X or Z-type queries, represented by functions $\hat{X}$ and $\hat{Z}$, are close to linear *when considered separately*. The analysis of the test follows the arguments from [IV12, IKM09]. We make the important observation that property (i) of *exact linearity* described above can be guaranteed to hold at the level of the observables themselves for a pair of functions $\mathcal{X}$, $\mathcal{Z}$ that are close to the provers' $\hat{X}$ and $\hat{Z}$. The property, although already implicit in [IV12, IKM09], did not play a significant role in their analysis. For us it is crucial, as it allows certain relations to hold with *zero error*, at the level of operators rather than being state-dependent (for an example where the property is needed, see the beginning of the proof of Lemma 16).

For our purposes the linearity obtained from the entangled-prover linearity test alone is far from sufficient, and indeed there is nothing "quantum" about it: provers not sharing any entanglement, and replying deterministically, can of course succeed in the test as long as their answers are given according to a pair of linear functions.

Therefore a second ingredient is used to turn the classical linearity test above into a "quantum" linearity. This makes use of the uniquely quantum "complementarity" between X and Z Pauli observables: the two operators anti-commute. This is exploited by the famous *CHSH test*, a test that can only be successfully passed by provers sharing entanglement. Here we follow closely the work of [Ji15] and combine the CHSH test with a "stabilizer test" that relies on certain properties of the code used to distribute qubits between the provers (of course malicious provers need not a priori be using this encoding at all — it is a consequence of the test that they should). This part of the test could not be passed by provers who do not share any entanglement: there is no classical randomized strategy that enables the provers to simultaneously sample from the output distribution generated by the application of $\hat{X}(a)$ and $\hat{Z}(b)$ on a code state, as required by the protocol.[3]

1.3 An exponential quantum PCP

We explain the role played by the quantum linearity test in the proof of Theorem 1. Our starting point are the results [FV15, Ji15], which provide a "multiplayer games" analog of the Cook-Levin theorem for QMA-complete problems (in contrast to Kitaev's theorem, which provides the "constraint satisfaction" analog). In particular, Ji [Ji15] gives a five-prover one-round classical interactive proof system for the local Hamiltonian problem such that the verifier's maximum acceptance probability is $1 - K\lambda_{min}(H)n^{-\kappa}$ for constants K, κ (see [Ji15, Theorem 23] for a precise statement). This is sufficient to establish an inverse-polynomial soundness-completeness gap for instances whose minimum eigenvalue are separated by an inverse polynomial.

Moving forward, the most direct route to a proof of the "multiplayer games" variant of the quantum PCP conjecture, i.e. an interactive proof system for the local Hamiltonian problem with *constant* completeness-soundness gap, faces at least two substantial difficulties. First, one should establish a *gap-preserving* reduction, whereby the verifier's maximum acceptance probability is related to $\lambda_{min}(H)$ up to constant, instead of inverse polynomial, factors. Second, it appears like one would still need to establish a "constraint satisfaction" variant of the quantum PCP conjecture in order to allow the gap-preserving reduction to start from a family of instances of the local Hamiltonian problem for which constant approximations to the minimum energy are QMA-hard. Strictly speaking however this step may not be needed, as the transformation provided in the first

[3] It is an interesting question whether the test can be used to verify a high degree of entanglement between the provers; see Section 1.5 for further discussion.

step may ultimately be *gap-introducing*, as is the case in e.g. Dinur's proof of the classical PCP theorem.

Our quantum linearity test provides a step towards resolving the first obstacle: we use it to derive a gap-preserving reduction, albeit at the cost of an exponential blow-up in question length. Indeed, Theorem 1 relies on a protocol with questions of length polynomial in the number of qubits, or variables, of the local Hamiltonian instance, where ideally the *number* of questions would be polynomial (and their length logarithmic).[4] Taking stock of this loss, however, allows us to provide a simple solution to the second obstacle, thereby providing an unconditional constant-gap interactive proof system for the local Hamiltonian problem. Our method is straightforward: we expand an initial inverse-polynomial promise gap on the smallest eigenvalue by taking appropriate tensor powers of the Hamiltonian.[5] While this destroys the locality of the Hamiltonian, it is not an issue for us as our protocol is able to handle any Hamiltonian that is a linear combination of terms made of the tensor product of an arbitrary number of X and Z Pauli operators.[6]

In order to connect the provers' answers, when asked to perform measurements involving both Pauli X and Z operators, to the conclusion of the quantum linearity test, which characterizes their operations on queries that involve only X or Z operators separately, we introduce an additional *consistency test*. In this test one of the provers is asked to measure a term involving both type of Paulis, while the others are asked for a single type. The results are then checked for consistency, providing the desired consistency between different types of queries.

An interesting feature of our protocol, already implicit in [Ji15], arises from the use of the stabilizer encoding that underlies the protocol. As already mentioned the r honest provers should distribute an arbitrary state between themselves by encoding it one qubit at a time using an r-qubit stabilizer code. This allows them to operate jointly on the encoded state by applying any *transversal gate*, which are logical gates that can be implemented by performing a local operation on each of the encoded qubits. Pauli X and Z unitaries are transversal gates for the code we employ, but more complicated codes can support other types of gates. Thus using entanglement between the provers the verifier is able to "orchestrate" certain operations on the entangled state without any prover knowing what gate is being performed (as two different logical gates may in general be implemented transversally with a subset of the provers still performing the same operation in both cases). In our analysis the property is leveraged by treating the r "physical" provers as two "logical" provers, each consisting of a subset of the physical provers. This effectively lets us formulate, and analyze, most of the protocol as being performed with two provers only; see Section 3 for more details.

1.4 Related work

We build on a number of previous works in quantum information and complexity theory.

[4] It is this exponential blow-up which makes Theorem 1 follow from results in the literature on classical and quantum interactive proofs, as already mentioned. See Section 1.5 below for further discussion.

[5] We emphasize that the main contribution of our work, and the respect in which it provides a *gap-preserving* reduction, is *not* due to this simple *gap-amplification* trick. Rather, it is in our quantum linearity test and accompanying interactive proof system which provide a gap-preserving reduction irrespective of whether amplification has been performed or not.

[6] In order for the verifier to remain polynomial-time we do need to be able to sample from the distribution implied by the modulus of the corresponding coefficients, which in general precludes starting from an arbitrary Hamiltonian and expanding it in the Pauli basis. This is not a problem for the amplification procedure described here.

First we mention that motivation for the problem we consider goes back to a question of Aharonov and Ben-Or (personal communication, 2013), who asked how a quantum generalization of the exponential classical PCP could look like if it was not derived through the "circuitous route" obtained as the compilation of known but complex results from the theory of classical and quantum interactive proof systems (as described earlier). In this respect we point to [AAV13, Section 5] for a very different approach to the same question based on a "quantum take" on the arithmetization technique.

More directly, our work builds on the already-mentioned works [FV15, Ji15] initiating the study of entangled-prover interactive proof systems for the local Hamiltonian problem. The idea of using a distributed encoding of the ground state in order to obtain a multiprover interactive proof system for the ground state energy is introduced in [FV15]. In that work the protocol required the provers to return qubits; the possibility for making the protocol purely classical was uncovered by Ji [Ji15]. Our use of stabilizer codes, and the stabilizer test which forms part of our protocol, originate in his work. In addition we build upon ideas introduced in the study of quantum multiprover interactive proofs with entangled provers [KM03, CHTW04], and especially the three-prover linearity test of [IV12] and the use of oracularization from [IKM09] to make it into a two-prover test. Finally we draw upon important results from the quantum self-testing literature; in particular, self tests for the graph states [McK13].

Compared to the works mentioned above, and [FV15, Ji15] in particular, our result differs in two important respects, making it incomparable in general. First, the question size in our protocol is much larger: $\mathrm{poly}(n)$ bits instead of $O(\log n)$ for [Ji15]. Second, the dependence of the verifier's acceptance probability on the ground state energy is much better: while our dependence is of a constant factor, in [Ji15] there is a polynomial scaling.[7] Interpreting all three results as steps towards a quantum PCP theorem, [FV15, Ji15] propose a first step that is *size-preserving* (the number of questions is polynomial in the instance size) but has only an inverse polynomial gap; in contrast we take the route of a *gap-preserving* construction, but the number of questions becomes exponential in the instance size.

Our results are also related to work in quantum property testing [MdW13], and in particular testing EPR pairs [MYS12] and more general entangled states [McK14]. In this setting state-of-the-art results [RUV13] essentially show how the presence of n EPR pairs between two provers can be certified via a protocol using queries and answers of length polynomial in n, with inverse-polynomial completeness-soundness gap. Thus here again no "constant-gap" results are known, where the gap would remain constant as the number of EPR pairs tested grows. Our work is incomparable: our protocol has constant gap but *does not* by itself suffice to certify that the provers share a large number of EPR pairs that are in *tensor product* form.

Very recently and independently of our work, McKague [McK15] has proposed protocols for self-testing many-qubit states that achieve a guarantee similar to ours, i.e. the protocol certifies that there exists an isometry acting on the provers' state for which the expectation values of Pauli operators on the output are close to the expectation values of the provers' measurements (see e.g. Lemma 17). However, his protocol is not directly comparable to ours since it requires $\mathrm{poly}(n)$-bit answers and $\log(n)$-bit questions, whereas we use $O(1)$-bit answers and $\mathrm{poly}(n)$-bit questions;

[7] One could attempt to recover our result by repeating the protocol in [Ji15] a polynomial number of times. Provided there existed an appropriate parallel repetition theorem, this would amplify the soundness to a constant. However, the answer length would now be polynomial, and it is unclear whether this could be reduced to a constant without having to go once more through the complicated reductions of [Vid13], defeating the purpose.

in addition and more importantly for us the completeness-soundness gap in his protocol scales polynomially with n, whereas in our case the scaling is independent of n.

1.5 Directions for future work

Improving the question length from polynomial in n (in fact, linear in n: the polynomial dependency only enters our result through the amplification procedure, but is not needed for the gap-preserving reduction itself) to linear in $\log n$ would give a proof of the "multiplayer games" variant of the quantum PCP conjecture stated in [FV15]. We expect this to present a significant challenge (note that it would recover, and strengthen, the inclusion NEXP $\subseteq$ MIP*), but it forms the motivation behind our work. In the classical setting, the key ingredient in the proof of NEXP $\subseteq$ MIP consists in replacing the linearity test with a test for *multilinear* functions, or more generally for low-degree multivariate polynomials. There has also been recent work in the context of direct-sum testing, which directly achieves a linearity test with reduced query length [DDG$^+$14]. It is an interesting open question whether these tests can be generalized to the quantum setting, extending our quantum linearity test.

More generally, the area of *device-independent* quantum property testing has many interesting open problems [MdW13], which however almost systematically suffer from inverse-polynomial completeness-soundness gaps as soon as the property tested scales in size. Our results may suggest novel approaches to some of these problems.

Organization of the paper. In Section 2 we introduce some notation used throughout as well as basic definitions on stabilizer codes and local Hamiltonians. In Section 3 we describe the protocol used for the proof of Theorem 1. In Section 4 we analyze the quantum linearity test performed as part of the protocol as a stand-alone test. In Section 5 we conclude the analysis of the protocol, leading to the proof of Theorem 1.

2 Preliminaries

We assume basic familiarity with quantum information but give all required definitions. We refer to the standard textbook [NC01] for additional background material.

2.1 Quantum states and measurements

A n-qubit quantum state is represented by a unit vector $|\psi\rangle \in \mathbb{C}^2 \otimes \cdots \otimes \mathbb{C}^2 = (\mathbb{C}^2)^{\otimes n} \approx \mathbb{C}^{2^n}$, where the ket notation $|\cdot\rangle$ is used to signify a column vector. A bra $\langle\psi|$ is used for the conjugate-transpose $\langle\psi| = |\psi\rangle^\dagger$, which is a row vector. We use $\||\psi\rangle\|^2 = |\langle\psi|\psi\rangle|$ to denote the Euclidean norm, where $\langle\psi|\phi\rangle$ is the skew-Hermitian inner product between vectors $|\phi\rangle$ and $|\psi\rangle$. For a matrix X, $\|X\|$ will refer to the operator norm, the largest singular value. When the Hilbert space can be decomposed as $\mathcal{H} = \mathcal{H}_A \otimes \mathcal{H}_B$ for some $\mathcal{H}_A$ and $\mathcal{H}_B$, and X is an operator on $\mathcal{H}_A$, we often write X as well for the operator $X \otimes \mathbb{I}_{\mathcal{H}_B}$ on $\mathcal{H}$. It will always be clear from context which space an operator acts on.

A density matrix on n qubits is a positive semi-definite matrix $\rho \in \mathbb{C}^{2^n} \times \mathbb{C}^{2^n}$ of trace 1. The density matrix associated to $|\psi\rangle$ is the rank-1 projection $|\psi\rangle\langle\psi|$.

A n-qubit measurement (also called POVM, for projective operator-valued measurement) with k outcomes is specified by k positive matrices $M = \{M_1, \ldots, M_k\}$ in $\mathbb{C}^{2^n} \times \mathbb{C}^{2^n}$ such that $\sum_i M_i = \mathbb{I}$.

