Retracted: Two-Player Entangled Games are NP-Hard

Anand Natarajan[1]
Center for Theoretical Physics, MIT, Cambridge, USA
anand@natarajans.edu
https://orcid.org/0000-0003-3648-3844

Thomas Vidick[2]
Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, USA
vidick@cms.caltech.edu
https://orcid.org/0000-0002-6405-365X

Abstract

The article, published on June 4th, 2018 in the CCC 2018 proceedings, has been retracted by agreement between the authors, the editor(s), and the publisher Schloss Dagstuhl / LIPIcs. The retraction has been agreed due to an error in the proof of the main result. This error is carried over from an error in the referenced paper "Three-player entangled XOR games are NP-hard to approximate" by Thomas Vidick (SICOMP '16). That paper was used in an essential way to obtain the present result, and the error cannot be addressed through an erratum. See Retraction Notice on page 19.

We show that it is NP-hard to approximate, to within an additive constant, the maximum success probability of players sharing quantum entanglement in a two-player game with classical questions of logarithmic length and classical answers of constant length. As a corollary, the inclusion $\mathsf{NEXP} \subseteq \mathsf{MIP}^*$, first shown by Ito and Vidick (FOCS'12) with three provers, holds with two provers only. The proof is based on a simpler, improved analysis of the low-degree test of Raz and Safra (STOC'97) against two entangled provers.

2012 ACM Subject Classification Theory of computation → Quantum complexity theory, Theory of computation → Interactive proof systems, Theory of computation → Complexity classes

Keywords and phrases low-degree testing, entangled nonlocal games, multi-prover interactive proof systems

Digital Object Identifier 10.4230/LIPIcs.CCC.2018.20

1 Introduction

Interactive proofs are a fundamental concept in theoretical computer science, with applications to complexity theory, cryptography, and more. A classic result [19, 24] shows that interaction is a powerful resource: the class IP of problems that a polynomial-time verifier can solve with access to a single, untrusted prover is equal to PSPACE. A subsequent line of works culminating in [3] showed that even more power can be gained by interacting with *multiple* provers: the class MIP of problems decidable by a polynomial-time verifier interacting with multiple non-communicating provers is equal to NEXP. This result was an important catalyst in the discovery of the PCP theorem [2, 1], a seminal result in complexity theory that has had broad-ranging implications for hardness of approximation [9, 12]. More recently, increasingly efficient probabilistically checkable proofs (PCPs) have played a major role in the design of protocols for delegated computation of space [11] or time-bounded [18] circuits.

[1] Supported by NSF CAREER Grant CCF-1452616

[2] Supported by NSF CAREER Grant CCF-1553477, AFOSR YIP award number FA9550-16-1-0495, and the IQIM, an NSF Physics Frontiers Center (NSF Grant PHY-1125565) with support of the Gordon and Betty Moore Foundation (GBMF-12500028).

33rd Computational Complexity Conference (CCC 2018).
Editor: Rocco A. Servedio; Article No. 20; pp. 20:1–20:19
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany
COMPUTATIONAL COMPLEXITY CONFERENCE

What happens when one considers a verifier that has the increased power of quantum polynomial-time computation, or provers that may use the non-local properties of quantum entanglement? In the single-prover setting, it is a highly non-trivial result that a quantum verifier, having the ability to exchange quantum messages with the prover, cannot decide more languages than a classical polynomial-time verifier: $\mathsf{QIP} = \mathsf{IP} = \mathsf{PSPACE}$ [15].

The story for multi-prover interactive proof systems is more complex. Cleve et al. [6] were the first to explore the consequences of entanglement for complexity theory. The class MIP^* is the class of languages having multi-prover interactive proofs between a classical polynomial-time verifier and quantum provers who may share entanglement. (The class QMIP^* allows a quantum verifier and quantum messages; it is known that $\mathsf{QMIP}^* = \mathsf{MIP}^*$ [23].) It has been known since the early days of quantum mechanics [8], and more specifically the work of Bell [4], that allowing spatially isolated provers to perform local measurements on a shared entangled state may allow them to generate correlations between their (classical) outputs that cannot be reproduced by any local model, even using shared randomness. In general, quantum strategies have a higher success probability than classical ones, and this can affect both the completeness and soundness parameters of a proof system. As a result, the only trivial lower bound on MIP^* is IP, since the verifier can ignore one of the provers, and there are no trivial upper bounds, as the size of entangled-prover strategies can be arbitrary. Cleve et al. [6] showed that entanglement could at least in some cases lead to a collapse of a complexity class based on an interactive proofs: they studied XOR proof systems and showed that $\oplus\mathsf{MIP}^* \subseteq \mathsf{PSPACE}$ (for any constant completeness-soundness gap), while it follows from Håstad's work [12] that $\oplus\mathsf{MIP} = \mathsf{NEXP}$ (for some choice of constant completeness and soundness parameters).

Nevertheless, a sequence of works established techniques to limit the power of entangled provers, eventually leading to a proof that $\mathsf{MIP} \subseteq \mathsf{MIP}^*$ [14] for proof systems involving four provers, a single round of interaction, and sufficiently large, but constant, answer size. The result is a corollary of the inclusion $\mathsf{NEXP} \subseteq \mathsf{MIP}^*$, whose proof follows the same structure as Babai et al.'s celebrated proof [3] that $\mathsf{NEXP} \subseteq \mathsf{MIP}$. The main technical component of the proof is an analysis of the soundness of Babai et al.'s multilinearity test with entangled provers. The result was later refined in [26], who obtained a scaled-down version that applies to multiplayer games specified in explicit form: the main result of [26] is that it is NP-hard to approximate the value of a three-player entangled game specified in explicit form (in contrast to an interactive proof system, which is specified by a family of circuits for the verifier). The proof rests on an analysis of the soundness of the "plane-vs-point" low-degree test [22], an improvement over Babai et al.'s multilinearity test, with entangled provers.

A rather intriguing limitation of the results in [14, 26] is that they only apply to games, or interactive proof systems, with three or more entangled players, or provers. Even though in any interaction the verifier in the proof systems considered in those works only exchanges messages with two out of the three provers,[3] the proof seems to crucially require that the joint Hilbert space supporting the provers' strategies can be decomposed in at least three tensor factors. Most importantly, this requirement is used in the proof of the "self-improvement lemma" that is key to control the accumulation of approximation errors in the inductive analysis of both the multilinearity and low-degree tests. Intuition for the requirement that there are three players is based on the phenomenon of monogamy of entanglement: it has been known at least since the work of Toner [25] that this kind of "embedding" of a two-player

[3] More precisely, all tests considered, including the low-degree test, take the form: (i) the verifier selects two provers at random, and calls them "Alice" and "Bob"; (ii) the verifier plays a two-prover game with Alice and Bob.

game in a three-player game can effectively limit the players' ability to take advantage of their shared entanglement, in some cases drastically lowering their maximum success probability in the game. Could it be that the two-prover entangled value of the game can be approximated in polynomial time, while the three-player entangled value is NP-hard?

We answer this question by showing that the same plane-vs-point low-degree test analyzed in [26] remains sound even when it is played with two, instead of three, entangled provers. As a consequence, we obtain the first non-trivial hardness results for the class $\mathsf{MIP}^*(2,1)$ of two-prover one-round entangled proof systems. (The best prior result is hardness for inverse-exponential completeness-soundness gap [13], which cannot be amplified by a polynomial-time verifier using e.g. parallel repetition.)

▶ **Theorem 1.** *The inclusion* $\mathsf{NEXP} \subseteq \mathsf{MIP}^*(2,1)$ *holds. Furthermore, it still holds when* $\mathsf{MIP}^*(2,1)$ *is restricted to one-round proof systems with constant answer size.*

Theorem 1 is obtained by scaling up a stronger NP-hardness result for two-player entangled projection games,[4] see Theorem 15 and Corollary 16 in Section 4.

Theorem 1 shows that allowing the provers to share entanglement does not weaken the power of two-prover one-round interactive proof systems. As mentioned earlier, entanglement may also have the effect of increasing the complexity of such proof systems, by allowing the verifier to implement protocols whose completeness can only be achieved by provers sharing entanglement. In fact, this is known to occur when the completeness-soundness gap is allowed to be exponentially small. In this regime, it was shown by [10] that the class QMIP^* of multi-prover interactive proof systems with a quantum verifier and messages contain $\mathsf{QMA}_{\mathsf{EXP}}$, the quantum analogue of NEXP, and subsequent works by Ji [16, 17] improved this result to show that MIP^* with exponentially small gap contains NEEXP (nondeterministic doubly-exponential time). However, it remained an open question whether a similar phenomenon occurs when the completeness-soundness gap is a constant.

