Erratum: Three-Player Entangled XOR Games are NP-hard to Approximate

Thomas Vidick

Abstract

This note indicates an error in the proof of Theorem 3.1 in [T. Vidick, SIAM J. Comp. 45(3):1007–1063, 2016]. Due to an induction step in the soundness analysis not being carried out correctly, the analysis fails to prove the claimed result. The error invalidates the proofs of the main computational hardness results claimed in the paper.

We discuss implications for subsequent works. In some cases results can be partially recovered by applying a weakened version of Theorem 3.1 shown in [Ji et al., arXiv:2009.12982] subsequently to the discovery of the error. The validity of Theorem 3.1 as stated in the paper remains an open question.

The main result of the paper [Vid16] is a proof of soundness of a "plane-vs-point" low-degree test introduced in [RS97] when the test is executed with three quantum players sharing entanglement. Soundness of the test is used to establish the main results of the paper, which are lower bounds on the complexity of approximating the entangled value of multiplayer games with quantum players sharing entanglement.

John Wright has discovered an error in the soundness analysis given in [Vid16] which invalidates the main results of the paper. In this note we describe the error (Section 1), explain how it affects the results claimed in [Vid16] and how they can be partially recovered using [JNV$^+$20b] (Section 2), and discuss consequences for subsequent works (Section 3).

1 Description of the error

The main technical result in [Vid16] is formulated as Theorem 3.1, which states soundness of the low-degree plane-vs-point test against entangled players. For a description of the test, see [Vid16, Section 3.1.1]. For convenience we recall the statement of the theorem, using the same notation as [Vid16].

(Retracted) Theorem 3.1. *Let $0 < \varepsilon \leq 1/2$, $d \geq 1, m \geq 2, r \geq 3$ integers, and $\mathbb{F}_q$ a finite field of size $|\mathbb{F}_q| = q$ such that $q \geq (dm/\varepsilon)^{d_1}$, where $d_1 \geq 1$ is a universal constant. Let $(|\Psi\rangle, A, C)$ be a strategy with success $1-\varepsilon$ in the $(d,m,r,\mathbb{F}_q)$-low-degree test. Then there exists a measurement $\{M^g\}$ with outcomes $g \in pol_d(\mathbb{F}_q^m)$ such that*

$$\mathrm{E}_{x\in\mathbb{F}_q^m} \sum_{g\in pol_d(\mathbb{F}_q^m)} \sum_{a\in\mathbb{F}_q:\, g(x)\neq a} \langle A_x^a,\, M^g\rangle_\Psi \;\leq\; C_1\varepsilon^{c_1}, \tag{1}$$

where $0 < c_1 \leq 1, C_1 > 0$ are universal constants.

In the theorem statement the notation $\mathrm{pol}_d(\mathbb{F}_q^m)$ refers to the set of multivariate polynomials on $\mathbb{F}_q$ with m variables and of total degree at most d. The other notation is not important for our purposes, and we refer to [Vid16] for details.

The proof of Theorem 3.1 given in [Vid16] contains a mistake. The remainder of this section explains the mistake. We note that no counter-example is known to the statement of Theorem 3.1, so that to the best of our present knowledge either proving or disproving the (retracted) theorem is an open problem.

The proof of Theorem 3.1 is by induction on the number of variables m. The induction hypothesis is essentially identical to the theorem statement for a given m and all d, ε and q satisfying the conditions stated in the theorem. (For our purposes the parameter r can be ignored.)

Fix d, q and ε and suppose that one wants to show that the induction hypothesis is true for some m, assuming that it is true for $(m-1)$ and all d', q' and ε' that satisfy the conditions of the theorem. The first step in doing so consists in applying the hypothesis for $(m-1)$ (and the same d, q and a closely related ε) a total of $(d+1)$ times to various "restrictions" of the strategy $(|\Psi\rangle, A, C)$ to $(m-1)$-dimensional hyperplanes in $\mathbb{F}_q^m$. Each such inductive call returns a measurement $\{M_i^g\}$ of the form given in the theorem statement, where here i ranges in $\{1, \ldots, d+1\}$. The second steps consists in combining together these $(d+1)$ measurements into a single measurement $\{M_{pasted}^g\}$ (this step is called "pasting" in the paper) that is shown to satisfy the requirements for the induction hypothesis for this m and d, q and ε.