	1	2	3	4	5	6	7
Stabilizers	I	I	I	X	X	X	X
	I	X	X	I	I	X	X
	X	I	X	I	X	I	X
	I	I	I	Z	Z	Z	Z
	I	Z	Z	I	I	Z	Z
	Z	I	Z	I	Z	I	Z
Logical X	X	X	X	X	X	X	X
Logical Z	Z	Z	Z	Z	Z	Z	Z

Table 1: Stabilizer table for the 7-qubit Steane code

The measurement is *projective* if each M_i is a projector, i.e. $M_i^2 = M_i$. The probability of obtaining the i-th outcome when measuring state ρ with M is $\mathrm{Tr}(M_i\rho)$. By Naimark's dilation theorem, any POVM can be simulated by a projective measurement acting on an enlarged state; that is, for every POVM $M = \{M_i\}_i$ acting on state $|\psi\rangle \in \mathcal{H}$ there exists a projective measurement $M' = \{P_i\}_i$ and a state $|\psi\rangle \otimes |\phi\rangle \in \mathcal{H} \otimes \mathcal{H}_{\text{ancilla}}$ with the same outcome probabilities as M. Moreover, the post-measurement state after performing M is the same as the *reduced* post-measurement state obtained after performing M' and tracing out the ancilla subsystem $\mathcal{H}_{\text{ancilla}}$.

An n-qubit observable is a Hermitian matrix $O \in \mathbb{C}^{2^n} \times \mathbb{C}^{2^n}$ that squares to identity. O is diagonalizable with eigenvalues ± 1, $O = P_+ - P_-$, and $P = \{P_+, P_-\}$ is a projective measurement. For any state ρ, $\mathrm{Tr}(O\rho)$ is the expectation of the ± 1 outcome obtained when measuring ρ with P. If $\rho = |\psi\rangle\langle\psi|$ we abbreviate this quantity, $\mathrm{Tr}(O\rho) = \mathrm{Tr}(P_+\rho) - \mathrm{Tr}(P_-\rho) = \langle\psi|O|\psi\rangle$ as $\langle P\rangle_\psi$.

A convenient orthogonal basis for the real vector space of n-qubit observables is given by the set $\{I, X, Y, Z\}^{\otimes n}$, where $\{I, X, Y, Z\}$ are the four single-qubit Pauli observables

$$I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \tag{1}$$

We often consider operators that are tensor products of just I and X, or just I and Z. We denote these by $X(a), Z(b)$, where the strings $a, b \in \{0,1\}^n$ indicate which qubits to apply the X or Z operators to: a 0 in position i indicates an I on qubit i, and a 1 indicates an X or Z.

2.2 Stabilizer codes

Stabilizer codes are the quantum analogue of linear codes. For an introduction to the theory of stabilizer codes we refer to [Got97]. We will only use very elementary properties of such codes.

The codes we consider are *Calderbank-Shor-Steane (CSS) codes* [CS96, Ste96]. For an r-qubit code the codespace, the vector space of all valid codewords, is the subspace of $(\mathbb{C}^2)^{\otimes r}$ that is the simultaneous $+1$ eigenspace of a set $\{S_1, \ldots, S_k\}$ of r-qubit pairwise commuting Pauli observables called the stabilizers of the code. The stabilizers form a group under multiplication. Unitary operations, such as a Pauli X or Z operators, on the logical qubit are implemented on the codespace by logical operators $X_{logical}$ and $Z_{logical}$. The smallest CSS code is Steane's 7-qubit code [Ste96]. Table 1 lists a set of stabilizers that generate the stabilizer group of the code.

Every CSS code satisfies certain properties which will be useful for us. Firstly, both the stabilizer generators and the logical operators can be written as tensor products of only I, X, and Z

operators — there are no Y. This simplifies our protocol, allowing us to consider only two distinct basis settings. Secondly, every CSS code has the following symmetry: for every index $i \in [r]$ there exists stabilizers S_X, S_Z such that S_X is a tensor product of only X and I operators and has an X at position i, and S_Z is equal to S_X with all X operators replaced by Z operators.

These properties imply the following simple observation, which will be important for us. For every Pauli operator $P \in \{I, X, Z\}$ acting on the i-th qubit of the code there is a tensor product $\bar{P}$ of Paulis acting on the remaining $(r-1)$ qubits such that $P \otimes \bar{P}$ is a stabilizer operator on the whole state, and moreover each term in the tensor product is either identity or P. Indeed, the choice of $\bar{P}$ is not unique. Henceforth, we use the notion $\bar{P}$ to denote *any* such operator, unless otherwise specified.

2.3 Local Hamiltonians

A n-qubit local Hamiltonian is a Hermitian, positive semidefinite operator H on $(\mathbb{C}^2)^{\otimes n}$ that can be decomposed as a sum $H = \sum_{i=1}^{m} H_i$ with each H_i is local, i.e. H_i can be written as $H_i = I \otimes \cdots I \otimes h_i \otimes I \otimes \cdots \otimes I$, where h_i is a Hermitian operator on $(\mathbb{C}^2)^{\otimes k}$ with norm (largest singular value) at most 1. The smallest k for which H admits such a decomposition is called the locality of H. The terms are normalized such that $\|H_i\| \leq 1$ for all i. A family of Hamiltonians $\{H_i\}$ acting on increasing numbers of qubits is called local if all H_i are k-local for some k independent of n (for us k will always be 2).

The local Hamiltonian problem is the prototypical QMA-complete problem, as 3SAT is for NP.

Definition 2. *Let $k \geq 2$ be an integer. The* k-local Hamiltonian problem *is to decide, given a family of k-local Hamiltonians $\{H_n\}_{n\in\mathbb{N}}$ such that H_n acts on n qubits, and functions $a, b : \mathbb{N} \to (0, 1)$ such that $b - a = \Omega(\mathrm{poly}^{-1}(n))$, if the smallest eigenvalue of H_n is less than $a(n)$ or greater than $b(n)$.*

Here we restrict our attention to Hamiltonians

$$H = \frac{1}{m} \sum_{i=1}^{m} H_i,$$

for which each term H_i can be written as a linear combination of tensor products of Pauli X and Z observables only (no Y). Such Hamiltonians are known to be QMA complete; in particular we consider a restricted class of 2-local Hamiltonians, called Hamiltonians of XZ form, for which each H_i can be written as $H_i = \alpha_{i_1 i_2}(X_{i_1} \otimes X_{i_2} + Z_{i_1} \otimes Z_{i_2})$, where $i_1, i_2 \in \{1, \ldots, n\}$ indicate the qubits on which a Pauli X or Z acts and the coefficients $\alpha_{i_1 i_2} \in \mathbb{R}$ satisfy $|\alpha_{i_1 i_2}| \leq 1$.

Theorem 3 (Cubitt and Montanaro [CM14], lemma 21)**.** *The local Hamiltonian problem for XZ-Hamiltonians is QMA-complete.*[8]

2.4 State-dependent distance measure

We make extensive use of a state-dependent distance between measurements that has been frequently used in the context of entangled-prover interactive proof systems (see e.g. [IV12, Ji15]).

[8] In [CM14, Lemma 21] this is stated for the XY Hamiltonian, to which XZ is equivalent by local rotation.

Let $\{M^a\}$ and $\{N^a\}$ be two POVMs with the same number of possible outcomes, indexed by a, and let $|\psi\rangle$ be a quantum state. The *state-dependent distance* between M and N on $|\psi\rangle$ is defined as

$$d_\psi(M,N) = \Big(\sum_a \big\|\sqrt{M^a}|\psi\rangle - \sqrt{N^a}|\psi\rangle\big\|^2\Big)^{1/2}.$$

To simplify the notation, let $A^a = \sqrt{M^a}$ and $B^a = \sqrt{N^a}$. Then this distance can be rewritten as:

$$\begin{aligned} d_\psi(M,N)^2 &= \sum_a \|A^a|\psi\rangle - B^a|\psi\rangle\|^2 \\ &= \sum_a \langle\psi|(A^a - B^a)^2|\psi\rangle \\ &= \sum_a \big(\|A^a|\psi\rangle\|^2 + \|B^a|\psi\rangle\|^2 - \langle\psi|(A^aB^a + B^aA^a)|\psi\rangle\big) \\ &= 2 - \sum_a \langle\psi|(A^aB^a + B^aA^a)|\psi\rangle \\ &= 2 - 2\sum_a \mathrm{Re}\big(\langle\psi|A^aB^a|\psi\rangle\big). \end{aligned}$$

In the last line we have used the fact that A^a and B^a are Hermitian. If we specialize to the case of projective measurements with binary outcomes, we get the following relations (here $A = A^0 - A^1$ and $B = B^0 - B^1$ are the observables associated to the measurements):

$$\begin{aligned} d_\psi(M,N)^2 &= 2 - \langle\psi|(A^0B^0 + A^1B^1 + B^0A^0 + B^1A^1)|\psi\rangle \\ &= 2 - \frac{1}{4}\langle\psi|((\mathbb{I}+A)(\mathbb{I}+B) + (\mathbb{I}-A)(\mathbb{I}-B) + (\mathbb{I}+B)(\mathbb{I}+A) + (\mathbb{I}-B)(\mathbb{I}-A))|\psi\rangle \\ &= 2 - \frac{1}{4}\langle\psi|(4\mathbb{I}+2AB+2BA)|\psi\rangle \\ &= 1 - \frac{1}{2}\langle\psi|(AB+BA)|\psi\rangle \\ &= \frac{1}{2}\langle\psi|(A-B)^2|\psi\rangle. \end{aligned} \tag{2}$$

This distance measure has the following useful property:

Lemma 4. *Let $|\psi\rangle$ be a quantum state, $\{C_a\}$ a family of operators such that $\|\sum_a C_aC_a^\dagger\| \le K$ and $\{M^a\}$ and $\{N^a\}$ POVMs. Then*

$$\Big|\sum_a \langle\psi|C_a\sqrt{M^a}|\psi\rangle - \sum_a \langle\psi|C_a\sqrt{N^a}|\psi\rangle\Big| \le \sqrt{K}\,d_\psi(M,N).$$

Proof. Let $A^a = \sqrt{M^a}$ and $B^a = \sqrt{N^a}$. Applying the Cauchy-Schwarz inequality,

$$\begin{aligned} \Big|\sum_a \langle\psi|C_a(A^a - B^a)|\psi\rangle\Big| &\le \Big|\langle\psi|\sum_a C_aC_a^\dagger|\psi\rangle\Big|^{1/2}\Big|\langle\psi|\sum_a (A^a-B^a)^2|\psi\rangle\Big|^{1/2} \\ &\le \sqrt{K}\,d_\psi(M,N), \end{aligned}$$

as claimed. □

A second measure of proximity that is often convenient is the *consistency*. As before, let $\{M^a\}$ and $\{N^a\}$ be POVMs with the same number of outcomes. Then their consistency is defined as

$$\mathrm{CON}_\psi(M,N) = \mathrm{Re}\Big(\sum_a \langle\psi|M^a N^a|\psi\rangle\Big).$$

The following lemma relates the consistency and the state-dependent distance.

Lemma 5 (Lemma 10 in [Ji15])**.** *Let ψ be a state and $\{M^a\}$ and $\{N^a\}$ be two POVMs with equal numbers of outcomes. If $\mathrm{CON}_\psi(M,N) = 1-\delta$ for some $\delta \geq 0$ then $d_\psi(M,N) \leq \sqrt{\delta}$.*

A useful property of the consistency is that if M and N are POVMs acting on two separate subsystems of $|\psi\rangle$, applying Naimark dilation to each of them results in projective measurements M' and N' and a state $|\psi'\rangle$ such that $\mathrm{CON}_\psi(M,N) = \mathrm{CON}_{\psi'}(M',N')$.

3 Description of the protocol

In this section we describe the protocol used in the proof of Theorem 1. The input to the protocol is an n-qubit local Hamiltonian H in XZ form, as described in Section 2.3. The verifier interacts with r provers. One should think of "honest" provers as sharing a qubit-by-qubit encoding of the n-qubit ground state $|\Gamma\rangle$ of H according to a CSS code, such as Steane's 7-qubit code in which case $r = 7$, as described in Section 2.2, and of performing the Pauli measurement indicated by the query (a complete description of the honest strategy is given in Definition 6).

In the protocol, the verifier asks the provers to perform one of a series of tests chosen according to some pre-specified distribution. In each test the prover must answer with one or two answer bits, which are encoded as ± 1 to match the convention that quantum observables have ± 1 eigenvalues. There are several possible types of queries that each prover may receive:

1. An *X-query*, represented by (X, a, b), where a, b are uniformly random strings in $\{0,1\}^n$.[9] The expected answer is two bits $\alpha, \beta \in \{-1, 1\}$.

2. A *Z-query*, represented by (Z, a, b). Same as an X-query, except with Z instead of X.

3. An *XZ-query*, represented by (X, a, Z, b) where a and b are arbitrary binary strings such that $a \wedge b = 0^n$. This type of query is used only in the energy test, and the distribution on a and b depends on the Hamiltonian. The expected answer is two bits $\alpha, \beta \in \{-1, 1\}$.