In a subsequent work [21], building on the soundness analysis of the two-player low-degree test presented in this paper, we were able to answer (a version of) this question in the affirmative, showing the first constant-gap QMA-hardness results for entangled-player games. Specifically, we show that it is QMA-hard, under randomized reductions, to give a constant additive approximation to the maximum success probability of a players sharing entanglement in a multiplayer game specified in explicit form. The reduction in [21] yields a game with 7 players and one round of interaction. Interestingly, the analysis of this 7 player game, which uses the quantum error-correcting code framework of [10, 16], relies essentially on the soundness of the low-degree test with *two* entangled players. This is a further application of the techniques of this work, beyond the hardness for two-player games achieved in Theorem 1.

The main ingredient needed to obtain Theorem 1, and our main technical contribution, is a soundness analysis of the plane-vs-point low-degree test in the presence of two entangled provers. The analysis that we provide is both conceptually and technically simpler than the analysis in [26]. Although our proof relies on elementary reductions from [26], we present it in a modular way which, we hope, will make it more easily accessible, and more conveniently re-usable, than the proof in [26]. In the following subsection we describe the low-degree test and give a high-level overview of our analysis.

[4] The reduction proceeds in a standard way by using a succinctly represented instance of the 3-SAT problem as starting point; we omit the details.

Out of the two provers, choose one at random to be Alice and the other to be Bob.

1. Let d, m be integer and q a prime power given as input.
2. Select a random point $x \in \mathbb{F}_q^m$ and two random directions $y_1, y_2 \in \mathbb{F}_q^m$. If y_1 and y_2 are not linearly independent, accept; otherwise, let s be the plane spanned by the two lines parallel to y_1, y_2 passing through x.
3. Send s to Alice and x to Bob. Receive g, a specification of a degree-d polynomial restricted to s, from Alice, and $a \in \mathbb{F}_q$ from Bob.
4. Accept if and only if $g(x) = a$.

Figure 1 The (d, m, q)-low-degree test.

1.1 The low-degree test

We recall the "plane-vs-point" low-degree test from [26] in Figure 1. The test is essentially the same as the classical test from [22]. It asks one prover for the restriction of a low-degree m-variate polynomial g to a random two-dimensional subspace s of $\mathbb{F}_q^m$, where $\mathbb{F}_q$ is the finite field with q elements, q a prime power, and the other prover for the evaluation of g at a random $x \in s$; the prover's answers are checked for consistency.

Since the test treats both provers symmetrically, for the purposes of the soundness analysis we may reduce to the case where the provers share a permutation-invariant state and use the same collection of measurement operators. The following states the result of our analysis of the test. It extends Theorem 3.1 in [26] to the case of two provers.[5] In the theorem, we use the notation $\langle A, B\rangle_\Psi$ for $\langle \Psi | A \otimes B | \Psi\rangle$.

▶ **Theorem 2.** *There exists a $\delta = \mathrm{poly}(\varepsilon)$ and a constant $c > 0$ such that the following holds. Let $\varepsilon > 0$, m, d integers, and q a prime power such that $q \geq (dm/\varepsilon)^c$. For any strategy for the players using entangled state $|\Psi\rangle$ and projective measurements $\{A_s^r\}_r$ that succeeds in the (d, m, q)-low-degree test with probability at least $1 - \varepsilon$, there exists a POVM $\{S^g\}_g$, where g ranges over m-variate polynomials over $\mathbb{F}_q$ of total degree at most d, such that the following hold:*

1. *Approximate consistency with A:*
$$\mathop{\mathrm{E}}_{s} \sum_{g} \sum_{r \neq g|_s} \langle A_s^r, S^g\rangle_\Psi \leq \delta ,$$
where the expectation is over a random two-dimensional subspace s of $\mathbb{F}_q^m$, as chosen by the verifier in the test;
2. *Self-consistency:*
$$\sum_g \langle S^g, (\mathrm{Id} - S^g)\rangle_\Psi \leq \delta .$$

The proof of Theorem 2 follows the same structure as the proof of Theorem 3.1 in [26]. The proof is by induction on the number of variables m. The base case $m = 2$ is trivial, since there is a single subspace s, and the provers' associated POVM $\{A^r\}$ can directly play the role of $\{S^g\}$ in the theorem. Suppose then that the theorem is true for a value $(m-1)$ such that $m - 1 \geq 2$. To show that the theorem holds for m there are three main steps, which mirror the classical analysis of the low-degree test:

[5] The self-consistency condition is not explicitly stated in [26] but (as we will show) it follows easily from the proof.

1. (Section 6.3 of [26]) By the induction hypothesis, for every $(m-1)$-dimension hyperplane s in $\mathbb{F}_q^m$ there is a POVM $\{Q_s^g\}_g$ with outcomes g in the set of degree-d polynomials on s, such that on average over the choice of a uniformly random s and $x \in s$ the POVM $\{Q_s^g\}$ is consistent with $\{A_x^a\}$.

2. (Section 6.4 of [26]) For any $k \geq 1$, measurements $\{Q_s^g\}_g$ associated with k parallel subspaces $s_1, \ldots, s_k$ are "pasted" together to yield a combined measurement $\{Q_{(s_i)}^{(g_i)}\}$ that returns a k-tuple of degree-d polynomials g_i defined on s_i. This is proved by induction on k.

3. (Section 6.5 of [26]) Finally, taking k to be sufficiently large compared to d, the measurement $\{Q_{(s_i)}^{(g_i)}\}$ is consolidated into a single global measurement $\{S^g\}$ that satisfies the conclusion of the theorem for the m-variate case.

These three steps remain unchanged in the current proof. At only very few places in [26] is the presence of three provers used; in most cases this is only a matter of convenience and is easily avoided. For completeness, in Appendix A we explicitly list those places and how the use of three provers can be avoided.

As already mentioned the critical point in the proof where three provers, or rather the existence of three tensor factors in the provers' Hilbert space, is used, is to control the error increase throughout the induction. As shown by the analysis, if the measurements $\{Q_s^g\}$ produced by the induction hypothesis are δ-consistent with $\{A_x^a\}$, then the resulting S^g will be $O(\delta^c)$-consistent with the same $\{A_x^a\}$, for some constant $c < 1$. For poly-logarithmic m such an increase is unmanageable.

The key step in the analysis consists in establishing a "self-improvement lemma", which resets the consistency error to some constant baseline at each step of the induction. This is called the "consolidation procedure" in [26]. A similar self-improvement was already at the heart of Babai et al.'s proof of $\mathsf{MIP} \subseteq \mathsf{NEXP}$; variants thereof have found uses outside of complexity theory, such as in property testing.

Our main technical contribution is a simpler, self-contained proof of a variant of the consolidation procedure from [26] (stated as Proposition 5.8 in that paper), which applies to strategies with two provers only. The procedure shows that the consistency error sustained by any POVM, when measured against a structure called a "robust triple" in [26], can be automatically improved. Our variant is based on a simpler notion than the robust triples from [26], that we call "global consistency". We believe that our formulation of self-improvement, and its analysis (which crucially relies on semidefinite duality), should be of broad interest. At a high level, the result relies on a procedure that, given a collection of positive semidefinite operators $\{A_i\}$, identifies a measurement $\{T_i\}$, i.e. $T_i \geq 0$ and $\sum_i T_i = \mathrm{Id}$, that "optimally coincides" with the $\{A_i\}$ (see Lemma 13 for a precise formulation).

Throughout we assume familiarity with the notation and proof structure from [26], though we recall the most important notions in Section 2. In particular we formally define robust triples and global consistency, and show that the former notion implies the latter, so that our result can be directly used in lieu of Proposition 5.8 in the analysis of [26]. In Section 3 we prove our replacement for Proposition 5.8, Proposition 12. The proof of (the scaled-down version of) Theorem 1 follows from the analysis of the test using similar reductions as in [26]; we briefly explain how in Section 4.