As discussed in detail in the proof given in [Vid16] the second step induces an increase in "error", where by error we mean the best right-hand side that can be shown for (1). Whenever the inductive call returns measurements $\{M_i^g\}$ that satisfy the claimed bound, the combined measurement $\{M_{pasted}^g\}$ will satisfy (1) with a weaker error bound. For concreteness, let's say that the error obtained is twice what is desired. The key step that bootstraps the worse bound into the desired one is a so-called "consolidation procedure" that restores any "good enough" error to some given target value independent of m and d; it is this step that allows the theorem to claim constants C_1 and c_1 that do not depend on the other parameters of the problem.[1]

The consolidation procedure is described in [Vid16, Section 5]. The key step in its analysis is Proposition 5.8. Proposition 5.8 as stated in the paper is wrong. To explain the error, for convenience we re-state the proposition.

Proposition 1.1 (Proposition 5.8 in [Vid16]). *There exists a constant $K > 0$ such that the following holds. Let $r \geq 3$, H a symmetric r-player game, X a finite set, and for every $x \in X$, $G_x = (V_x, E_x)$ a graph, S_x a set, and $\mathcal{G}_x \subseteq \{g : V_x \to S_x\}$.*

Suppose that for any $0 < \varepsilon < K$ and strategy $(P, |\Psi\rangle)$ for the players that has success $1-\varepsilon$ in the game H and is ε-self-consistent there exists a collection $A_x = \{A_{x,v}^a\}_{a \in S}$ of projective measurements defined for every $v \in V_x$, possibly depending on P but independent of $|\Psi\rangle$, such that that for all $x \in X$, $(G_x, A_x, \mathcal{G}_x)_\Psi$ is a (δ, μ)-robust triple for some $\delta, \mu > 0$ such that $\eta'(\delta, \mu) < 1/4$, where η' is as defined in Lemma 5.4.

Suppose further that for any $\varepsilon' > 0$ there exists $\eta = \eta(\varepsilon')$ such that $\eta \to 0$ as $\varepsilon' \to 0$, and whenever $(P', |\Psi'\rangle)$ is a strategy with success $1 - \varepsilon'$ in H, there exists a family of sub-measurements $\{Q_x^g\}_{g \in \mathcal{G}}$ that is η-consistent with A'_x (obtained from P') and η-complete, on average over $x \in X$.

Then for any small enough $0 < \varepsilon < K$ and strategy $(P, |\Psi\rangle)$ for the players that has success $1-\varepsilon$ in the game H and is ε-self-consistent there exists a family of (complete) measurements $\{R_x^g\}_{g \in \mathcal{G}}$ that is η_c-consistent with A_x for some $\eta_c = O(r(\eta')^{1/4})$, on average over $x \in X$.

For our purposes the contents of the first "Suppose that" in the proposition statement can be ignored, except for the fact that the assumption is required to hold for all $0 < \varepsilon < K$. The "Suppose further" part formalizes the result of the induction and pasting steps as discussed above, which is that one is able to show the desired equation (1) with some dependence $\eta(\varepsilon')$ that goes to 0 with ε' but may not be as good as the

[1] For intuition, a combinatorial analogue of this step is the isoperimetric inequality, which states that a sufficiently small subset of an expander graph that does not expand much must necessarily be extremely small, much smaller than it was necessary to assume in the first place.

one claimed in (1) (say, it is twice too big). (For our purposes, one may equate the condition of being "η-consistent with A'_x" to satisfying (1) with a right-hand side of η.)

In order to reach the conclusion claimed in the proposition, the proof makes use of the "Suppose further" assumption twice. In the first application the strategy (P', Ψ') is (essentially) identical to the original strategy,[2] and ε' is closely related to ε. This application does not create any issues. In the second application the strategy $(P', |\Psi'\rangle)$ is derived from $(P, |\Psi\rangle)$ after large modifications, including some form of "conditioning" step that may blow up the error by some factor, i.e. in that application it may be the case that ε' is as large as 2ε.

In the text these two applications appear in the "Proof of Proposition 5.8" on p.1043. In the second paragraph of the proof it is assumed that the error ε of the original strategy is not too large so that a certain "self-improvement" lemma can be applied. Essentially, this requires that the quantity $\eta(\varepsilon')$ from the "Suppose further" statement is smaller than some universal constant such as $\frac{1}{10}$. The text claims that, by taking ε small enough, ε' can be made sufficiently small for $\eta(\varepsilon')$ to be small enough.