4. A *W-query*, represented by (N, a, b) where $N \in \{X', Z'\}$ and a, b are uniformly random strings in $\{0,1\}^n$. The expected answer is a single bit $\alpha \in \{-1, 1\}$.

To each query is associated an intended behavior of the prover, which is specified as part of the *honest strategy* given in the following definition.

Definition 6. *The* honest strategy *for the r provers consists of the following. The provers share an (rn)-qubit state $|\psi\rangle$ which is obtained as the qubit-by-qubit encoding, using a CSS code as described in Section 2.2, of the ground state $|\Gamma\rangle$ of the local Hamiltonian H. Each prover holds one qubit from the encoding of each qubit of $|\Gamma\rangle$.*

Upon receiving a query, any prover performs the following depending on the type of the query:

[9] We will always assume the strings a, b are sent to the prover in lexicographic order.

- *X-query (X, a, b): measure the compatible observables $X(a)$ and $X(b)$ on its share of the encoded state, and return the two outcomes.*
- *Z-query (Z, a, b): same as X-query but with observables $Z(a)$ and $Z(b)$.*
- *XZ-query (X, a, Z, b): measure the compatible observables $X(a)$ and $Z(b)$, and return the two outcomes.*
- *W-query (N, a, b): measure the observable $(X(a) + Z(b))/\sqrt{2}$ if $N = X'$, $(X(a) - Z(b))/\sqrt{2}$ if $N = Z'$, and return the outcome.*

The protocol is to be performed with r "physical" provers, but we formulate all but one of the tests (the energy test) as a two-prover test. In this case we call the two provers "logical" provers. A query to the two logical provers can be mapped to a query to the r physical provers as follows. One of the physical provers is chosen at random to play the role of the first logical prover, called the *special prover*. The remaining $(r-1)$ physical provers together play the role of the second logical prover, called the *composite prover*.[10] For a given query Q to the special prover of a type among those specified above we define a *complementary query* $\overline{Q}$ for the composite prover as per the following lemma.

Lemma 7. *For any X-query or Z-query, there exists a complementary query $\overline{Q}$ such that*

1. *The query associated to each physical prover forming the composite prover in $\overline{Q}$ is of the same type as Q. In particular the distribution on query strings is as specified by the query type.*
2. *If all provers apply the honest strategy and provide answers α, β to Q and $\overline{\alpha}, \overline{\beta}$ to $\overline{Q}$ respectively, where $\overline{\alpha}$ and $\overline{\beta}$ are each obtained as the product of the answer to the corresponding query coming from each of the physical provers making up the composite prover, it holds that $\alpha\overline{\alpha} = \beta\overline{\beta} = +1$.*

Proof. Both items follow from the properties of CSS codes described in Section 2.2. We give the proof for an X query (X, a, b). Let the index of the special prover be i, and let S_X be a stabilizer of the code, such that S_X consists only of X and I Paulis and has an X in position i. For each physical prover $j \neq i$ associated with the composite prover, if the operator in position j of S_X is X, prover j is sent the query (X, a, b). Otherwise, prover j is sent a uniformly random X-query (X, c, d).

Composite answers $\overline{\alpha}, \overline{\beta}$ to the complementary query are determined by taking the product of the answers from all provers who did not receive random strings; using that S_X is a stabilizer of the code ensures that item 2 is satisfied.

In the composite query, for a given choice of S_X each prover receives a query that is either identical to the original query, or is a uniformly random string; since the original query is chosen at random this is also the case for each of the physical provers associated with the composite prover. This proves item 1. □

[10] The physical provers remain isolated throughout the protocol and are never allowed to communicate; it is only for purposes of analysis that we group $(r-1)$ physical provers into a single logical prover. In particular the physical provers are never told which logical prover they are associated with, and the distribution of queries to any physical prover is the same whether it plays the role of the special or composite prover.

The complete protocol is described in Figure 1. It is based on four tests. In the *energy test*, the verifier asks the provers to measure a randomly chosen term in the Hamiltonian. This test is described in more detail in Section 5.1. The remaining three tests are called the *encoding tests*; together these tests form our quantum linearity test. The *two-query linearity test*, a variant of the classical linearity test of Blum, Luby and Rubinfeld, is designed to show the existence of exactly linear (in a sense to be made precise in Section 4.1) observables $\mathcal{X}$ and $\mathcal{Z}$ that are close to the provers' actual measurement operators. In the *stabilizer test*, the provers are asked to measure a random generator of the stabilizer group associated with the code, and the verifier accepts their answers if their product is $+1$. In the *anticommutation test*, a variant of the CHSH game is played between the verifier and the two logical provers.

Given a local Hamiltonian H in XZ form, the verifier performs the following one-round interaction with r provers. The probability $p \in (0,1)$ is a parameter of the protocol that can be specified freely.

- Choose one of the r provers uniformly at random to be the *special prover*. The other provers form the *composite prover*.
- With probability p, perform the *energy test* described in Section 5.1.
- With probability $(1-p)/3$ each, perform one of the following three *encoding tests*:
  1. **Linearity test:** The verifier chooses a basis setting $N \in \{X, Z\}$ and strings $a, b, c \in \{0,1\}^n$ uniformly at random. He sends the special prover (N, a, b) and the composite prover either $\overline{(N, a, c)}$, $\overline{(N, b, c)}$, or $\overline{(N, a+b, c)}$, each with probability $1/3$.
     The verifier accepts if answers associated with the same query string match, and the product of the answers associated to a, b and $a+b$ is $+1$.
  2. **Anticommutation test:** The verifier chooses basis settings $N \in \{X, Z\}$ and $N' \in \{X', Z'\}$ and strings $a, b, c \in \{0,1\}^n$ uniformly at random. He sends the special prover (N', a, b) and the composite prover $\overline{(X, a, c)}$ if $N = X$, and $\overline{(Z, b, c)}$ if $N = Z$.
     The verifier ignores the second answer bit from each prover, and accepts or rejects according to the following rule: if the inner product $a \cdot b = 0 \mod 2$ (i.e. the bit-wise AND $a \wedge b$ has even Hamming weight), then he automatically accepts; otherwise, if the two basis settings were Z' and Z the verifier accepts if the product of the provers' answers is -1; otherwise, he accepts if the product is $+1$.
  3. **Stabilizer test:** The verifier chooses a basis setting $N \in \{X, Z\}$ and three strings $a, b, c \in \{0,1\}^n$ uniformly at random. He sends the special prover (N, a, b) and the composite prover $\overline{(N, a, c)}$. The verifier accepts if the product of the answers associated to the query string a is $+1$.

Figure 1: The protocol

4 The quantum linearity test

In this section we analyze the part of the protocol described in Section 3 that consists of the three encoding tests. Note that these tests do not depend on the local Hamiltonian. They can thus be considered as an independent game to be played with the r provers, where each test is chosen with equal probability.

The following lemma states the main result of this section. (To understand the notation used in the lemma it may be useful to first read the ensuing paragraphs on modeling arbitrary strategies for the provers in the protocol.)

Lemma 8. *Assume a strategy $(N, |\psi\rangle)$ for the provers succeeds in the encoding tests with probability at least $\omega^*_{encode} - \epsilon$, where ω^*_{encode} is the success probability of the strategy described in Definition 6.*

Then for any prover $i \in \{1, \ldots, r\}$ and $a \in \{0,1\}^n$ there exists observables $\mathcal{X}(a)$ and $\mathcal{Z}(a)$ acting on the i-th prover's register[11] *such that*

$$\forall a, b \in \{0,1\}^n \qquad \mathcal{X}(a)\mathcal{X}(b) = \mathcal{X}(a+b) \qquad \text{and} \qquad \mathcal{Z}(a)\mathcal{Z}(b) = \mathcal{Z}(a+b),$$

and

$$\frac{1}{2^{2n}} \sum_{a,b \in \{0,1\}^n} \big\| (\mathcal{X}(a)\mathcal{Z}(b) - (-1)^{a\cdot b}\mathcal{Z}(b)\mathcal{X}(a))|\psi\rangle \big\|^2 \;=\; O\big(\epsilon^{1/4}\big). \tag{3}$$

Moreover, if $\hat{X}(a)$ (resp. $\hat{Z}(b)$) is the observable that prover i performs when asked an X-query (X, a, b), and the outcome β associated to b is ignored, (resp. Z-query (Z, a, b) with the outcome α associated to a ignored), then

$$\frac{1}{2^n} \sum_{a \in \{0,1\}^n} \big\| (\mathcal{X}(a) - \hat{X}(a))|\psi\rangle \big\|^2 \;=\; O\big(\epsilon^{1/4}\big) \qquad \text{and} \qquad \frac{1}{2^n} \sum_{b \in \{0,1\}^n} \big\| (\mathcal{Z}(b) - \hat{Z}(b))|\psi\rangle \big\|^2 \;=\; O\big(\epsilon^{1/4}\big).$$

We note that the constant ω^*_{encode} is given by

$$\omega^*_{encode} = \frac{2}{3} + \frac{1}{3}\omega^*_{\text{anti-com}}, \tag{4}$$

where $\omega^*_{\text{anti-com}} \in (0,1)$ is specified in the proof of Lemma 12. This is because an honest strategy passes the linearity and stabilizer tests with probability 1, and the anticommutation test with probability $\omega^*_{\text{anti-com}}$.

Before proceeding with the analysis of the encoding tests we introduce some notation associated with arbitrary strategies for the provers in the protocol. We specify a strategy using the shorthand $(N, |\psi\rangle)$. Here $|\psi\rangle$ denotes the r-partite state shared by the provers, and N the collection of POVM that the provers apply in response to the different types of queries they can be asked. Given a query (N, a, b), where $N \in \{X, Z\}$ we denote by $\{N^{\alpha\beta}_{ab}\}_{\alpha,\beta}$ the two-outcome POVM that is applied by a given prover. Although these operators may differ from one prover to another it will usually not be necessary to specify explicitly the index $i \in \{1, \ldots, r\}$ associated with the prover. Instead we will only differentiate between the special prover, whose operators will be denoted $\hat{N}$, and the composite prover, for whom the resulting operator, obtained by taking the tensor product of operators applied by each of the $(r-1)$ associated physical provers, will be denoted $\overline{N}$. These

[11] We allow extending this register by adding ancilla qubits initialized to $|0\rangle$.

operators are local to one and $(r-1)$ provers respectively, but we will usually omit tensor product signs and write them as operators acting on the whole Hilbert space, keeping in mind that an operator of the form $\hat{N}$ always commutes with an operator of the form $\overline{N}$.

By taking appropriate marginals over the answers we can define associated observables for the provers, $\hat{X}(a)$ and $\hat{Z}(b)$ for the special prover and $\overline{X}(a)$ and $\overline{Z}(b)$ for the composite prover, where

$$\hat{X}(a) = \frac{1}{2^n} \sum_{b\in\{0,1\}^n} \sum_{\beta\in\{\pm 1\}} \big(N_{ab}^{1\beta} - N_{ab}^{-1\beta}\big), \qquad \hat{Z}(b) = \frac{1}{2^n} \sum_{a\in\{0,1\}^n} \sum_{\alpha\in\{\pm 1\}} \big(M_{ab}^{\alpha 1} - M_{a,b}^{\alpha -1}\big), \quad (5)$$

for the POVM $\{N_{ab}^{\alpha\beta}\}$ and $\{M_{ab}^{\alpha\beta}\}$ associated to the queries (X,a,b) and (Z,a,b) respectively. Observables $\bar{X}(a)$ and $\bar{Z}(b)$ are defined similarly.

With the notation in place the proof of Lemma 8 follows from the analysis of the encoding tests given in the following subsections.

Proof of Lemma 8. Fix an arbitrary strategy $(N, |\psi\rangle)$ for the provers. For $a, b \in \{0,1\}^n$ let $\mathcal{X}(a)$ and $\mathcal{Z}(b)$ be the observables introduced in Definition 10. When $a \cdot b = 1 \mod 2$ the anticommutation property implied by (3) is proven in Lemma 12. When $a \cdot b = 0 \mod 2$ the corresponding commutation is proved in Lemma 13. Finally the relation between the observables $\mathcal{X}, \mathcal{Z}$ and the provers' original strategy follows from the definition and Lemma 9 analyzing the linearity test. □

4.1 Two-query linearity test

In [IV12] it is shown that the classical 3-query linearity test of Blum, Luby, and Rubinfeld [BLR93] (BLR) is sound against entangled provers. The proof is an adaptation of the Fourier-analytic proof due to Haståd to the matrix-valued setting. Here we analyze a two-query version of the test, again using Fourier analysis to prove soundness. The test is based on the idea of oracularization with a dummy question introduced in [IKM09]. We note that the use of two provers, rather than three as in the original test, is essential for us. This is because our quantum linearity test relies on simultaneously testing for linearity of two functions that are obtained (in the honest case) by applying tensor products of X and Z operators respectively. For the "linearity" part of the test to be compatible with the other subtests, such as the anticommutation test, it is necessary that the special prover and the composite prover share a state that is equivalent to an EPR pair. Monogamy of entanglement thus prevents us from designing a protocol that would enforce both the anticommutation test and a three-prover variant of the classical linearity test; this is a key difference between our *quantum* linearity test and the entangled-prover *classical* linearity test of [IV12].