2 Preliminaries

2.1 Notation

We use $\mathcal{H}$ to denote a finite-dimensional Hilbert space, and $\mathrm{L}(\mathcal{H})$ for the linear operators on $\mathcal{H}$. Subscripts $\mathcal{H}_A$, $\mathcal{H}_B$ indicate distinct spaces. For $|\Psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$ and $A \in \mathrm{L}(\mathcal{H}_A)$, $B \in \mathrm{L}(\mathcal{H}_B)$ we write $\langle A, B\rangle_\Psi = \langle \Psi | A \otimes B | \Psi\rangle$. Note that we do not conjugate A or B. Given two families of operators $\{A_x^a\}$ and $\{B_x^a\}$ on $\mathcal{H}_A$, where $x \in \mathcal{X}$ and $a \in \mathcal{A}$ range over finite sets, and $0 \le \delta \le 1$, we write $A_x^a \approx_\delta B_x^a$ for

$$\mathop{\mathrm{E}}_x \sum_a \langle (A_x^a - B_x^a)^2, \mathrm{Id}\rangle_\Psi \;=\; O(\delta)\;.$$

The expectation over x will usually be taken with respect to the uniform distribution. The distinction between taking an expectation (over x) or a summation (over a) will always be clear from context.

2.2 Measurements

Throughout, we consider a bipartite state $|\Psi\rangle \in \mathcal{H} \otimes \mathcal{H}$ assumed to be invariant under permutation of the two registers. All operators we consider act on the finite-dimensional space $\mathcal{H}$.

▶ **Definition 3.** A *sub-measurement* $\{M^a\}_a$ is a collection of positive semidefinite operators satisfying $M = \sum_a M^a \le \mathrm{Id}$. We say that a sub-measurement is η-complete if

$$\langle M, \mathrm{Id}\rangle_\Psi \ge 1 - \eta\;;$$

η is called the *completeness error*. If $M = \mathrm{Id}$ then we say that $\{M^a\}_a$ is a measurement, in which case the completeness error is zero.[6]

The following definition appears in [26].

▶ **Definition 4.** Let $\mathcal{X}$ and $\mathcal{A}$ be finite sets. Let $\{M_x^a\}_a$ be a family of sub-measurements indexed by $x \in \mathcal{X}$ and with outcomes $a \in \mathcal{A}$. For each x, let $M_x = \sum_a M_x^a$. We say that $\{M_x^a\}$ is

- ε*-self-consistent* if

$$\mathop{\mathrm{E}}_x \sum_{a \ne a'} \langle M_x^a, M_x^{a'}\rangle_\Psi \le \varepsilon\;,$$

- γ*-projective* if

$$\mathop{\mathrm{E}}_x \langle M_x, (\mathrm{Id} - M_x)\rangle_\Psi \le \gamma\;.$$

- Let $\{T^g\}$ be a sub-measurement with outcomes in the set of all functions $g : \mathcal{X} \to \mathcal{A}$. We say that $\{M_x^a\}$ and $\{T^g\}$ are δ*-consistent* if

$$\mathop{\mathrm{E}}_x \sum_{g,a:\, a \ne g(x)} \langle T^g, M_x^a\rangle_\Psi \le \delta\;.$$

[6] The converse does not necessarily hold, as $|\Psi\rangle$ may not have full support.

We consider families of functions such that distinct functions have few points of intersection. The following definition is the reformulation of the definition of an error-correcting code, that is adapted to our notation using functions (where the codeword associated to a function is the evaluation table of the function, and vice-versa).

▶ **Definition 5.** Let $\mathcal{X}$ and $\mathcal{A}$ be finite sets, $\mathcal{G}$ a set of functions from $\mathcal{X}$ to $\mathcal{A}$, and $0 \leq \kappa \leq 1$. We say that $(\mathcal{X}, \mathcal{A}, \mathcal{G})$ is κ-structured if for any two distinct $g, g' \in \mathcal{G}$,

$$\Pr_{x\in\mathcal{X}}\big(g(x) = g'(x)\big) \leq \kappa \;,$$

where the probability is taken under the uniform distribution on $\mathcal{X}$.

The following lemma states useful properties of consistency.

▶ **Lemma 6.** *Let $(\mathcal{X}, \mathcal{A}, \mathcal{G})$ be κ-structured. Let $\{A_x^a\}_{a\in\mathcal{A}}$ be a family of measurements indexed by $x \in \mathcal{X}$ that is ε-self-consistent. Let $\{T^g\}_{g\in\mathcal{G}}$ be a sub-measurement that is δ-consistent with $\{A_x^a\}$. Then*

- *$\{T^g\}$ is δ'-self-consistent, for $\delta' = O(\sqrt{\varepsilon} + \sqrt{\delta} + \kappa)$;*
- *Let $T = \sum_g T^g$, and suppose $\{T^g\}$ is γ-projective. Then*

$$T A_x^a \approx_{\sqrt{\varepsilon}+\sqrt{\delta}+\gamma+\kappa} A_x^a T \;.$$

Proof. We sketch the proof. For the first item,

$$\begin{aligned}
\sum_{g\neq g'} \langle T^g, T^{g'}\rangle_\Psi &= \mathop{\mathrm{E}}_x \sum_a \sum_{g\neq g'} \langle T^g, T^{g'} A_x^a\rangle_\Psi\\
&\approx_{\sqrt{\delta}} \mathop{\mathrm{E}}_x \sum_{g\neq g'} \langle T^g, T^{g'} A_x^{g(x)}\rangle_\Psi\\
&\approx_{\sqrt{\varepsilon}} \mathop{\mathrm{E}}_x \sum_{g\neq g'} \langle T^g A_x^{g(x)}, T^{g'}\rangle_\Psi\\
&\approx_{\sqrt{\delta}} \mathop{\mathrm{E}}_x \sum_{g\neq g'} \mathbf{1}_{g(x)=g'(x)} \langle T^g A_x^{g(x)}, T^{g'}\rangle_\Psi\\
&\approx_{\sqrt{\delta}} \mathop{\mathrm{E}}_x \sum_{g\neq g'} \mathbf{1}_{g(x)=g'(x)} \langle T^g, T^{g'}\rangle_\Psi\\
&\leq \kappa \;.
\end{aligned}$$

For the second item, it suffices to lower bound

$$\begin{aligned}
\mathop{\mathrm{E}}_x \sum_a \langle T A_x^a T A_x^a, \mathrm{Id}\rangle_\Psi &\approx_{\sqrt{\varepsilon}} \mathop{\mathrm{E}}_x \sum_a \sum_g \langle T A_x^a T^g, A_x^a\rangle_\Psi\\
&\approx_{\sqrt{\delta}} \mathop{\mathrm{E}}_x \sum_a \sum_g \langle T A_x^a T^g, A_x^{g(x)}\rangle_\Psi\\
&\approx_{\sqrt{\delta}} \mathop{\mathrm{E}}_x \sum_{a,a'} \sum_g \langle T A_x^a T^g, A_x^{a'}\rangle_\Psi\\
&= \langle T^2, \mathrm{Id}\rangle_\Psi \;.
\end{aligned}$$

The claimed bound then follows by expanding $\mathrm{E}_x \sum_a (T A_x^a - A_x^a T)^2$ and regrouping terms. ◀

2.3 Global consistency

The analysis of the low-degree test amounts to arguing that a set of measurement operators which produce outcomes that are locally consistent can be combined into a single measurement which returns a global object consistent with each of the local measurements: it is possible to recombine local views. In [26] the notion of local consistency used is called a "robust triple". For convenience we recall the definition.

▶ **Definition 7** (Definition 5.2 in [26]). Let $G = (V, E)$ be a graph, S a finite set, $\mathcal{G} \subseteq \{g : V \to S\}$ a set of functions and for every $v \in V$, $\{A_v^a\}_{a\in S}$ a measurement with outcomes in S. Given $\delta > 0$ and $0 < \mu \leq 1$, we say that $(G, \{A_v^a\}, \mathcal{G})$ is a (δ, μ)-robust triple if:

1. (self-consistency) The family of measurements $\{A_v^a\}$ is δ-self-consistent;
2. (small intersection) $(V, S, \mathcal{G})$ is δ-structured;
3. (stability) For any sub-measurement $\{R^g\}_{g\in\mathcal{G}}$ it holds that
$$\mathop{\mathrm{E}}_{v\in V}\ \mathop{\mathrm{E}}_{v'\in N(v)} \sum_g \langle R^g,\, (A_v^{g(v)} - A_{v'}^{g(v')})^2\rangle_\Psi \leq \delta\ ,$$
where $N(v)$ is the set of neighbors of v in G;
4. (expansion) G has mixing time $O(\mu^{-1})$. Precisely, if for any $v \in V$ we let $p_k(v)$ denote the distribution on V that results from starting a k-step random walk at v, then for any $\delta > 0$ and some $k = O(\log(1/\delta)\log(1/\mu))$ it holds that $\mathrm{E}_{v\in V} \|p_k(v) - |V|^{-1}\|_1 \leq \delta$.