We describe the test as a test to be performed with two provers. In the actual protocol one of the two provers is the special prover and the other is the composite prover that consists of the combination of $(r-1)$ out of the r "physical" provers. We identify the composite prover with the first prover in the test as described below, and the special prover with the second prover in the test. When the test is performed the role of special prover is assigned uniformly at random among the r possibilities.[12]

The test is specified in Figure 1. For convenience we repeat it here. First the verifier chooses a random basis setting $N \in \{X, Z\}$ that is sent to both provers. The accompanying strings are determined as follows:

[12] As previously mentioned, a prover is never told if it is playing the role of the special prover or of one of the provers making up the composite prover.

1. Choose two strings $a, b \in \{0,1\}^n$ uniformly at random. Send the lexicographically ordered pair $\{a, b\}$ to the first prover.

2. Let c be with equal probability either a, b, or $a + b$, and let c' be a random string. Send the lexicographically ordered pair $\{c, c'\}$ to the second prover.

3. The provers reply with $\alpha, \beta \in \{\pm 1\}$ and $\gamma, \gamma' \in \{\pm 1\}$ respectively. Depending on the value of c the verifier performs one of the following two tests:

   (a) *Consistency test*: if $c = a$ (resp. b), accept if both provers return the same value as their corresponding answer: $\gamma = \alpha$ (resp. $\gamma = \beta$).

   (b) *Linearity test*: if $c = a + b$, accept if $\gamma = \alpha\beta$.

We show the following.

Lemma 9. *Suppose two provers sharing entangled state $|\psi\rangle$ and making measurements $\{M_{ab}^{\alpha\beta}\}_{\alpha,\beta}$, $\{N_{ab}^{\alpha\beta}\}_{\alpha,\beta}$ respectively succeed in the oracularized linearity test with probability $1 - \epsilon$. Then there exists a projective measurement $\{C^u\}_{u\in\{0,1\}^n}$ and $\epsilon_{lin} = O(\sqrt{\epsilon})$ such that*

$$\mathbf{E}_a \, \mathrm{CON}_{\psi''}(\tilde{N}_a, C_a) \; = \; 1 - \epsilon_{lin},$$

where the expectation is taken with respect to the uniform distribution on $a \in \{0,1\}^n$, $|\psi''\rangle$ is $|\psi\rangle$ tensored with local ancilla registers on the second prover's register, $\tilde{N}_a^\alpha = \mathbf{E}_b \sum_\beta N_{ab}^{\alpha\beta}$ and $C_a = \sum_u (-1)^{u\cdot a} C^u$.

Moreover, the honest strategy (see Definition 6) succeeds in the test with probability 1.

Before proceeding with the proof of the lemma we introduce useful notation.

Definition 10 (Exactly linear observables)**.** *Let $(N, |\psi\rangle)$ be a strategy for the provers in the protocol. For any prover i and $a, b \in \{0,1\}^n$ define observables $\mathcal{X}(a)$ and $\mathcal{Z}(b)$ as the observables C_a and C_b from Lemma 9 when the second prover's measurements in the linearity test are prover i's measurements on an X-query and Z-query respectively. We say that $\mathcal{X}$ and $\mathcal{Z}$ are* exactly linear *observables because they automatically satisfy*

$$\forall a, b \in \{0,1\}^n, \qquad \mathcal{X}(a) = \mathcal{X}(b)\mathcal{X}(a+b) \qquad \textit{and} \qquad \mathcal{Z}(b) = \mathcal{Z}(a)\mathcal{Z}(a+b). \tag{6}$$

Proof of Lemma 9. Assume without loss of generality that the provers' POVMs $\{M_{ab}^{\alpha\beta}\}_{\alpha,\beta}$ and $\{N_{ab}^{\alpha\beta}\}_{\alpha,\beta}$ respectively are projective. Note that here the subscript ab should always be understood as a lexicographically ordered pair $\{a, b\}$, and the superscript indicates that α is the answer associated to a and β to b. Consider the marginalized operators obtained by averaging out over the second question:

$$\tilde{M}_a^\alpha = \mathbf{E}_b \sum_\beta M_{ab}^{\alpha\beta}, \qquad \tilde{N}_a^\alpha = \mathbf{E}_b \sum_\beta N_{ab}^{\alpha\beta}.$$

These operators are in general *not* projectors. It will also be useful to consider the following conditional measurement operator, which *is* a projector:

$$M_{a|ab}^\alpha = \sum_\beta M_{ab}^{\alpha\beta}.$$

Suppose that the provers' acceptance probability conditioned on the verifier performing the consistency part of the test (i.e. $c = a$ or $c = b$) is $1 - \epsilon_c$, while conditioned on the verifier performing the linearity part of the test (i.e. $c = a + b$) it is $1 - \epsilon_l$, so that $\epsilon = 2\epsilon_c/3 + \epsilon_l/3$. These probabilities can be translated into statements on the provers' measurements as follows.

$$\begin{aligned} 1 - \epsilon_c &= \mathbf{E}_a \sum_\alpha \big\langle \tilde{M}_a^\alpha \tilde{N}_a^\alpha \big\rangle_\psi \\ &= \mathbf{E}_a \, \mathrm{CON}_\psi(\tilde{M}_a, \tilde{N}_a) \\ &= \mathbf{E}_{ab} \, \mathrm{CON}_\psi(M_{a|ab}, \tilde{N}_a), \qquad (7) \\ 1 - \epsilon_l &= \mathbf{E}_{ab} \sum_{\alpha,\beta} \big\langle M_{ab}^{\alpha\beta} \tilde{N}_{a+b}^{(\alpha\beta)} \big\rangle_\psi \\ &= \mathbf{E}_{ab} \sum_{\alpha,\beta} \big\langle M_{a|ab}^\alpha M_{b|ab}^\beta \tilde{N}_{a+b}^{(\alpha\beta)} \big\rangle_\psi, \qquad (8) \end{aligned}$$

where the last equality uses that the POVM elements $M_{ab}^{\alpha\beta}$ are projectors. Using again Naimark's dilation theorem the POVM $\{\tilde{N}_a^\alpha\}$ acting on the second register of $|\psi\rangle$ can be simulated by a projective measurement $\{Y_a^\alpha\}$ acting on the state $|\psi'\rangle = |\psi\rangle \otimes 2^{-n/2} \sum_a |a\rangle$. Next we construct a set of "exactly linear" observables using the projectors Y_a^α. Let $d(a|ab) = d_{\psi'}\big(M_{a|ab}, Y_a\big)$, so that by Lemma 5,

$$d(a|ab)^2 = O(\mathrm{CON}_{\psi'}(M_{a|ab}, Y_a)),$$

and taking expectation values and applying Jensen's inequality

$$\begin{aligned} \mathbf{E}_{ab} \, d(a|ab) &\leq \sqrt{\mathbf{E}_{ab} \, d(a|ab)^2} \\ &= O\Big(\sqrt{\mathbf{E}_{ab} \, \mathrm{CON}_\psi(M_{a|ab}, \tilde{N}_a)}\Big) \\ &= O(\sqrt{\epsilon_c}). \qquad (9) \end{aligned}$$

Introduce observables $Y_a = Y_a^{+1} - Y_a^{-1}$. For every $u \in \{0,1\}^n$ consider the Fourier transform $\hat{Y}_u = \mathbf{E}_a (-1)^{a \cdot u} Y_a$. Define measurement operators $B^u = (\hat{Y}_u)^2$. By Parseval's identity, these operators form a POVM. Using Naimark's theorem there exists an extended state $|\psi''\rangle$ and a projective measurement $\{C^u\}$ that simulates B^u. Introduce the binary projective measurement

$$C_a^\alpha = \sum_{u:(-1)^{u \cdot a} = \alpha} C^u,$$

and the corresponding observable

$$C_a = C_a^{+1} - C_a^{-1} = \sum_u (-1)^{u \cdot a} C^u.$$

From the orthogonality of the projectors C^u it follows that $C_a C_b = C_{a+b}$, so that $\{C_a\}$ is perfectly linear. It remains to show that the operators C_a^α are consistent with the second prover's operators $\tilde{N}_a^\alpha$, on the state $|\psi''\rangle$ (where we extend $\tilde{N}_a^\alpha$ by making it act as the identity on the ancilla registers).

Write

$$\begin{aligned}
\mathbf{E}_a \,\mathrm{CON}_{\psi''}(\tilde{N}_a, C_a) &= \mathbf{E}_a \sum_{\alpha} \mathrm{Re}\big(\langle \tilde{N}_a^{\alpha} C_a^{\alpha} \rangle_{\psi''}\big) \\
&= \mathbf{E}_a \sum_{\alpha} \sum_{u:(-1)^{u\cdot a}=\alpha} \mathrm{Re}\big(\langle Y_a^{\alpha} B^u \rangle_{\psi'}\big) \\
&= \mathbf{E}_a \sum_{\alpha} \sum_{u:(-1)^{u\cdot a}=\alpha} \Big\langle \frac{1}{2}\big(1+(-1)^{\alpha} Y_a\big)(\hat{Y}_u)^2 \Big\rangle_{\psi'} \\
&= \frac{1}{2} + \frac{1}{2}\,\mathbf{E}_a \sum_{u} (-1)^{u\cdot a} \big\langle Y_a (\hat{Y}_u)^2 \big\rangle_{\psi'} \\
&= \frac{1}{2} + \frac{1}{2}\,\mathbf{E}_a \,\big\langle \hat{Y}_u^3 \big\rangle_{\psi'}.
\end{aligned}$$

To conclude, this last expression can be bounded as

$$\begin{aligned}
\sum_u \big\langle \hat{Y}_u^3 \big\rangle_{\psi'} &= \sum_u \big\langle (\mathbf{E}_{abc} (-1)^{u\cdot(a+b+c)} Y_a Y_b Y_c) \big\rangle_{\psi'} \\
&= \mathbf{E}_{ab} \,\big\langle (Y_a Y_b Y_{a+b}) \big\rangle_{\psi'} \\
&= \mathbf{E}_{ab} \sum_{\alpha\beta} \big\langle (Y_a^{\alpha} Y_b^{\beta} Y_{a+b}^{\alpha\beta} - Y_a^{\alpha} Y_b^{\beta} Y_{a+b}^{-\alpha\beta}) \big\rangle_{\psi'} \\
&= \mathbf{E}_{ab} \sum_{\alpha_1\alpha_2} \big\langle (2 Y_a^{\alpha} Y_b^{\beta} Y_{a+b}^{\alpha\beta} - Y_a^{\alpha} Y_b^{\beta}) \big\rangle_{\psi'} \\
&= 2\,\mathbf{E}_{ab} \sum_{\alpha\beta} \big\langle (Y_a^{\alpha} Y_b^{\beta} Y_{a+b}^{\alpha\beta}) \big\rangle_{\psi'} - 1 \\
&\geq 2\,\mathbf{E}_{ab} \Big(\sum_{\alpha\beta} \big\langle (M_{a|ab}^{\alpha} M_{b|ab}^{\beta} Y_{a+b}^{\alpha\beta}) \big\rangle_{\psi'} - O\big(d(a|ab) + d(b|ab)\big) \Big) - 1 \\
&= 1 - O\big(\epsilon_l + \sqrt{\epsilon_c}\big),
\end{aligned}$$

where the inequality uses Lemma 4 twice and the last line is by (8) and (9).

□

4.2 Stabilizer Test

The stabilizer test is described in Figure 1. The following lemma states the main consequence we will use.

Lemma 11. *Suppose the strategy $(N, |\psi\rangle)$ succeeds in the stabilizer test with probability $1-\epsilon$. Then there exists $\epsilon_{stab} = O(\sqrt{\epsilon})$ such that*

$$\Big(\mathbf{E}_a \big\| (\hat{X}(a) - \bar{X}(a))|\psi\rangle \big\|^2 \Big)^{1/2} \leq \epsilon_{stab} \qquad \textit{and} \qquad \Big(\mathbf{E}_a \big\| (\hat{Z}(a) - \bar{Z}(a))|\psi\rangle \big\|^2 \Big)^{1/2} \leq \epsilon_{stab},$$

where $\hat{X}(a)$, $\hat{Z}(b)$ and $\bar{X}(a)$, $\bar{Z}(b)$ are observables defined from the provers' strategies in (5).

Moreover, the honest strategy succeeds in the test with probability 1.

Proof. It follows from the definition of CON_ψ that any strategy $(N, |\psi\rangle)$ succeeding in the test with probability $1-\epsilon$ satisfies

$$\mathbf{E}_a \, \mathrm{CON}_\psi(\hat{X}(a), \bar{X}(a)) \geq 1 - O(\epsilon)$$
$$\mathbf{E}_a \, \mathrm{CON}_\psi(\hat{Z}(a), \bar{Z}(a)) \geq 1 - O(\epsilon).$$

By applying Lemma 5 to the above relations, we obtain:

$$\mathbf{E}_a \, d_\psi(\hat{X}(a), \bar{X}(a)) = O(\sqrt{\epsilon}) \qquad \text{and} \qquad \mathbf{E}_a \, d_\psi(\hat{Z}(a), \bar{Z}(a)) = O(\sqrt{\epsilon}),$$

Expression (2) for the state-dependent distance d_Ψ between two observables yields a bound on the squared norm. □

4.3 Anticommutation Test

The anticommutation test is described in Figure 1. The goal of the test is to certify that

$$\hat{X}(a)\hat{Z}(b)|\psi\rangle \approx (-1)^{a \cdot b} \hat{Z}(a)\hat{X}(b)|\psi\rangle,$$

where $\hat{X}(a)$ and $\hat{Z}(b)$ are the observables defined from the provers' strategies in (5). There are two cases: if a and b overlap on an even number of positions, then the two operators should commute; otherwise, they should anti-commute. The anticommutation test enforces the latter property. In Section 4.4 we show how the former can be derived as a consequence.

Lemma 12. *There exists $\omega^*_{\textit{anti-com}} \in (0,1)$ such that the following holds. Suppose the strategy $(N, |\psi\rangle)$ succeeds in the anticommutation test with probability $\omega^*_{\textit{anti-com}} - \epsilon$ and in the stabilizer test with probability $1 - \epsilon_{stab}$. Then there exists $\epsilon_{ac} = O(\sqrt{\epsilon}) + O(\sqrt{\epsilon_{stab}})$ such that*

$$\Big(\mathbf{E}_{a,b:a\cdot b=1} \left\| \big(\hat{X}(a)\hat{Z}(b) - (-1)^{a\cdot b}\hat{Z}(b)\hat{X}(a)\big)|\psi\rangle \right\|^2 \Big)^{1/2} = \epsilon_{ac}.$$

Moreover, the honest strategy succeeds in this test with probability $\omega^*_{\textit{anti-com}}$.

We remark that the constant $\omega^*_{\text{anti-com}}$ is closely related to the maximum quantum winning probability of the CHSH game.

Proof. The analysis follows very closely that of the "special-player" stabilizer game in [Ji15, Section 3.2], which is in turn based on the CHSH game. The key idea is to show the existence of "rotated" operators $\hat{X}'$ and $\hat{Z}'$ acting on the special prover such that $\frac{\hat{X}'+\hat{Z}'}{\sqrt{2}}$ (resp. $\frac{\hat{X}'-\hat{Z}'}{\sqrt{2}}$) is consistent with $\bar{X}(a)$ (resp. $\bar{Z}(b)$) on the composite prover. Since we know from the analysis of the stabilizer test (Lemma 11) that $\bar{X}(a)$ and $\bar{Z}(b)$ are consistent with $\hat{X}(a)$ and $\hat{X}(b)$ respectively, this allows us to conclude that $\frac{\hat{X}'+\hat{X}'}{\sqrt{2}}$ is consistent with $\hat{X}(a)$ and $\frac{\hat{X}'-\hat{Z}'}{\sqrt{2}}$ is consistent with $\hat{Z}(b)$. Since the two operators $\frac{\hat{X}'\pm\hat{Z}'}{\sqrt{2}}$ anticommute *exactly* with each other by construction, the operators $\hat{X}(a)$ and $\hat{Z}(b)$ must thus *approximately* anti-commute. Note that the rotated operators $\hat{X}'$ and $\hat{Z}'$ are allowed to depend on both a and b; we leave this dependence implicit in the notation and similarly suppress it from $\bar{X}(a)$ and $\bar{Z}(b)$.

We now analyze the test in more detail. First, recall that as stated in the protocol, a and b are chosen independently and uniformly at random, so they have a chance of $1/2$ of having inner

product 0, and in this case the test always accepts. Let p_{suc} be the success probability of the test and define the bias β as:

$$\begin{aligned}\beta = 16p_{suc} - 12 &= 8\,\mathbf{E}_{a,b:a\cdot b=1}\left(\frac{1}{2} + \frac{1}{8}\langle \hat{X}'\bar{X} + \hat{Z}'\bar{X} + \hat{X}'\bar{Z} - \hat{Z}'\bar{Z}\rangle_\Psi\right) - 4 \\ &= \langle \hat{X}'\bar{X} + \hat{Z}'\bar{X} + \hat{X}'\bar{Z} - \hat{Z}'\bar{Z}\rangle_\Psi.\end{aligned}$$

The bias can be decomposed as a sum of squares as follows:

$$\frac{2\beta}{\sqrt{2}} = \mathbf{E}_{a,b:a\cdot b=1}\left(4 - \left\langle\left(\frac{\hat{X}'+\hat{Z}'}{\sqrt{2}} - \bar{X}\right)^2\right\rangle_\Psi - \left\langle\left(\frac{\hat{X}'-\hat{Z}'}{\sqrt{2}} - \bar{Z}\right)^2\right\rangle_\Psi\right).$$

Let $\omega^*_{\text{anti-com}} = \frac{3}{4} + \frac{\sqrt{2}}{8}$. If the success probability is $p_{suc} = \omega^*_{\text{anti-com}} - \epsilon$, then $\beta = 2\sqrt{2} - 16\epsilon$. From this and the sum-of-squares decomposition above, we deduce

$$\mathbf{E}_{a,b:a\cdot b=1}\left\langle\left(\frac{\hat{X}'+\hat{Z}'}{\sqrt{2}} - \bar{X}\right)^2\right\rangle_\Psi \leq \frac{32}{\sqrt{2}}\epsilon, \qquad \mathbf{E}_{a,b:a\cdot b=1}\left\langle\left(\frac{\hat{X}'-\hat{Z}'}{\sqrt{2}} - \bar{Z}\right)^2\right\rangle_\Psi \leq \frac{32}{\sqrt{2}}\epsilon. \tag{10}$$

Now, for *any* observables $\hat{X}', \hat{Z}'$, the following anti-commutation relation holds:

$$\begin{aligned}\{\hat{X}'+\hat{Z}', \hat{X}'-\hat{Z}'\} &= (\hat{X}'+\hat{Z}')(\hat{X}'-\hat{Z}') + (\hat{X}'-\hat{Z}')(\hat{X}'+\hat{Z}') \\ &= I - \hat{X}'\hat{Z}' + \hat{Z}'\hat{X}' - I + I + \hat{X}'\hat{Z}' - \hat{Z}'\hat{X}' - I \\ &= 0.\end{aligned}$$

Thus, the operators $(\hat{X}'+\hat{Z}')/\sqrt{2}, (\hat{X}'-\hat{Z}')/\sqrt{2}$ are exactly anticommuting. The anti-commutator of $\hat{X}(a)$ and $\hat{Z}(b)$ is

$$\begin{aligned}&\mathbf{E}_{a,b:a\cdot b=1}\left\|(\hat{X}(a)\hat{Z}(b) - \hat{Z}(b)\hat{X}(a))|\psi\rangle\right\|^2 \\ &\quad= \mathbf{E}_{a,b:a\cdot b=1}\left\|(\bar{X}(a)\bar{Z}(b) - \bar{Z}(b)\bar{X}(a))|\psi\rangle\right\|^2 + O(\sqrt{\epsilon_{stab}}) \\ &\quad= \frac{1}{4}\,\mathbf{E}_{a,b:a\cdot b=1}\left\|((\hat{X}'+\hat{Z}')(\hat{X}'-\hat{Z}') - (\hat{X}'-\hat{Z}')(\hat{X}'+\hat{Z}'))|\psi\rangle\right\|^2 + O(\sqrt{\epsilon}) + O(\sqrt{\epsilon_{stab}}) \\ &\quad= O(\sqrt{\epsilon}) + O(\sqrt{\epsilon_{stab}}),\end{aligned}$$

where we replaced the X and Z operators first by $\bar{X}, \bar{Z}$, and then by $(\hat{X}' \pm \hat{Z}')/\sqrt{2}$, and bounded the error using (10) and the Cauchy-Schwarz inequality. This proves the lemma. □

4.4 Commutation

The protocol does not involve a test for commutation, as the required property can be derived as a consequence of the encoding tests.

Lemma 13. *Suppose the strategy $(N, |\psi\rangle)$ succeeds in the anti-commutation, linearity and stabilizer tests with probability $1-\epsilon$ each. Then there exists $\epsilon_{com} = O(\sqrt{\epsilon})$ such that*

$$\left(\mathbf{E}_{a,b:a\cdot b=0}\left\|\big(\hat{X}(a)\hat{Z}(b) - \hat{Z}(b)\hat{X}(a)\big)|\psi\rangle\right\|^2\right)^{1/2} = \epsilon_{com}.$$

Proof. We combine the anticommutation, linearity, and stabilizer tests through the following sequence of approximate identities, where the symbol $\approx$ means that two states are $O(\epsilon_{lin} + \epsilon_{stab} + \epsilon_{anti-com})$-close in squared Euclidean norm. In particular, we use the stabilizer test to "push" an operator from the special prover onto the other provers, which allows us to commute it past other operators acting on the special prover. First we relate the observables associated with the special prover's strategy to the exactly linear observables $\mathcal{X}$, $\mathcal{Z}$ obtained from the linearity test (see Definition 10).

$$\begin{aligned}
\mathbf{E}_{a,b:a\cdot b=0}\, \hat{X}(a)\hat{Z}(b)|\psi\rangle &\approx \mathbf{E}_{a,b:a\cdot b=0}\, \hat{X}(a)\bar{Z}(b)|\psi\rangle \\
&\approx \mathbf{E}_{a,b:a\cdot b=0}\, \mathcal{X}(a)\bar{Z}(b)|\psi\rangle \\
&\approx \mathbf{E}_{a,b:a\cdot b=0}\, \mathcal{X}(a)\hat{Z}(b)|\psi\rangle \\
&\approx \mathbf{E}_{a,b:a\cdot b=0}\, \mathcal{X}(a)\mathcal{Z}(b)|\psi\rangle,
\end{aligned}$$

where the first and third lines follow from the stabilizer test (Lemma 11) and the second and fourth from the linearity test (Lemma 9). Next write

$$\begin{aligned}
\mathbf{E}_{a,b:a\cdot b=0}\, \mathcal{X}(a)\mathcal{Z}(b)|\psi\rangle &= \mathbf{E}_{a,b:a\cdot b=0}\, \mathbf{E}_{c:c\cdot a=1}\, \mathcal{X}(a)\mathcal{Z}(c)\mathcal{Z}(c+b)|\psi\rangle \\
&\approx \mathbf{E}_{a,b:a\cdot b=0}\, \mathbf{E}_{c:c\cdot a=1}\, \bar{Z}(c+b)\mathcal{X}(a)\mathcal{Z}(c)|\psi\rangle \\
&\approx -\,\mathbf{E}_{a,b:a\cdot b=0}\, \mathbf{E}_{c:c\cdot a=1}\, \bar{Z}(c+b)\mathcal{Z}(c)\mathcal{X}(a)|\psi\rangle \\
&\approx -\,\mathbf{E}_{a,b:a\cdot b=0}\, \mathbf{E}_{c:c\cdot a=1}\, \mathcal{Z}(c)\mathcal{X}(a)\mathcal{Z}(c+b)|\psi\rangle \\
&\approx \mathbf{E}_{a,b:a\cdot b=0}\, \mathbf{E}_{c:c\cdot a=1}\, \mathcal{Z}(c)\mathcal{Z}(c+b)\mathcal{X}(a)|\psi\rangle \\
&= \mathbf{E}_{a,b:a\cdot b=0}\, \mathcal{Z}(b)\mathcal{X}(a)|\psi\rangle \\
&\approx \mathbf{E}_{a,b:a\cdot b=0}\, \hat{Z}(b)\hat{X}(a)|\psi\rangle.
\end{aligned}$$

Here the first equality uses the exact linearity (6) of $\mathcal{X}$ and $\mathcal{Z}$. The second line uses the linearity test and the stabilizer test. The third line uses approximate anticommutation (Lemma 12). The fourth line again uses the stabilizer and linearity tests, and the fifth line uses approximate anticommutation. The sixth line uses exact linearity, and the last is obtained from the linearity and stabilizer tests. □

5 A game for the local Hamiltonian problem

In this section we complete the analysis of the protocol described in Section 3 and prove our main theorem. The encoding tests have already been considered in the previous section, and it remains to define and analyze the energy test; this is done in the following subsection. In Section 5.2 we introduce an isometry that will let us "extract" an n-qubit state, destined to play the role of the QMA witness for the local Hamiltonian instance under consideration, from the strategy of the provers. Theorem 19 in Section 5.3 summarizes the result of the analysis so far, stating it in terms of a gap-preserving reduction. Finally Theorem 1 is proved in Section 5.4 by combining Theorem 19 with a gap amplification procedure.

5.1 Energy Test

In the energy test the verifier asks the provers to measure a local Hamiltonian on their shared encoded state. The test is made of two subtests, each to be performed with probability half: the

energy measurement test and the *consistency test*.