We observe that the only way in which items 3. and 4. from the definition are used for the self-improvement results is through [26, Claim 5.3], which states the following.

▶ **Claim 8** (Claim 5.3 in [26]). *Suppose $(G, A, \mathcal{G})_\Psi$ is a (δ, μ)-robust triple. Then there exists a $\delta' = O\big(\delta^{1/2}\log^2(1/\delta)\log^2(1/\mu)\big)$ such that for any sub-measurement $\{R^g\}_{g\in\mathcal{G}}$,*
$$\sum_g \langle R^g, A^g - (A^g)^2\rangle_\Psi \,\leq\, \delta'\ , \tag{1}$$
where $A^g = E_{v\in V} A_v^{g(v)}$.

It is more direct, and more general, to use condition (1) directly as part of the definition, as this allows us to set aside any notion of an expanding graph.

▶ **Definition 9.** Let $(\mathcal{X}, \mathcal{A}, \mathcal{G})$ be κ-structured. Let $\{A_x^a\}_{a\in\mathcal{A}}$ be a family of measurements indexed by $x \in \mathcal{X}$ and with outcomes $a \in \mathcal{A}$. For $g \in \mathcal{G}$, let $A^g = \mathrm{E}_x A_x^{g(x)}$. Let $|\Psi\rangle$ be a permutation-invariant bipartite state. For $0 \leq \varepsilon, \delta \leq 1$ we say that $(\{A_x^a\}, \mathcal{G})$ is (ε, δ)-globally consistent on $|\Psi\rangle$ if:

1. $\kappa = O(\varepsilon)$;
2. The family $\{A_x^a\}$ is ε-self-consistent;
3. There exists a positive semidefinite operator Z such that
$$\forall g \in \mathcal{G},\ 0 \leq A^g - (A^g)^2 \leq Z, \qquad \text{and} \qquad \langle Z, \mathrm{Id}\rangle_\Psi \leq \delta.$$

It is not hard to verify that condition 3. in the definition is equivalent to (1). This can be seen by writing the bound δ in the condition as the optimum of a semidefinite program, and taking the dual. This is done in a similar way to the analysis of the semidefinite program (2). The only difference is that the latter considers consistency when the state $|\Psi\rangle$ is maximally entangled. Formally, we have the following lemma.

▶ **Lemma 10.** *Let $|\Psi\rangle \in \mathcal{H}\otimes\mathcal{H}$ be a state invariant under permutation of its two registers, such that the reduced density of $|\Psi\rangle$ on either register has full support. Let $\{A_i\}$ be a family of positive semidefinite operators on $\mathcal{H}$ with $A_i \leq \mathrm{Id}$ for all i. Then the following primal and dual semidefinite program satisfy strong duality, and hence have the same optimum value:*

Primal SDP

$$\begin{aligned} \sup \quad & \sum_i \langle T_i, A_i\rangle_\Psi \\ \textit{s.t.} \quad & T_i \geq 0 \qquad \forall i\,, \\ & \sum_i T_i \leq \mathrm{Id}\ . \end{aligned}$$

Dual SDP

$$\begin{aligned} \inf \quad & \langle Z, \mathrm{Id}\rangle_\Psi \\ \textit{s.t.} \quad & Z \geq A_i \qquad \forall i\,, \\ & Z \geq 0\,. \end{aligned}$$

Proof. Both the primal and dual are strictly feasible, as can be seen by taking e.g. $T_i \propto \mathrm{Id}$ such that $\sum_i T_i = \mathrm{Id}/2$, and $Z = 2\,\mathrm{Id}$. ◀

Taking A_i in Lemma 10 to equal $A^g-(A^g)^2$, the primal value being less than δ' is equivalent to (1), while the dual value being less than δ' is equivalent to item 3. in Definition 9.

For later use we note that self-consistency of $\{A^a_x\}$ implies self-consistency of the operators A^g introduced in Definition 9, in the following sense.

▶ **Lemma 11.** *Let $\{A^a_x\}$ be a family of measurements that is ε-self-consistent. Then for any sub-measurement $\{R^g\}$,*

$$\sum_g \langle A^g,\, R^g\rangle_\Psi \approx_{\sqrt{\varepsilon}} \sum_g \langle \mathrm{Id},\, R^g A^g\rangle_\Psi .$$

Proof. Write

$$\begin{aligned} \sum_g \langle A^g,\, R^g\rangle_\Psi &= \sum_g \operatorname*{E}_x \langle A^{g(x)}_x,\, R^g\rangle_\Psi \\ &= \sum_{g,a} \operatorname*{E}_x \langle A^{g(x)}_x,\, R^g A^a_x\rangle_\Psi \\ &\approx_{\sqrt{\varepsilon}} \sum_g \operatorname*{E}_x \langle A^{g(x)}_x,\, R^g A^{g(x)}_x\rangle_\Psi \\ &\approx_{\sqrt{\varepsilon}} \sum_g \operatorname*{E}_x \langle \mathrm{Id},\, R^g A^{g(x)}_x\rangle_\Psi\ . \end{aligned}$$ ◀

3 Self-improvement with two provers

The main result on self-improvement from [26] is stated as Proposition 5.8 in that paper. Our main technical result, Proposition 12 below, improves upon Proposition 5.8 in the following respects:

- Proposition 12 allows performing self-improvement with two provers only;
- Proposition 12 only requires the notion of consistency introduced in Definition 9, which as argued in Section 2.3 is less restrictive than the notion of robust triple used in [26];

- The proof of Proposition 12 is simpler and yields better parameters.

We state the proposition and give its proof here. In Section 4 we show how the proposition is used to obtain the hardness results.

▶ **Proposition 12.** *There exists universal constants $\varepsilon_0, \delta_0, t_0 > 0$ such that the following holds. Let $(\mathcal{X}, \mathcal{A}, \mathcal{G})$ be κ-structured. Let $\{A_x^a\}_{a\in\mathcal{A}}$ be a family of measurements indexed by $x \in \mathcal{X}$, and $|\Psi\rangle$ a bipartite permutation-invariant state. Suppose that the following conditions hold:*

1. *$(\{A_x^a\}, \mathcal{G})$ is (ε, δ)-globally consistent on $|\Psi\rangle$, for some $0 \le \varepsilon \le \varepsilon_0$, $0 \le \delta \le \delta_0$;*
2. *There exists a function $t = t(\varepsilon', \delta')$ and $\varepsilon_0', \delta_0' > 0$ such that for any $0 \le \varepsilon' \le \varepsilon_0'$ and $0 \le \delta' \le \delta_0'$ it holds that $t(\varepsilon', \delta') \le t_0$, and such that the following holds. For any (ε', δ') and state $|\Phi\rangle$ such that $(\{A_x^a\}, \mathcal{G})$ is (ε', δ')-globally consistent on $|\Phi\rangle$, there exists a measurement $\{Q^g\}_{g\in\mathcal{G}}$ that is $t(\varepsilon', \delta')$-consistent with $\{A_x^a\}$.*

Then there exists a measurement $\{R^g\}_{g\in\mathcal{G}}$ that is δ'-consistent with $\{A_x^a\}$, for some $\delta' = O(\sqrt{r(\varepsilon, \delta)})$, where $r(\varepsilon, \delta)$ is the function defined in Lemma 13.

The key "improvement" provided by the proposition is that, while the function t is only assumed to be bounded by a fixed constant for sufficiently small values of the arguments, the proposition returns a measurement $\{R^g\}$ that has an explicit consistency δ' with $\{A_x^a\}$, where δ' is polynomial in ε and δ, irrespective of t (indeed t need not approach 0 as ε, δ approach 0).

We note that, in our language, [26, Proposition 5.8] considers a family of globally consistent pairs $(\{A_{t,x}^a\}, \mathcal{G}_t)$, parametrized by some finite set $t \in T$, and makes both the assumptions and the conclusions of Proposition 12 in an averaged sense, for uniformly random $t \in T$. For simplicity we state and prove the proposition for $|T| = 1$. The case of general T is needed for the inductive application of the Proposition towards the proof of Theorem 2. We sketched the inductive step in the introduction. We refer to [26] for details of the derivation of Theorem 2 from Proposition 12, which is identical to the derivation of [26, Theorem 3.1] from [26, Proposition 5.8], up to minor modifications that we review in Appendix A.

The main step in the proof of the proposition is provided by the following lemma, which is analogous to [26, Claim 5.4]. The semidefinite program considered in the proof of the lemma, and its analysis, are our main points of departure from the proof in [26]. Indeed, the proof of an upper bound on the completeness error of the sub-measurement $\{S^g\}$ constructed in the proof of the lemma is the main point where the existence of a three-fold tensor product decomposition of the Hilbert space is most crucially used in [26].

▶ **Lemma 13.** *There exists a function $r(\varepsilon, \delta) = O(\sqrt{\varepsilon} + \sqrt{\delta})$ such that the following holds for all $0 \le \varepsilon, \delta, \eta \le 1$. Let $(\mathcal{X}, \mathcal{A}, \mathcal{G})$ be κ-structured. Let $\{A_x^a\}_{a\in\mathcal{A}}$ be a family of measurements indexed by $x \in \mathcal{X}$. Let $|\Psi\rangle$ be a permutation-invariant bipartite state and assume $(\{A_x^a\}, \mathcal{G})$ are (ε, δ)-globally consistent on $|\Psi\rangle$. Let $\{Q^g\}_{g\in\mathcal{G}}$ be a sub-measurement that is η-consistent with $\{A_x^a\}$ on $|\Psi\rangle$. Then there exists a sub-measurement $\{S^g\}$ that is $r(\varepsilon, \delta)$-consistent with $\{A_x^a\}$ and projective and has completeness error*

$$\langle \mathrm{Id} - S, \mathrm{Id}\rangle_\Psi \;\le\; \langle \mathrm{Id} - Q, \mathrm{Id}\rangle_\Psi + \eta + r(\varepsilon, \delta) \,.$$

Proof. For $g \in \mathcal{G}$, let $A^g = \mathrm{E}_x\, A_x^{g(x)}$. We consider the following primal and dual semidefinite program, obtained from the semidefinite program in Lemma 10 by setting A_i to A^g and formally replacing the state $|\Psi\rangle$ appearing in the SDP by the maximally entangled state[7].

[7] Note that we are not assuming that the state $|\Psi\rangle$ appearing in the hypothesis of Lemma 13 is maximally entangled. The purpose of defining the SDP (2) without reference to the state $|\Psi\rangle$ is to make the resulting complementary slackness conditions (4) easier to work with.

The primal becomes

$$\begin{aligned}\omega \;=\; \sup \quad & \sum_g \mathrm{Tr}(T^g A^g) \\ \text{s.t.} \quad & T^g \geq 0 \qquad \forall g \in \mathcal{G}\,, \\ & \sum_g T^g \leq \mathrm{Id}\,,\end{aligned} \tag{2}$$

and the dual

$$\begin{aligned}\min \quad & \mathrm{Tr}(Z) \\ \text{s.t.} \quad & Z \geq A^g \qquad \forall g \in \mathcal{G}\,, \\ & Z \geq 0\,.\end{aligned} \tag{3}$$

As shown in Lemma 10 both the primal and dual are strictly feasible, so that strong duality holds. Let $\{T^g\}$ be an optimal primal solution. Without loss of generality, $\sum_g T^g = \mathrm{Id}$, as any solution such that $(\mathrm{Id} - \sum_g T^g)A^{g'} \neq 0$ for any g' is clearly not optimal. The complementary slackness conditions, which follow from the KKT conditions for optimality, immediately imply

$$T^g Z \;=\; T^g A^g \qquad \forall g \in \mathcal{G}\,. \tag{4}$$

For each $g \in \mathcal{G}$ let

$$S^g = \mathop{\mathrm{E}}_x A_x^{g(x)} T^g A_x^{g(x)}\,.$$

Then $\{S^g\}$ is a sub-measurement. We show that S^g satisfies the desired consistency, projectivity and completeness properties.

(i) *Consistency:* We have that

$$\mathop{\mathrm{E}}_x \sum_g \sum_{a \neq g(x)} \langle S^g,\, A_x^a \rangle_\Psi = \sum_g \langle S^g,\, (\mathrm{Id} - A^g)\rangle_\Psi\,.$$

Using self-consistency of $\{A_x^a\}$,

$$\begin{aligned}\sum_g \langle S^g,\, \mathrm{Id}\rangle_\Psi &= \mathop{\mathrm{E}}_x \sum_g \langle A_x^{g(x)} T^g A_x^{g(x)},\, \mathrm{Id}\rangle_\Psi \\ &\approx_{\sqrt{\varepsilon}} \mathop{\mathrm{E}}_x \sum_g \langle T^g,\, A_x^{g(x)}\rangle_\Psi \\ &= \sum_g \langle T^g, A^g\rangle_\Psi\,.\end{aligned} \tag{5}$$

Similarly,

$$\begin{aligned}\sum_g \langle S^g,\, A^g\rangle_\Psi &= \mathop{\mathrm{E}}_x \sum_g \langle A_x^{g(x)} T^g A_x^{g(x)},\, A^g\rangle_\Psi \\ &\approx_{\sqrt{\varepsilon}} \mathop{\mathrm{E}}_x \sum_g \langle T^g,\, A_x^{g(x)} A^g A_x^{g(x)}\rangle_\Psi\,.\end{aligned} \tag{6}$$

Using the Cauchy-Schwarz inequality,

$$\begin{aligned}\mathop{\mathrm{E}}_x \sum_g \langle T^g,\, \big(A^g - A_x^{g(x)}\big) A^g A_x^{g(x)}\rangle_\Psi &\le \Big(\mathop{\mathrm{E}}_x \sum_g \langle T^g,\, A_x^{g(x)} (A^g)^2 A_x^{g(x)}\rangle_\Psi\Big)^{\frac{1}{2}}\\ &\qquad \cdot \Big(\mathop{\mathrm{E}}_x \sum_g \langle T^g,\, \big(A^g - A_x^{g(x)}\big)^2\rangle_\Psi\Big)^{\frac{1}{2}}\\ &\le \Big(\sum_g \langle T^g,\, \big(A^g - (A^g)^2\big)\rangle_\Psi\Big)^{\frac{1}{2}}\\ &\le \sqrt{\delta}\,, \end{aligned} \tag{7}$$

where the second inequality uses $A_x^{g(x)}(A^g)^2 A_x^{g(x)} \le \mathrm{Id}$ for the first term, and expands the square and uses $(A_x^{g(x)})^2 \le A_x^{g(x)}$ for the second term, and the last inequality follows from item 3. in the definition of globally consistent. Combined with (5) and (6), we have shown

$$\mathop{\mathrm{E}}_x \sum_g \sum_{a \ne g(x)} \langle S^g,\, A_x^a\rangle_\Psi \approx_{\sqrt{\varepsilon}+\sqrt{\delta}} \sum_g \langle T^g,\, (A^g - (A^g)^3)\rangle_\Psi\ . \tag{8}$$

Writing

$$\begin{aligned} A^g - (A^g)^3 &= A^g - (A^g)^2 + \sqrt{A^g}\big(A^g - (A^g)^2\big)\sqrt{A^g}\\ &\le 2\big(A^g - (A^g)^2\big)\ , \end{aligned}$$

since all terms commute and $(A^g)^2 \le A^g \le \mathrm{Id}$, using item 3. in the definition of globally consistent the right-hand side of (8) is at most 2δ.