5.1.1 Energy measurement test

The goal of the energy measurement test is to estimate the energy of a randomly chosen term in the Hamiltonian. We analyze the test for the general case of a Hamiltonian that is not necessarily local, but such that H can be decomposed as

$$H = \frac{1}{m}\sum_{\ell=1}^{m} \alpha_\ell P_\ell, \tag{11}$$

where each P_ℓ is an n-qubit operator consisting of a tensor product of single-qubit I, X and Z Pauli operators and the real coefficients α_ℓ satisfy $|\alpha_\ell| \leq 1$ for all $\ell \in \{1, \ldots, m\}$. The energy measurement test proceeds as follows:

- For each $\ell \in \{1, \ldots, m\}$ define an operator Q_ℓ acting on rn qubits by replacing each Pauli X in P_ℓ with $X_{logical}$ on the r-qubit code state, and each Pauli Z by $Z_{logical}$.
- Send each prover an XZ-query (see Section 3) representing the associated share of Q_ℓ.
- Each prover replies with two values in $\{-1, 1\}$. Take the product of all values received, and compare it to the sign of α_ℓ. If the signs disagree, accept. If the signs agree, reject with probability $|\alpha_\ell|$ and accept with probability $1 - |\alpha_\ell|$.

Lemma 14. *The acceptance probability of the energy measurement test, when the correct Pauli operators are applied by each prover on its respective register of an (rn)-qubit state $|\psi\rangle$, is*

$$\begin{aligned} w^*_{energy}(H, |\psi\rangle) &= 1 - \Big(\frac{1}{2m}\sum_{\ell=1}^{m} \frac{|\alpha_\ell| + \alpha_\ell \langle\psi|P_\ell|\psi\rangle}{2}\Big) \\ &= 1 - \Big(\frac{1}{4}\langle\psi|H|\psi\rangle + \frac{1}{2m}\sum_{\ell} |\alpha_\ell|\Big). \end{aligned}$$

Proof. The proof is a simple calculation in all points similar to that performed in [Ji15, Section 4]; see in particular the discussion that precedes Theorem 23 in that paper. We omit the details. □

Note that this lemma only describes the behavior of the honest provers. The corresponding soundness result for dishonest provers is essentially our main theorem (Theorem 19).

5.1.2 Consistency test

The goal of the consistency test is to ensure that the measurement operators used by the provers in the energy test are consistent with the operators $\mathcal{X}$ and $\mathcal{Z}$ defined from their strategies in (5). In the test the verifier first selects a local term P_ℓ from the Hamiltonian uniformly at random. The energy measurement test considers an associated term P^i_ℓ, $i = 1, \ldots, r$, for each prover, which can be written as a product $X(a_i)Z(b_i)$ of X and Z operators for non-overlapping strings a_i and b_i. Let i be the index of the special prover. The verifier performs one of the following tests, each chosen with the indicated probability.

- With probability $1/2$, send the special prover (X, a, Z, b), and the composite prover $\overline{(X, c, c+a)}$, where $c \in \{0,1\}^n$ is chosen uniformly at random. Accept if the special prover's X-answer agrees with the product of the composite prover's two answers.
- With probability $1/4$, send the special prover (X, c, d), and the composite prover $\overline{(X, c, c+a)}$, where $c, d \in \{0,1\}^n$ are chosen uniformly at random. Accept if the special prover and composite prover agree on $X(c)$.
- With probability $1/4$, send the special prover $(X, c+a, d)$, and the composite prover $\overline{(X, c, c+a)}$, where $c, d \in \{0,1\}^n$ are chosen uniformly at random. Accept if the special prover and composite prover agree on $X(c+a)$.

The same tests are performed with the role of X and Z (to the composite prover) interchanged, the X and Z-variants being selected by the verifier with probability $1/2$ each.

Lemma 15. *Suppose the strategy $(N, |\psi\rangle)$ for the provers succeeds in the consistency test with probability $1 - \epsilon_{cons}$ and in the encoding tests with probability $1 - \epsilon_{encode}$, then*

$$\frac{1}{m}\sum_{\ell=1}^{m} \left\|(\hat{P}_\ell^i - \mathcal{X}(a)\mathcal{Z}(b))|\psi\rangle\right\|^2 \;=\; O\big(\epsilon_{cons}^{1/4}\big) + O\big(\epsilon_{encode}^{1/4}\big),$$

where a and b are strings such that $P_\ell^i = X(a)Z(b)$, and $\hat{P}_\ell^i$ is the measurement applied by the special prover upon receiving the query (X, a, Z, b) in the energy test.

Moreover, honest provers (see Definition 6) pass the test with probability 1.

Proof. We show that XZ-queries, X-queries, and Z-queries on the special prover are all consistent with $\overline{(X, c, c+a)}$ and $\overline{(Z, c, c+b)}$ queries to the composite prover. The analysis uses similar techniques to the analysis of the linearity test. First, let us analyze the X case. Let the POVM applied by the composite prover be $\{M_{c,c+a}^{\alpha\alpha'}\}$ and define marginalized operators

$$M_{c|c,c+a}^{\alpha} = \sum_{\alpha'} M_{c,c+a}^{\alpha\alpha'}.$$

Likewise, define marginalized operators for the special prover:

$$\hat{P}_{a|\ell}^{\alpha} = \sum_{\alpha'} P_\ell^{\alpha\alpha'}, \qquad \hat{P}_{b|\ell}^{\alpha'} = \sum_{\alpha} P_\ell^{\alpha\alpha'}.$$

Here we have suppressed the superscript i indicating the index of the special prover.

The following consistency relations follow from the assumption that the provers succeed with probability $1 - \epsilon_{cons}$ in the test. We use the notation $\mathbf{E}_{a\sim P_\ell}$ to indicates that the string a is chosen from the distribution of queries induced by the Hamiltonian, in contrast to $\mathbf{E}_a$ which indicates a uniformly random string.

$$\mathbf{E}_{a\sim P_\ell}\, \mathbf{E}_c\, \mathrm{CON}(M_{c|c,c+a}^{\alpha}, \hat{X}(c)^{\alpha}) \geq 1 - O(\epsilon_{cons}) \tag{12}$$

$$\mathbf{E}_{a\sim P_\ell}\, \mathbf{E}_c\, \mathrm{CON}(M_{c+a|c,c+a}^{\alpha}, \hat{X}(c+a)^{\alpha}) \geq 1 - O(\epsilon_{cons}) \tag{13}$$

$$\mathbf{E}_{a\sim P_\ell}\, \mathbf{E}_c\, \mathrm{CON}(\hat{P}_{a|\ell}^{\alpha}, \sum_{\beta\cdot\beta'=\alpha} M_{c,c+a}^{\beta\beta'}) \geq 1 - O(\epsilon_{cons}). \tag{14}$$

We now use these relations to show that the special prover's marginalized measurement $\hat{P}^{\alpha}_{a|\ell}$ is close to the operator $\mathcal{X}^{\alpha}(a)$ produced by the linearity test. We show this in two steps. First, we relate $\hat{P}^{\alpha}_{a|\ell}$ to the composite prover's measurement:

$$\begin{aligned}
&\mathbf{E}_{a\sim P_\ell} \operatorname{CON}\Big(\hat{P}^{\alpha}_{a|\ell}, \mathcal{X}^{\alpha}(a)\Big) \\
&\quad\geq \mathbf{E}_{a\sim P_\ell}\, \mathbf{E}_c \Big[\operatorname{CON}\Big(\sum_{\beta\cdot\beta'=\alpha} M^{\beta\beta'}_{c,c+\alpha}, \mathcal{X}^{\alpha}(a)\Big) - d_\Psi\Big(\hat{P}^{\alpha}_{a|\ell}, \sum_{\beta\cdot\beta'=\alpha} M^{\beta\beta'}_{c,c+a}\Big)\Big] \\
&\quad\geq \mathbf{E}_{a\sim P_\ell}\, \mathbf{E}_c \operatorname{CON}\Big(\sum_{\beta\cdot\beta'=\alpha} M^{\beta\beta'}_{c,c+a}, \mathcal{X}^{\alpha}(a)\Big) - O(\sqrt{\epsilon_{cons}}) \\
&\quad= \mathbf{E}_{a\sim P_\ell}\, \mathbf{E}_c \sum_{\alpha} \sum_{\beta\cdot\beta'=\alpha} \big\langle M^{\beta\beta'}_{c,c+a} \mathcal{X}^{\alpha}(a)\big\rangle_\Psi - O(\sqrt{\epsilon_{cons}}).
\end{aligned}$$

In the above, we used Lemma 4 to go from the first to the second line, and then lemma 5 and (14) to go to the third line. Next we relate M to a product of two measurements $\hat{X}$:

$$\begin{aligned}
\mathbf{E}_{a\sim P_\ell} \operatorname{CON}\Big(\hat{P}^{\alpha}_{a|\ell}, \mathcal{X}^{\alpha}(a)\Big) &= \mathbf{E}_{a\sim P_\ell}\, \mathbf{E}_c \sum_{\alpha} \sum_{\beta\cdot\beta'=\alpha} \big\langle M^{\beta}_{c|c,c+a} M^{\beta'}_{c+a|c,c+a} \mathcal{X}^{\alpha}(a)\big\rangle_\Psi - O(\sqrt{\epsilon_{cons}}) \\
&\geq \mathbf{E}_{a\sim P_\ell}\, \mathbf{E}_c \sum_{\alpha} \sum_{\beta\cdot\beta'=\alpha} \big\langle \hat{X}^{\beta}(c) \hat{X}^{\beta'}(c+a) \mathcal{X}^{\alpha}(a)\big\rangle_\Psi - O(\sqrt{\epsilon_{cons}}).
\end{aligned}$$

Here, we used equations (12) and (13), together with Lemmas 4 and 5. Finally, we use the encoding test to relate $\hat{X}$ to the exactly linear observable $\mathcal{X}$:

$$\begin{aligned}
&\mathbf{E}_{a\sim P_\ell} \operatorname{CON}\Big(\hat{P}^{\alpha}_{a|\ell}, \mathcal{X}^{\alpha}(a)\Big) \\
&\quad\geq \mathbf{E}_{a\sim P_\ell}\, \mathbf{E}_c \sum_{\alpha} \sum_{\beta\cdot\beta'=\alpha} \big\langle \bar{X}^{\beta}(c) \hat{X}^{\beta'}(c+a) \mathcal{X}^{\alpha}(a)\big\rangle_\Psi - O(\sqrt{\epsilon_{cons}} + \sqrt{\epsilon_{encode}}) \\
&\quad\geq \mathbf{E}_{a\sim P_\ell}\, \mathbf{E}_c \Big[\sum_{\alpha} \sum_{\beta\cdot\beta'=\alpha} \big\langle \bar{X}^{\beta}(c) \mathcal{X}^{\beta'}(c+a) \mathcal{X}^{\alpha}(a)\big\rangle_\Psi \\
&\qquad\qquad - d_\Psi(\hat{X}(c+a), \mathcal{X}(c+a))\Big] - O(\sqrt{\epsilon_{cons}} + \sqrt{\epsilon_{encode}}) \\
&\quad\geq \mathbf{E}_{a\sim P_\ell}\, \mathbf{E}_c \Big[\sum_{\alpha} \sum_{\beta\cdot\beta'=\alpha} \big\langle \mathcal{X}^{\beta}(c) \mathcal{X}^{\beta'}(c+a) \mathcal{X}^{\alpha}(a)\big\rangle_\Psi - d_\Psi(\hat{X}(c), \mathcal{X}(c))\Big] - O(\sqrt{\epsilon_{cons}} + \sqrt{\epsilon_{encode}}) \\
&\quad\geq \mathbf{E}_{a\sim P_\ell}\, \mathbf{E}_c \sum_{\alpha} \sum_{\beta\cdot\beta'=\alpha} \big\langle \mathcal{X}^{\beta}(c) \mathcal{X}^{\beta'}(c+a) \mathcal{X}^{\alpha}(a)\big\rangle_\Psi - O(\sqrt{\epsilon_{cons}} + \sqrt{\epsilon_{encode}}) \\
&\quad= 1 - O(\sqrt{\epsilon_{cons}} + \sqrt{\epsilon_{encode}}).
\end{aligned}$$

Here we used Lemma 11 for the first inequality, Lemma 4 for the second, Lemma 9 for the third, again Lemma 11 and Lemma 4 for the third and Lemma 9 for the fourth. The last equality follows by exact linearity (Definition 10). Note that we could use Lemma 11 here because the marginal distributions of c and $c+a$ are both completely uniform.