(ii) *Completeness:*

$$\begin{aligned}\sum_g \langle S^g,\, \mathrm{Id}\rangle_\Psi &= \mathop{\mathrm{E}}_x \sum_g \langle A_x^{g(x)} T^g A_x^{g(x)},\, \mathrm{Id}\rangle_\Psi\\ &\approx_{\sqrt{\varepsilon}} \mathop{\mathrm{E}}_x \sum_g \langle T^g, A_x^{g(x)}\rangle_\Psi\\ &\approx_{\sqrt{\varepsilon}} \sum_g \langle T^g A^g,\, \mathrm{Id}\rangle_\Psi\\ &= \sum_g \langle T^g Z,\, \mathrm{Id}\rangle_\Psi\\ &= \langle Z,\, \mathrm{Id}\rangle_\Psi\ , \end{aligned}$$

where the third line uses Lemma 11 and the penultimate equality follows from (4), and for the last we used $\sum_g T^g = \mathrm{Id}$. We establish a lower bound on this last expression by introducing $\{Q^g\}$:

$$\begin{aligned}\langle Q,\, \mathrm{Id}\rangle_\Psi - \eta &\le \sum_g \langle Q^g,\, A^g\rangle_\Psi\\ &\le \sum_g \langle Q^g,\, Z\rangle_\Psi\\ &\le \langle \mathrm{Id},\, Z\rangle_\Psi\ , \end{aligned}$$

where the second inequality uses the dual constraint (3), and the third uses $\sum_g Q^g \le \mathrm{Id}$. It follows that

$$\sum_g \langle S^g,\, \mathrm{Id}\rangle_\Psi \ \ge\ \langle Q,\, \mathrm{Id}\rangle_\Psi - \eta - O\big(\sqrt{\varepsilon}\big).$$

(iii) *Projectivity:* By proceeding exactly as in (7), we can show

$$\begin{aligned}\langle S,\, S\rangle_\Psi &= \sum_g \operatorname*{E}_x \langle A_x^{g(x)} T^g A_x^{g(x)},\, S\rangle_\Psi\\ &\approx_{\sqrt{\varepsilon}+\sqrt{\delta}} \sum_g \langle A^g T^g A^g,\, S\rangle_\Psi\\ &= \sum_{g,g'} \operatorname*{E}_x \langle A^g T^g A^g,\, A_x^{g'(x)} T^{g'} A_x^{g'(x)}\rangle_\Psi\\ &\approx_{\sqrt{\varepsilon}+\sqrt{\delta}} \sum_{g,g'} \operatorname*{E}_x \langle A_x^{g(x)} T^g A_x^{g(x)},\, A_x^{g'(x)} T^{g'} A_x^{g'(x)}\rangle_\Psi\ .\end{aligned}$$

Using self-consistency of $\{A_x^a\}$, from the above we get

$$\begin{aligned}\langle S,\, S\rangle_\Psi &\approx_{\sqrt{\varepsilon}+\sqrt{\delta}} \sum_{g,g'} \operatorname*{E}_x \langle T^g A_x^{g(x)},\, A_x^{g'(x)} T^{g'} A_x^{g'(x)}\rangle_\Psi\\ &\approx_{\sqrt{\varepsilon}+\sqrt{\delta}} \sum_{g,g'} \langle T^g A^g,\, S^{g'}\rangle_\Psi\\ &= \langle Z,\, S\rangle_\Psi\ ,\end{aligned}\tag{9}$$

where the second line again uses similar arguments as (7) and the last line uses (4) and $\sum_g T^g = \mathrm{Id}$. Using the dual constraint (3), we deduce

$$\begin{aligned}\langle S,\, S\rangle_\Psi &\geq \sum_g \langle A^g,\, S^g\rangle_\Psi - O(\sqrt{\varepsilon}+\sqrt{\delta})\\ &\approx_{\sqrt{\varepsilon}+\sqrt{\delta}} \langle S, \mathrm{Id}\rangle_\Psi\ ,\end{aligned}$$

where the second line follows from consistency of $\{S^g\}$ and $\{A_x^a\}$ shown in item (i). ◀

Based on Lemma 13, we give the proof of Proposition 12.

Proof of Proposition 12. Let ε, δ be as in condition 1., and $\{Q^g\}$ be the measurement whose existence follows from condition 2. in the proposition, when $|\Phi\rangle = |\Psi\rangle$ and $\varepsilon', \delta' = \varepsilon, \delta$. By applying Lemma 13 to the state $|\Psi\rangle$ and measurements $\{A_x^a\}$ and $\{Q^g\}$ we obtain a sub-measurement $\{S^g\}$ that is $\xi = r(\varepsilon, \delta)$-projective and consistent with $\{A_x^a\}$. Among all sub-measurements that are ξ-projective and consistent with $\{A_x^a\}$, let $\{T^g\}$ be one that minimizes the completeness error $\theta = \langle \mathrm{Id} - T, \mathrm{Id}\rangle$. Provided ε_0, δ_0 are small enough we may assume $\theta \leq t(\varepsilon, \delta) + r(\varepsilon, \delta) \leq 1/4$. If $\theta = 0$ the measurement T is perfectly complete, and we are done as we can take the measurements R^g in the conclusion of the proposition to be equal to T^g. So, for the rest of the proof, we can assume that $\theta > 0$. To complete the proof we need to prove a better upper bound on θ. Towards this, introduce a state

$$|\Phi\rangle = \frac{|\tilde{\Phi}\rangle}{\||\tilde{\Phi}\rangle\|}, \qquad \text{where} \qquad |\tilde{\Phi}\rangle = (\mathrm{Id} - T) \otimes (\mathrm{Id} - T)|\Psi\rangle\ .$$

Given the assumption that $\theta > 0$, it follows that $\||\tilde{\Phi}\rangle\| > 0$, and hence this state is well defined. Moreover we can estimate the norm of $|\tilde{\Phi}\rangle$ as follows:

$$\begin{aligned}\||\tilde{\Phi}\rangle\|^2 &= \langle (\mathrm{Id} - T)^2,\, (\mathrm{Id} - T)^2\rangle_\Psi\\ &= \langle \mathrm{Id} - 2T + T^2, \mathrm{Id} - 2T + T^2\rangle_\Psi\\ &= 1 - 4\langle T, \mathrm{Id}\rangle_\Psi + 4\langle T, T\rangle_\Psi + 2\langle T^2, \mathrm{Id}\rangle_\Psi - 4\langle T^2, T\rangle_\Psi + \langle T^2, T^2\rangle_\Psi\\ &= 1 - 4\langle T, (\mathrm{Id} - T)\rangle_\Psi + 2\langle T^2, (\mathrm{Id} - T)\rangle_\Psi - \langle T^2, T(\mathrm{Id} - T)\rangle_\Psi - \langle T^2, T\rangle_\Psi\\ &\approx_{\sqrt{\xi}} 1 - \langle T,\, T^2\rangle_\Psi\\ &\approx_{\sqrt{\xi}} 1 - \langle T, \mathrm{Id}\rangle_\Psi\ ,\end{aligned}\tag{10}$$

where the last two approximations use the projectivity assumption on T.

▶ **Claim 14.** *There are $\varepsilon' = O(\varepsilon + \sqrt{\xi})$ and $\delta' = O(\delta + \sqrt{\xi})$ such that $(\{A_x^a\}, \mathcal{G})$ is (ε', δ')-globally consistent on $|\Phi\rangle$.*

Proof. We verify the properties in Definition 9. Item 1. is automatic. For item 2., self-consistency of $\{A_x^a\}$ on $|\Phi\rangle$, write

$$\begin{aligned}\mathop{\mathrm{E}}_x \sum_a \langle A_x^a,\, A_x^a\rangle_{\tilde{\Phi}} &= \mathop{\mathrm{E}}_x \sum_a \langle A_x^a(\mathrm{Id} - T) - TA_x^a(\mathrm{Id} - T),\, (\mathrm{Id} - T)A_x^a - (\mathrm{Id} - T)A_x^a T\rangle_\Psi \\ &\approx_{\sqrt{\xi}} \mathop{\mathrm{E}}_x \sum_a \langle A_x^a,\, A_x^a\rangle_\Psi - \langle T,\, A_x^a\rangle_\Psi + \langle A_x^a(\mathrm{Id} - T),\, T\rangle_\Psi \\ &\approx_{\sqrt{\xi}} 1 - \varepsilon - \langle T,\, \mathrm{Id}\rangle_\Psi \ .\end{aligned}$$

Together with (10), it follows that $\{A_x^a\}$ is ε'-self-consistent on $|\Phi\rangle$, for some $\varepsilon' = O(\varepsilon + \sqrt{\xi})$. For item 3. in the definition, let Z be such that $A^g - (A^g)^2 \leq Z$ for all $g \in \mathcal{G}$, and $\langle Z, \mathrm{Id}\rangle_\Psi \leq \delta$. Then

$$\begin{aligned}\langle Z, \mathrm{Id}\rangle_{\tilde{\Phi}} &\approx_{\sqrt{\xi}} \langle Z, (\mathrm{Id} - T)\rangle_\Psi \\ &\leq \delta \ ,\end{aligned}$$

and the property follows using (10). ◀

Applying condition 2. in the proposition to $|\Phi\rangle$ and $(\{A_x^a\}, \mathcal{G})$ we obtain a measurement $\{Q^g\}$ that is $\xi' = t(\varepsilon', \delta')$-projective and consistent with $\{A_x^a\}$ on $|\Phi\rangle$. Define a sub-measurement $\{R^g\}$ by

$$R^g := TT^gT + (1 - T)Q^g(1 - T) \ .$$

The completeness of this measurement on $|\Psi\rangle$ is

$$\begin{aligned}\langle R,\, \mathrm{Id}\rangle_\Psi &= \langle T^3,\, \mathrm{Id}\rangle_\Psi + \langle (1 - T)^2,\, \mathrm{Id}\rangle_\Psi \\ &\approx_{\sqrt{\xi}} 1 \ ,\end{aligned} \tag{11}$$

since

$$\langle T^3,\, \mathrm{Id}\rangle_\Psi \ \approx_{\sqrt{\xi}} \ \langle T^2,\, \mathrm{Id}\rangle_\Psi \ \approx_{\sqrt{\xi}} \ \langle T,\, \mathrm{Id}\rangle_\Psi \ .$$

To evaluate consistency with $\{A_x^a\}$,

$$\begin{aligned}&\mathop{\mathrm{E}}_x \sum_g \sum_{a \neq g(x)} \langle R^g, A_x^a\rangle_\Psi \\ &\quad= \mathop{\mathrm{E}}_x \sum_g \sum_{a \neq g(x)} \big(\langle TT^gT,\, A_x^a\rangle_\Psi + \langle (1 - T)Q^g(1 - T),\, A_x^a\rangle_\Psi\big) \\ &\quad\approx_{\sqrt{\varepsilon} + \sqrt{\xi} + \kappa} \mathop{\mathrm{E}}_x \sum_g \sum_{a \neq g(x)} \big(\langle TT^g,\, TA_x^a\rangle_\Psi + \langle (1 - T)Q^g,\, (1 - T)A_x^a\rangle_\Psi\big) \\ &\quad= O(\sqrt{\xi}) + O(\sqrt{\xi'})\big\||\tilde{\Phi}\rangle\big\|^2 \ ,\end{aligned}$$

where the second line uses the second item in Lemma 6 and the last $\varepsilon = O(\xi)$, given the definition of the function r. Using (11), if we complete $\{R^g\}$ into a measurement $\{\tilde{R}^g\}$ by adding an arbitrary term, the latter will have consistency $\delta'' = O(\sqrt{\xi}) + O(\sqrt{\xi'})\||\tilde{\Phi}\rangle\|^2$ with

$\{A^a_x\}$. Applying Lemma 13 yields a sub-measurement $\{V^g\}$ that is $\xi = r(\varepsilon,\delta)$-projective and consistent with $\{A^a_x\}$, and for which

$$\langle(\mathrm{Id}-V),\, \mathrm{Id}\rangle_\Psi \;=\; O(\sqrt{\xi}) + O(\sqrt{\xi'})\||\tilde{\Phi}\rangle\|^2.$$

Recall that by assumption, $\{T^g\}$ is the most complete measurement that is ξ-projective and consistent with A_x. Hence, $\langle(\mathrm{Id}-V),\mathrm{Id}\rangle_\Psi \geq \langle(\mathrm{Id}-T),\mathrm{Id}\rangle_\Psi$, so that

$$\theta \;\leq\; O(\sqrt{\xi}) + O(\sqrt{\xi'})(\theta + O(\sqrt{\xi})).$$

Provided ε,δ are small enough that $O(\sqrt{\xi'}) = O(\sqrt{t(\varepsilon',\delta')})$, with ε',δ' as in Claim 14, is at most $1/4$, as can be assumed from the assumed upper bound $t(\varepsilon',\delta') \leq t_0$ for $\varepsilon' \leq \varepsilon_0$ and $\delta' \leq \delta_0$ provided t_0 is a small enough universal constant, we have obtained $\theta = O(\sqrt{\xi}) = O(\sqrt{r(\varepsilon,\delta)})$, as claimed. ◀

4 NP-hardness for two-player entangled games

Based on the result of the analysis of the low-degree test stated in Theorem 2 and following the same sequence of reductions — composition of the low-degree test with itself, to reduce answer size, and combination with the 3-SAT test — as in [26] we obtain the following analogue of [26, Theorem 4.1], which establishes NP-hardness for games with $\mathrm{poly}(\log\log n)$-bit answers.

▶ **Theorem 15.** *There is an $\varepsilon > 0$ such that the following holds. Given a 2-player game G in explicit form, it is NP-hard to distinguish between $\omega(G) = 1$ and $\omega^*(G) \leq 1-\varepsilon$. Furthermore, the problem is still NP-hard when restricting to games G of size n that are projection games for which questions and answers can be specified using $O(\log n)$ bits and $\mathrm{poly}(\log\log n)$ bits respectively.*

In [26] this result is improved to obtain hardness for games with constant-bit answers by reducing the 3-SAT test, on which the proof of Theorem 15 is based, to the three-player QUADEQ test for testing satisfiability of a system of quadratic equations in binary variables. This amounts to composing a PCP based on low-degree polynomials with the "exponential PCP" based on the three-query linearity test of [5], and yields hardness for three-player games with binary answers. The same steps can be completed with two players only by using the technique of oracularization to transform the QUADEQ and linearity tests into two-player games. The idea of oracularization is that for every triple of questions (q_1,q_2,q_3) to be sent to the three players in the original test, the verifier sends the entire triple to a single player, Alice, and receives a triple of answers. The verifier also sends a randomly selected question from the triple to a second player, Bob. The verifier accepts if and only if Bob's answer is consistent with Alice's, and the triple of answers provided by Alice would have been accepted in the original test. For concreteness, we summarize the oracularized QUADEQ test in Figure 2. (Note that the third element in each of Alice's question and answer triples is redundant and can be eliminated.)

It is easy to see that honest strategies pass the oracularized QUADEQ test with probability 1. To establish soundness of the test, i.e to show an analogue of Lemma 3.5 of [26], we can follow essentially the same steps as in the proof of that lemma. The key step of the proof is to argue that, due to the soundness of the linearity test against entangled provers, there exist measurements on each prover's space whose outcomes are linear functions that are consistent with the measurements applied in the test. For the oracularized test, we can perform this step using the soundness of the oracularized linearity test against entangled provers, which was analyzed in [20]. The rest of the proof proceeds unchanged. As a result we obtain the

Out of the two provers, choose one at random to be Alice and the other to be Bob.

1. With probability 1/4 each, do the following:
  - **a.** Send label ℓ_1 to the two players and perform the $(n/2)$-bit (oracularized) linearity test.
  - **b.** Same with label ℓ_2.
  - **c.** Send labels (ℓ_1, ℓ_2) to the two players and perform the n-bit linearity test.
  - **d.** Same but perform the n^2-bit linearity test.

2. Select random $u, v \in \mathbb{F}_2^{n/2}$ and $i \in [3]$, and generate the three queries $q_1 = (\ell_1, u)$, $q_2 = (\ell_2, v)$, $q_3 = (\ell_1, \ell_2, (u, v))$. Send q_1, q_2 to Alice, receiving answers a_1, a_2, and let $a_3 = a_1 + a_2$. Send q_i to Bob, receiving answer b. Accept if $b = a_i$.

3. Select random $u, v \in \mathbb{F}_2^n$ and $i \in [3]$, and generate the three queries $q_1 = (\ell_1, \ell_2, u)$, $q_2 = (\ell_1, \ell_2, v)$, $q_3 = (\ell_1, \ell_2, u \otimes v)$. Send q_1, q_2 to Alice, receiving answers a_1, a_2 and let $a_3 = a_1 \cdot a_2$. Send q_i to Bob, receiving answer b. Accept if $b = a_i$.

4. Select a random vector $v \in \mathbb{F}_2^K$ and let $w = \sum_k w_k a^{(k)} \in \mathbb{F}_2^{n^2}$. Send (ℓ_1, ℓ_2, w) to a randomly chosen player and check that the answer $a = \sum_k w_k c^{(k)}$.

Figure 2 The two-prover QUADEQ test. See Section [26, Section 3.4] for additional explanations regarding the notation.

following corollary, which establishes Theorem 1; it is completely analogous to [26, Corollary 4.3], except that due to the oracularization, the two provers now have to provide answers of two bits each instead of one.

▶ **Corollary 16.** *There is an $\varepsilon > 0$ such that the following holds. Given a two-player projection game G in explicit form in which answers from one player is restricted to 2 bits, and answers from the other player to a single bit, it is NP-hard to distinguish between $\omega(G) = 1$ and $\omega^*(G) \leq 1 - \varepsilon$.*

Using that the games G for which NP-hardness is shown in Corollary 16 are projection games, we may apply results on the parallel repetition of two-player entangled projection games [7] to amplify the completeness and soundness parameters from 1 and $1 - \varepsilon$ to 1 and δ respectively, for any $\delta > 0$, by repeating the game $\mathrm{poly}(\varepsilon^{-1} \log \delta^{-1})$ times and incurring a corresponding multiplicative factor blow-up in the length of questions and answers in the game.