Performing an analogous analysis for the Z operators,

$$\mathbf{E}_{b\sim P_\ell} \operatorname{CON}(\hat{P}^{\beta}_{b|\ell}, \mathcal{Z}(b)^{\beta}) \geq 1 - O(\sqrt{\epsilon_{cons}}) - O(\sqrt{\epsilon_{encode}}).$$

Thus the special prover's operators are close to $\mathcal{X}(a)$ and $\mathcal{Z}(b)$ used in the encoding tests. To put these results together it remains to apply the stabilizer property to these operators; while we cannot do this directly since a and b are not distributed uniformly, we can use the exact linearity to write $\mathcal{Z}(b) = \mathbf{E}_c \, \mathcal{Z}(b+c)\mathcal{Z}(c)$, and apply Lemma 11 to each term in the product:

$$\begin{aligned}
\mathbf{E}_\ell \, \hat{P}_\ell |\psi\rangle &= \mathbf{E}_\ell \, \hat{P}_{a|\ell} \hat{P}_{b|\ell} |\psi\rangle \\
&\approx \mathbf{E}_\ell \, \hat{P}_{a|\ell} \mathcal{Z}(b) |\psi\rangle \\
&= \mathbf{E}_\ell \, \mathbf{E}_c \, \hat{P}_{a|\ell} \mathcal{Z}(b+c)\mathcal{Z}(c) |\psi\rangle \\
&\approx \mathbf{E}_\ell \, \mathbf{E}_c \, \bar{Z}(c)\bar{Z}(b+c) \hat{P}_{a|\ell} |\psi\rangle \\
&\approx \mathbf{E}_\ell \, \mathbf{E}_c \, \bar{Z}(c)\bar{Z}(b+c) \mathcal{X}(a) |\psi\rangle \\
&\approx \mathbf{E}_\ell \, \mathbf{E}_c \, \mathcal{X}(a)\mathcal{Z}(b+c)\mathcal{Z}(c) |\psi\rangle \\
&= \mathbf{E}_\ell \, \mathcal{X}(a)\mathcal{Z}(b) |\psi\rangle
\end{aligned}$$

where the approximate equalities are measured using the distance d_ψ and we consider approximations of order $O(\epsilon_{cons}^{1/4}) + O(\epsilon_{encode}^{1/4})$ (the exponent $1/4$ arises when converting CON to d_ψ via lemma 5). □

5.2 The isometry

Suppose given a strategy $(N, |\psi\rangle)$ for the provers that succeeds with probability $1-\epsilon$ in the protocol described in Section 3. We define an isometry that maps $|\psi\rangle$ to an n-qubit state $|\varphi\rangle$ such that $|\varphi\rangle$ would succeed in the energy test with a probability that deviates from the actual provers' by $O(\epsilon^{1/16})$, if it was measured according to the correct Pauli operators associated with the local term that the verifier chooses in the test (for any possible such choice).

Our construction uses an isometry introduced by McKague [McK13], applied independently to each prover's register of $|\psi\rangle$. This isometry produces an (rn)-qubit state $|\varphi_1\rangle$, where n qubits are defined locally for each one of the r provers. The single-prover case is described and analyzed first, in Section 5.2.1; the full isometry acting on all r provers is described in Section 5.2.2. In particular, we show in that section that the expectation value of the *encoded* Hamiltonian acting on the output state of the isometry is close to the measured expectation value of the provers.

5.2.1 The single-prover isometry

We describe the isometry associated with the first provers' register; the isometries for the remaining provers are similar. The isometry appends an n-qubit maximally entangled state to $|\psi\rangle$, and "swaps" part of $|\psi\rangle$ from the first register to the first half of the maximally entangled state. The output of the circuit is given by

$$|\varphi_0\rangle = \Phi^1(|\psi\rangle) = \frac{1}{2^{3n/2}} \sum_{z,y,w \in \{0,1\}^n} (-1)^{y\cdot(z+w)} \mathcal{X}(w)\mathcal{Z}(y)\mathcal{X}(z)|\psi\rangle \otimes |wz\rangle,$$

where the last two registers correspond to the maximally entangled state used as ancilla. Here $\mathcal{X}$ and $\mathcal{Z}$ are the exactly linear operators obtained from the special provers' measurement operators on X and Z queries respectively (see Definition 10).

Lemma 16. *For any $a, b \in \{0,1\}^n$ it holds that*

$$\big|\langle\varphi_0|X(a)Z(b)|\varphi_0\rangle - \langle\psi|\mathcal{X}(a)\mathcal{Z}(b)|\psi\rangle\big| \;=\; O\big(\sqrt{\epsilon_{lin}} + \sqrt{\epsilon_{ac}} + \sqrt{\epsilon_{com}} + \sqrt{\epsilon_{stab}}\big),$$

where $X(a)Z(b)$ is the corresponding tensor product Pauli operator acting on the second register of $|\varphi_0\rangle$ (associated with the first half of the maximally entangled state used as ancilla), $\mathcal{X}(a), \mathcal{Z}(b)$ are the special prover's exactly linear operators (as defined in Definition 10), and ϵ_{lin}, ϵ_{ac}, ϵ_{com} and ϵ_{stab} are defined in Lemma 9, Lemma 12, Lemma 13 and Lemma 11 respectively.

Proof. Expand

$$\begin{aligned}\langle\varphi_0|X(a)Z(b)|\varphi_0\rangle &= \frac{1}{2^{3n}}\sum_{yzw}\sum_{y'z'w'}\langle\psi|\langle wz|(-1)^{y\cdot(z+w)}\mathcal{X}(z)\mathcal{Z}(y)\mathcal{X}(w)\\ &\qquad\qquad\cdot X_2(a)Z_2(b)(-1)^{y'\cdot(z'+w')}\mathcal{X}(w')\mathcal{Z}(y')\mathcal{X}(z')|\psi\rangle|w'z'\rangle\\ &= \frac{1}{2^{3n}}\sum_{yzw}\sum_{y'z'w'}\langle\psi|\langle wz|(-1)^{y\cdot(z+w)}\mathcal{X}(z)\mathcal{Z}(y)\mathcal{X}(w)\\ &\qquad\qquad\cdot(-1)^{w'\cdot b}(-1)^{y'\cdot(z'+w')}\mathcal{X}(w')\mathcal{Z}(y')\mathcal{X}(z')|\psi\rangle|(w'+a)z'\rangle,\end{aligned}$$

where in going to the second line we used the fact that the operators X, Z commute with $\mathcal{X}, \mathcal{Z}$ and then the fact that $X(a)$ and $Z(b)$ are true Pauli operators to perform their action on the state. The resulting expression can be simplified as follows:

$$\begin{aligned}\langle\varphi_0|X(a)Z(b)|\varphi_0\rangle &= \frac{1}{2^{3n}}\sum_{yzy'w}(-1)^{z(y+y')+w(y+y'+b)+a(y'+b)}\\ &\qquad\qquad\cdot\langle\psi|\mathcal{X}(z)\mathcal{Z}(y)\mathcal{X}(w)\mathcal{X}(w+a)\mathcal{Z}(y')\mathcal{X}(z)|\psi\rangle\\ &= \mathbf{E}_{yz}(-1)^{z\cdot b+a\cdot y}\langle\psi|\mathcal{X}(z)\mathcal{Z}(y)\mathcal{X}(a)\mathcal{Z}(y+b)\mathcal{X}(z)|\psi\rangle,\end{aligned}$$

where we used exact linearity of $\mathcal{X}$ to combine $\mathcal{X}(w)\mathcal{X}(w+a)$ into $\mathcal{X}(a)$. We then evaluated the sum over w: this sum vanishes unless the coefficient of w in the exponent of the phase (-1) is equal to 0. This gives the relation $y + y' + b = 0$, which allows to eliminate the index y' from the summation.

Our goal is to bound the difference

$$\delta \;=\; \Big|\,\mathbf{E}_{yz}(-1)^{z\cdot b+a\cdot y}\langle\psi|\mathcal{X}(z)\mathcal{Z}(y)\mathcal{X}(a)\mathcal{Z}(y+b)\mathcal{X}(z)|\psi\rangle - \langle\psi|\mathcal{X}(a)\mathcal{Z}(b)|\psi\rangle\Big|$$

between this quantity and the expectation value $\langle\psi|\mathcal{X}(a)\mathcal{Z}(b)|\psi\rangle$. Using the triangle inequality,

$$\begin{aligned}\delta \le\; &\Big|\,\mathbf{E}_{yz}(-1)^{(z+a)\cdot y}\langle\psi|\mathcal{X}(z)\mathcal{Z}(y)\mathcal{X}(a+z)\mathcal{Z}(y+b)|\psi\rangle - \langle\psi|\mathcal{X}(a)\mathcal{Z}(b)|\psi\rangle\Big|\\ &+\Big|\,\mathbf{E}_{yz}(-1)^{(z+a)\cdot y}\langle\psi|\mathcal{X}(z)\mathcal{Z}(y)\mathcal{X}(a)\big(\mathcal{X}(z)\mathcal{Z}(y+b) - \mathcal{Z}(y+b)\mathcal{X}(z)\big)|\psi\rangle\Big|.\end{aligned}$$

This last term can be bounded using the results of the commutation test, Lemma 13, first applying the Cauchy-Schwarz inequality as

$$\begin{aligned}&\Big|\,\mathbf{E}_{yz}(-1)^{(z+a)\cdot y}\langle\psi|\mathcal{X}(z)\mathcal{Z}(y)\mathcal{X}(a)\big(\mathcal{X}(z)\mathcal{Z}(y+b) - \mathcal{Z}(y+b)\mathcal{X}(z)\big)|\psi\rangle\Big|\\ &\le \Big(\mathbf{E}_{yz}\,\big\|\mathcal{X}(z)\mathcal{Z}(y)\mathcal{X}(a)|\psi\rangle\big\|^2\Big)^{1/2}\Big(\mathbf{E}_{yz}\,\big\|\big(\mathcal{X}(z)\mathcal{Z}(y+b) - \mathcal{Z}(y+b)\mathcal{X}(z)\big)|\psi\rangle\big\|^2\Big)^{1/2}\\ &= O\big(\sqrt{\epsilon_{com}}\big),\end{aligned}$$

where to write the last equality we used the fact that both z and $y+b$ are uniformly distributed. Proceeding similarly but using first the stabilizer test, Lemma 11, then the anticommutation test, Lemma 12, and the stabilizer test again,

$$\begin{aligned}
\delta &= \Big| \mathbf{E}_{yz} (-1)^{(z+a)\cdot y} \langle\psi|\mathcal{X}(z)\mathcal{Z}(y)\mathcal{X}(a+z)\mathcal{Z}(y+b)|\psi\rangle - \langle\psi|\mathcal{X}(a)\mathcal{Z}(b)|\psi\rangle \Big| + O\big(\sqrt{\epsilon_{com}}\big) \\
&= \Big| \mathbf{E}_{yz} (-1)^{(z+a)\cdot y} \langle\psi|\bar{Z}(y+b)\mathcal{X}(z)\mathcal{Z}(y)\mathcal{X}(a+z)|\psi\rangle - \langle\psi|\mathcal{X}(a)\mathcal{Z}(b)|\psi\rangle \Big| + O\big(\sqrt{\epsilon_{com}}\big) \\
&\quad + O\big(\sqrt{\epsilon_{stab}}\big) \\
&= \big| \mathbf{E}_{yz} \langle\psi|\bar{Z}(y+b)\mathcal{X}(a)\mathcal{Z}(y)|\psi\rangle - \langle\psi|\mathcal{X}(a)\mathcal{Z}(b)|\psi\rangle \big| + O\big(\sqrt{\epsilon_{ac}}\big) + O\big(\sqrt{\epsilon_{com}}\big) + O\big(\sqrt{\epsilon_{stab}}\big) \\
&= \big| \mathbf{E}_{yz} \langle\psi|\mathcal{X}(a)\mathcal{Z}(b)|\psi\rangle - \langle\psi|\mathcal{X}(a)\mathcal{Z}(b)|\psi\rangle \big| + O\big(\sqrt{\epsilon_{ac}}\big) + O\big(\sqrt{\epsilon_{com}}\big) + O\big(\sqrt{\epsilon_{stab}}\big) \\
&= O\big(\sqrt{\epsilon_{lin}} + \sqrt{\epsilon_{ac}} + \sqrt{\epsilon_{com}} + \sqrt{\epsilon_{stab}}\big).
\end{aligned}$$

□

The preceding lemma shows that the output of the isometry matches any single prover's measurement of any Pauli operator. In particular, we can apply it to the queries made in the energy test.

Lemma 17. *For queries (X, a, Z, b) chosen under the distribution used in the energy measurement test it holds that*

$$\mathbf{E}_{a,b} \big| \langle\varphi_0|X(a)Z(b)|\varphi_0\rangle - \langle\psi|\hat{P}|\psi\rangle \big| \;=\; O\big(\epsilon_{lin}^{1/4} + \epsilon_{ac}^{1/4} + \epsilon_{com}^{1/4} + \epsilon_{stab}^{1/4} + \epsilon_{lin}^{1/4} + \epsilon_{cons}^{1/4}\big),$$

where $X(a)Z(b)$ is the corresponding tensor product Pauli operator acting on the second register of $|\varphi_0\rangle$, associated with the first half of the maximally entangled state used as ancilla, $\hat{P}$ is the special prover's operator associated with the query (X, a, Z, b), and ϵ_{lin}, ϵ_{ac}, ϵ_{com}, ϵ_{stab}, and ϵ_{cons} are defined in Lemma 9, Lemma 12, Lemma 13, Lemma 11, and Lemma 15, respectively.

Proof. The lemma is a direct consequence of Lemma 16 and Lemma 15. Note in particular that the exponent of $1/4$ arises from Lemma 15. □

5.2.2 The full isometry

We define an isometry acting on the joint state of all provers by composing the single-prover isometries above:

$$|\varphi_1\rangle = \Phi(|\psi\rangle) = (\Phi^1 \otimes \cdots \otimes \Phi^r)(|\psi\rangle). \tag{15}$$

Recall that the constant ω^*_{encode} from (4) is the success probability of the honest strategy in the encoding tests.