References

1 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.

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

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

4 John S. Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1:195–200, 1964.

5 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.

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

7 Irit Dinur, David Steurer, and Thomas Vidick. A parallel repetition theorem for entangled projection games. *Computational Complexity*, 24(2):201–254, 2015. doi:10.1007/s00037-015-0098-3.

8 Albert Einstein, Boris Podolsky, and Nathan Rosen. Can quantum-mechanical description of physical reality be considered complete? *Physical review*, 47(10):777, 1935.

9 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.

10 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.

11 Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: Interactive proofs for muggles. *Journal of the ACM (JACM)*, 62(4):27, 2015.

12 Johan Håstad. Some optimal inapproximability results. *Journal of the ACM (JACM)*, 48(4):798–859, 2001.

13 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.

14 Tsuyoshi Ito and Thomas Vidick. A multi-prover interactive proof for NEXP sound against entangled provers. *Proc. 53rd FOCS*, pages 243–252, 2012. `arXiv:1207.0550`.

15 Rahul Jain, Zhengfeng Ji, Sarvagya Upadhyay, and John Watrous. QIP = PSPACE. *Communications of the ACM*, 53(12):102–109, 2010. `arXiv:0907.4737`.

16 Zhengfeng Ji. Classical verification of quantum proofs. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, pages 885–898. ACM, 2016.

17 Zhengfeng Ji. Compression of quantum multi-prover interactive proofs. 2016. `arXiv:1610.03133`.

18 Yael Tauman Kalai, Ran Raz, and Ron D. Rothblum. How to delegate computations: the power of no-signaling proofs. In David B. Shmoys, editor, *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 485–494. ACM, 2014. `doi:10.1145/2591796.2591809`.

19 Carsten Lund, Lance Fortnow, Howard Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *Journal of the ACM (JACM)*, 39(4):859–868, 1992.

20 Anand Natarajan and Thomas Vidick. A quantum linearity test for robustly verifying entanglement. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 1003–1015. ACM, 2017. `doi:10.1145/3055399.3055468`.

21 Anand Natarajan and Thomas Vidick. Low-degree testing for quantum states, and a quantum entangled games PCP. 2018. `arXiv:1801.03821v2`.

22 Ran Raz and Shmuel Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In Frank Thomson Leighton and Peter W. Shor, editors, *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing, El Paso, Texas, USA, May 4-6, 1997*, pages 475–484. ACM, 1997. `doi:10.1145/258533.258641`.

23 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.

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

25 Ben Toner. Monogamy of non-local quantum correlations. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 465(2101):59–69, 2009. `doi:10.1098/rspa.2008.0149`.

26 Thomas Vidick. Three-player entangled XOR games are NP-hard to approximate. In *Proc. 54th FOCS*, 2013. `arXiv:1302.1242`.

A Modified proofs from [26]

As noted in the introduction, the principal modifications to the soundness analysis of the low-degree test in [26] necessary to make it hold for two provers concern the self-improvement results of section 5. There are a few other steps of the proof of the main theorem in [26] that seem to require a tripartite tensor product factorization of the Hilbert space to be carried out. In all cases this is easily avoided by simple modification of the proof. Although they remain very elementary, in this appendix we describe the only two other non-trivial modifications needed. The first is in the proof of [26, Claim 6.10]. (We refer to the paper [26] for context, including an explanation of the notation; the following discussion is meant for a reader already familiar with the proofs in [26].)

▶ **Claim 17** (Claim 6.10 in [26]). *The measurements* $\{Q_s^g\}_{g\in\mathcal{P}_d(s)}$ *satisfy*

$$\mathop{\mathrm{E}}_{s\in\mathcal{S}_{m-1}(\mathbb{F}_q^m)} \sum_{g\in\mathcal{P}_d(s)} \langle Q_s^g, (\mathrm{Id}-Q_s^g)\rangle_\Psi = O(\varepsilon^{c_\ell}) \ .$$

Proof. The proof is the same as in [26], except the third tensor factor is not needed — the second can be used for the same purpose:

$$\begin{aligned}
\mathop{\mathrm{E}}_{s\in\mathcal{S}_{m-1}(\mathbb{F}_q^m)} \sum_{g,g'\in\mathcal{P}_d(s),g\neq g'} \langle Q_s^g, Q_s^{g'}\rangle_\Psi &\approx \mathop{\mathrm{E}}_{s\in\mathcal{S}_{m-1}(\mathbb{F}_q^m)} \mathop{\mathrm{E}}_{x\in S} \sum_{g,g'\in\mathcal{P}_d(s),g\neq g'} \langle Q_s^g, Q_s^{g'} A_x^{g(x)}\rangle_\Psi \\
&\qquad + O(\varepsilon^{c_\ell}) \\
&\approx_{\varepsilon^{c_\ell}} \mathop{\mathrm{E}}_{s\in\mathcal{S}_{m-1}(\mathbb{F}_q^m)} \mathop{\mathrm{E}}_{x\in S} \sum_{g,g'\in\mathcal{P}_d(s),g\neq g'} \langle Q_s^g A_x^{g(x)}, Q_s^{g'}\rangle_\Psi \\
&\qquad + O(\varepsilon^{c_\ell}) + O(\varepsilon) \\
&\approx O(\varepsilon^{c_\ell}) + O(\varepsilon) \ .
\end{aligned}$$

In the first line, we used the consistency between Q_s^g on the first prover and $A_x^{g(x)}$ on the second; in the second line, we used the self-consistency of A; and in the third, we used the consistency between $Q_s^{g'}$ on the second prover and $A_x^{g(x)}$ on the first prover. ◀

The second is in the proof of [26, Claim 6.14]. Here again, the use of a third tensor factor can be avoided by a simple modification. Specifically, the last set of centered equations on p.1056 (right below (6.22)) should be replaced with

$$\begin{aligned}
\mathop{\mathrm{E}}_{(s_i)} \sum_{g,\deg(g)>d} \langle R_{(s_i)}^g, \mathrm{Id}\rangle_\Psi &\approx_{\varepsilon^{c_\ell}} \mathop{\mathrm{E}}_{(s_i),z,\ell,\ell'\ni z} \sum_{g,\deg(g)>d} \sum_{\substack{h(\ell\cap s_i)=g(\ell\cap s_i)\\ h'(\ell'\cap s_i)=g(\ell'\cap s_i)}} \langle R_{(s_i)}^g, B_\ell^h B_{\ell'}^{h'}\rangle_\Psi \\
&\approx_{\varepsilon^{c_\ell}} \mathop{\mathrm{E}}_{(s_i),z,\ell,\ell'\ni z} \sum_{g,\deg(g)>d} \sum_{\substack{h(\ell\cap s_i)=g(\ell\cap s_i)\\ h'(\ell'\cap s_i)=g(\ell'\cap s_i)}} \langle R_{(s_i)}^g B_{\ell'}^{h'}, B_\ell^h\rangle_\Psi \\
&= O(\varepsilon^{d_c/2})
\end{aligned}$$

Retraction Notice

The article, published on June 4th, 2018 in the CCC 2018 proceedings (LIPIcs, volume 102, `https://www.dagstuhl.de/dagpub/978-3-95977-069-9`), has been retracted by agreement between the authors, the editor(s), and the publisher Schloss Dagstuhl / LIPIcs. The retraction has been agreed due to an error in the proof of the main result, arising from an error in the cited paper "Three-player entangled XOR games are NP-hard to approximate" by Thomas Vidick (SICOMP '16). The error in that paper is in the proof of soundness of the plane vs point low-degree test against 3 entangled provers. The main technical result of the present article is a modification of that proof to hold against 2 entangled provers, and is affected by the same error. For more details on the nature of the error, and a description of which results in the literature are invalidated and which have been recovered by other techniques, see the erratum by Thomas Vidick for the SICOMP '16 paper available at `http://users.cms.caltech.edu/~vidick/errata.pdf`. (At a high level, the results that survive are "scaled up" complexity results about the complexity of MIP*, whereas results in the "scaled down" setting such as the NP-hardness result of this article cannot be directly recovered with current techniques.)

Dagstuhl Publishing – January 5, 2021.