Proposition 18. *Suppose a strategy $(N, |\psi\rangle)$ succeeds in the encoding tests with probability $\omega^*_{encode} - \epsilon$, and in the energy test with probability ω_{energy}. Then there exists an (rn)-qubit state $|\varphi_1\rangle$ and a strategy in which each prover applies the honest strategy defined in Definition 6 to its share of $|\varphi_1\rangle$ that succeeds in the encoding tests with probability $\omega^*_{encode} - O(\epsilon^{1/16})$ and the energy test with probability at least $\omega^*_{energy} - O(\epsilon^{1/16})$.*

Proof. Let $\hat{P}_i(a_i, b_i)$ be the observable associated with prover i's measurement operator on query (X, a_i, Z, b_i), and let the corresponding true Pauli be $P_i(a_i, b_i) = X(a_i)Z(b_i)$. Let $|\varphi_1\rangle$ be the state obtained from $|\psi\rangle$ by applying the isometry to each prover, as in (15). Let $X(a) = X_1(a_1) \otimes X_2(a_2) \otimes \cdots \otimes X_r(a_r)$ be the tensor product of the r (true) n-qubit Pauli operators, and similarly $Z(b) = Z_1(b_1) \otimes Z_2(b_2) \otimes \cdots \otimes Z_r(b_r)$, where the a_i and b_i are derived from a and b in the energy measurement test. Applying Lemma 17 and the triangle inequality, we obtain

$$|\langle\varphi_1|X(a)Z(b)|\varphi_1\rangle - \langle\psi|\hat{P}_1(a_1, b_1) \otimes \cdots \otimes \hat{P}_r(a_r, b_r)|\psi\rangle| = O\big(\epsilon_{lin}^{1/4} + \epsilon_{ac}^{1/4} + \epsilon_{com}^{1/4} + \epsilon_{stab}^{1/4} + \epsilon_{lin}^{1/4} + \epsilon_{cons}^{1/4}\big),$$

where the first expression involves the true Pauli operators whereas the second is obtained from the provers' measurements. Using the dependence of each of ϵ_{lin}, ϵ_{ac}, ϵ_{com} and ϵ_{stab} on ϵ stated in the corresponding lemmas, this proves the proposition, as the newly defined strategy produces answers that have statistical distance $O(\epsilon^{1/16})$ from the real provers' answers in any of the tests of the protocol. □

5.3 Proof of the main theorem

We now show our main result: a game for the local Hamiltonian problem that has a constant completeness-soundness gap. As an intermediate step we first prove that the protocol described in Section 3 allows the verifier to estimate, up to a constant additive factor, the ground energy of *any* (nonlocal) Hamiltonian that can be expressed as a weighted sum of tensor products of single-qubit I, X and Z Pauli operators.

Theorem 19. *There exists constants $0 < c, c_1 < 1$ and $c_2 > 0$ such that the following holds. Let H be a (not necessarily local) Hamiltonian with m terms over n qubits of the form* (11)*, and $\lambda_{\min}(H)$ the smallest eigenvalue of H. Then for every $\delta > 0$ there is a choice $p = \Theta(\delta^c)$ for the probability of performing the energy test in the protocol described in Section 3 such that the maximum probability $\omega^*(H)$ with which any r-prover strategy can succeed in the protocol when the Hamiltonian is H is bounded as*

$$c_1 - c_2\delta^c\lambda_{min}(H) \leq \omega^*(H) \leq c_1 - c_2\delta^c\lambda_{min}(H) + \delta.$$

Proof. First we establish completeness. An honest quantum strategy (as described in Definition 6) acting on an encoded ground state $|\Gamma\rangle$ of H succeeds in the protocol with probability

$$\omega_{honest}(H) = (1-p)\omega^*_{encode} + p\,\omega^*_{energy}(H),$$

where

$$\begin{aligned}\omega^*_{energy}(H) &= \omega^*_{energy}(H, |\Gamma\rangle) \\ &= \frac{1}{2} + \frac{1}{2}\Big(1 - \frac{1}{4}\lambda_{\min}(H) - \frac{1}{2m}\sum_\ell |\alpha_\ell|\Big).\end{aligned}$$

(Recall that in the energy test with probability $1/2$ the verifier performs a consistency check, which passes with probability 1 for honest provers, and with probability $1/2$ the verifier performs the energy measurement test.) Next we establish soundness. Suppose a strategy for the provers succeeds with probability ω_{cheat}, passes the encoding tests with probability $\omega^*_{encode} - \epsilon$, and passes the energy test with probability ω_{energy}; thus

$$\omega_{cheat} = (1-p)(\omega^*_{encode} - \epsilon) + p\,\omega_{energy}.$$

Applying Proposition 18, there exists an (rn)-qubits state $|\varphi_1\rangle$ using which a strategy based on applying the true Pauli operators will succeed in the encoding and energy tests with probability at least $\omega^*_{encode} - O(\epsilon^{1/16})$ and $\omega_{energy} - O(\epsilon^{1/16})$ respectively. Since this strategy implements valid logical X and Z operators in the energy test, by lemma 14 it passes the energy test with probability at most $\omega^*_{energy}(H)$. Thus $\omega_{energy} \leq \omega^*_{energy}(H) + O(\epsilon^{1/16})$, and

$$\begin{aligned}\omega_{cheat} &= (1-p)(\omega^*_{encode} - \epsilon) + p\,\omega_{energy} \\ &\leq (1-p)(\omega^*_{encode} - \epsilon) + p\omega^*_{energy}(H) + O(p\,\epsilon^{1/16}) \\ &\leq \omega_{honest}(H) - (1-p)\epsilon + O(p\,\epsilon^{1/16})\end{aligned}$$

Choosing p to be a sufficiently small constant times $\delta^{15/16}$, for all $0 \leq \epsilon \leq 1$ this expression is less than or equal to $\omega_{honest}(H) + \delta$. □

5.4 Amplification

In this section we show how Theorem 19 can be used to obtain Theorem 1. The main idea consists in leveraging the fact that our protocol does not require locality of the Hamiltonian to first "brute-force" amplify the gap of the underlying instance of the local Hamiltonian problem to a constant, and then run the protocol on the amplified non-local instance. This is achieved by first shifting the Hamiltonian by the appropriate multiple of identity so that the energy in the yes-instance is less than or equal to 0. The gap is then amplified by taking sufficiently many tensor product copies of the Hamiltonian, resulting in a nonlocal instance.

Lemma 20 (Gap amplification)**.** *Let H be an n-qubit Hamiltonian with minimum energy $\lambda_{\min}(H) \geq 0$ and such that $\|H\| \leq 1$. Let $p(n), q(n)$ be polynomials such that $p(n) > q(n)$ for all n. Let*

$$H' = \mathbb{I}^{\otimes a} - (\mathbb{I} - (H - a^{-1}\,\mathbb{I}))^{\otimes a}, \qquad \textit{where} \qquad a = \Big(\frac{1}{q} - \frac{1}{p}\Big)^{-1}.$$

Then H' is a (non-local) Hamiltonian over qubits $an = O(np(n))$ qubits such that $\|H'\| = O(1)$ and if $\lambda_{\min}(H) \leq 1/p$ then $\lambda_{min}(H') \leq 1/2$ whereas if $\lambda_{\min}(H) \geq 1/q$ then $\lambda_{min}(H') \geq 1$.

Proof. The proof follows by observing that $\lambda_{\min}(H') = 1 - (1 - (\lambda_{min}(H) - a^{-1}))^a$, and $(1 \pm \delta)^k = 1 \pm k\delta + O(\delta^2)$ when $k\delta = O(1)$. □

Theorem 1 follows by applying the result of Theorem 19 to the Hamiltonian H' obtained from H as in Lemma 20.

Acknowledgements. AN was supported by the ARO grant Contract Number W911NF-12-0486. Parts of this work was completed while the second author was visiting the Institute for Quantum Information and Matter (IQIM) at Caltech, and both authors acknowledge funding provided by the IQIM, an NSF Physics Frontiers Center (NFS Grant PHY-1125565) with support of the Gordon and Betty Moore Foundation (GBMF-12500028).

References

[AALV09] Dorit Aharonov, Itai Arad, Zeph Landau, and Umesh Vazirani. The detectability lemma and quantum gap amplification. In *Proc. 41st STOC*, pages 417–426, New York, NY, USA, 2009. ACM.

[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.

[AB09] Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, 2009.

[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.

[Aro94] Sanjeev Arora. *Probabilistic checking of proofs and hardness of approximation problems*. PhD thesis, Princeton University, 1994.

[AS98] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of NP. *J. ACM*, 45(1):70–122, 1998.

[BFL91] László Babai, Lance Fortnow, and Carsten Lund. Non-deterministic exponential time has two-prover interactive protocols. *Computational Complexity*, 1:3–40, 1991.

[BLR93] Manuel Blum, Michael Luby, and Ronitt Rubinfeld. Self-testing/correcting with applications to numerical problems. *Journal of Computer and System Sciences*, 47:549–595, 1993.

[BVY15] Mohammad Bavarian, Thomas Vidick, and Henry Yuen. Anchoring games for parallel repetition. *arXiv preprint arXiv:1509.07466*, 2015.

[CHTW04] Richard Cleve, Peter Høyer, Ben Toner, and John Watrous. Consequences and limits of nonlocal strategies, 2004. arXiv:quant-ph/0404076.

[CM14] Toby Cubitt and Ashley Montanaro. Complexity classification of local hamiltonian problems. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 120–129. IEEE, 2014.

[CS96] A. R. Calderbank and Peter W. Shor. Good quantum error-correcting codes exist. *Phys. Rev. A*, 54:1098–1105, 1996. arXiv:quant-ph/9512032.

[DDG+14] Roee David, Irit Dinur, Elazar Goldberg, Guy Kindler, and Igor Shinkar. Direct sum testing. Technical report, ECCC Report TR14–002, 2014.

[Din07] Irit Dinur. The PCP theorem by gap amplification. *J. ACM*, 54(3), June 2007.

[FGL+96] Uriel Feige, Shafi Goldwasser, Laszlo Lovász, Shmuel Safra, and Mario Szegedy. Interactive proofs and the hardness of approximating cliques. *J. ACM*, 43(2):268–292, 1996.

[FV15] Joseph Fitzsimons and Thomas Vidick. A multiprover interactive proof system for the local Hamiltonian problem. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science*, pages 103–112. ACM, 2015.

[Got97] Daniel Gottesman. Stabilizer codes and quantum error correction, 1997. arXiv:quant-ph/9705052.

[IKM09] Tsuyoshi Ito, Hirotada Kobayashi, and Keiji Matsumoto. Oracularization and two-prover one-round interactive proofs against nonlocal strategies. In *Proceedings: Twenty-Fourth Annual IEEE Conference on Computational Complexity (CCC 2009)*, pages 217–228, July 2009.

[IV12] Tsuyoshi Ito and Thomas Vidick. A multi-prover interactive proof for NEXP sound against entangled provers. *Proc. 53rd FOCS*, pages 243–252, 2012.

[Ji15] Zhengfeng Ji. Classical verification of quantum proofs. Technical report, arXiv:1505.07432, 2015.

[KKMV09] Julia Kempe, Hirotada Kobayashi, Keiji Matsumoto, and Thomas Vidick. Using entanglement in quantum multi-prover interactive proofs. *Computational Complexity*, 18:273–307, 2009.

[KM03] Hirotada Kobayashi and Keiji Matsumoto. Quantum multi-prover interactive proof systems with limited prior entanglement. *JCSS*, 66:429–450, 2003. arXiv:cs/0102013.

[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.

[KV11] Julia Kempe and Thomas Vidick. Parallel repetition of entangled games. In *Proc. 43rd STOC*, pages 353–362, 2011.

[McK13] Matthew McKague. Interactive proofs for BQP via self-tested graphs tates, 2013. arXiv:1309.5675v2.

[McK14] Matthew McKague. Self-testing graph states. In *Theory of Quantum Computation, Communication, and Cryptography*, pages 104–120. Springer, 2014.

[McK15] Matthew McKague. Self-testing in parallel, 2015. arXiv:1511.04194.

[MdW13] Ashley Montanaro and Ronald de Wolf. A survey of quantum property testing. *arXiv preprint arXiv:1310.2035*, 2013.

[MYS12] M McKague, T H Yang, and V Scarani. Robust self-testing of the singlet. *Journal of Physics A: Mathematical and Theoretical*, 45(45):455304, 2012.

[NC01] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2001.

[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.

[Sha92] Adi Shamir. IP = PSPACE. *J. ACM*, 39(4):869–877, 1992.

[Ste96] Andrew Steane. Multiple-particle interference and quantum error correction. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, volume 452, pages 2551–2577. The Royal Society, 1996. arXiv:quant-ph/9601029.

[Vid13] Thomas Vidick. Three-player entangled XOR games are NP-hard to approximate. In *Proc. 54th FOCS*, 2013.