Unfortunately, in general this requires the "constant" K in the last paragraph to depend on the function η: the conclusion of the proposition only holds as long as $\eta(\varepsilon') \leq \frac{1}{10}$, where ε' can be thought of as being equal to 2ε. Unless $\eta(2x) < x$ for small x, a condition that is not shown in the paper, this requires the constant K to be chosen smaller in the conclusion part of the proposition than in the assumptions. As one unwinds the way the proposition is used in the inductive argument explained above, one realizes that this dependence eventually leads to a proof of the induction hypothesis that only applies for values of ε that are exponentially small in m and d.

To summarize, Proposition 5.8 in [Vid16] is wrong. A correct version of the proposition can be obtained where the constant K that appears in the last paragraph is made to depend on the function η from the second paragraph. However, this modification eventually leads to a statement of Theorem 3.1 where the bound on the right-hand side of (1) has an exponential dependence on m and d. Such a bound is not sufficient to yield the applications claimed in the paper.

2 How other results in [Vid16] are affected

As explained in the preceding section, the main technical result, Theorem 3.1, does not hold as claimed. Since all other results in the paper crucially rely on Theorem 3.1 they are invalidated as well. More precisely, the straightforward weakening of Theorem 3.1 that still holds requires a right-hand side in (1) that includes an exponential dependence on both m and d in the constant C_1. This dependence means that in order to obtain a meaningful hardness result one would have to employ a technique such as parallel repetition where the number of repetitions is exponential in m and d. In e.g. the protocol for verifying instances of 3SAT given in [Vid16, Figure 3] the parameters m and d are both poly-logarithmic in the instance size n. The number of repetitions required would then be (super-)polynomial in the instance size, yielding a game with questions and answers of (super-)polynomial length. That such a game can be used to verify 3-SAT is trivial, since with such communication only one player is needed, and the player can send a satisfying assignment as proof.

To the best of our knowledge the derivation of the hardness results presented in [Vid16, Section 4] from Theorem 3.1 is by itself not flawed. Therefore, were a replacement for Theorem 3.1 to be shown, the hardness results would follow. In work completed subsequently to the discovery of the error, the author together with Ji, Natarajan, Wright and Yuen established a weaker variant of the main result, available

[2] By "original strategy" we mean the one to which the proposition is applied: the $(P, |\psi\rangle)$ which appears in the last paragraph.

as [JNV$^+$20b]. We refer to [JNV$^+$20b] for a precise statement of their result and a comparison with [Vid16, Theorem 3.1] and give an informal comparison here.

The main result of [JNV$^+$20b] is a variant of Theorem 3.1 with two main differences. First, the low-degree test is replaced by a *low individual-degree test*. In this variant of the low-degree test the verifier only queries the players for the restriction of a polynomial to either an axis-parallel line or a point (in contrast the low-degree test queries them for the restriction to an arbitrary plane or a point). In the classical setting, soundness of this test is a straightforward extension of the multilinearity test from [BFL91].

The second, more important, difference is that the use of the low-individual degree test yields a weaker soundness guarantee. In particular, the $\mathrm{poly}(\varepsilon)$ bound claimed on the right-hand side of (1) is weakened to a bound that scales as $\mathrm{poly}(m) \cdot \mathrm{poly}(\varepsilon)$. As explained in [JNV$^+$20b], a polynomial dependence on m is unavoidable for the low individual-degree test.

Replacing the use of the low-degree test in [Vid16] by the low individual-degree test from [JNV$^+$20b] leads to a suitably weakened version of the hardness results claimed in the paper. In particular, results in [Vid16, Section 3] that compose the low-degree test first with itself, then with other tests to further reduce the answer size and eventually lead to an XOR game, can be applied starting from the low individual-degree test.

Recall that the NP-hardness results for the entangled value of certain restricted classes of 3-player entangled games derived in [Vid16, Section 4] are obtained by setting $m = O(\log n)$, with n the input size. Thus replacing the use of Theorem 3.1 by the main result in [JNV$^+$20b] and employing parallel repetition would enable one to recover the hardness results in a weakened form, where the games for which NP-hardness is claimed have questions of length $\mathrm{poly}\log(n)$ instead of $O(\log(n))$. Equivalently, the hardness for games of size n is weakened from $\mathrm{NTIME}(\mathrm{poly}(n))$ to $\mathrm{NTIME}(2^{O(\log^c n)})$ for some $0 < c < 1$. While we believe this application to be a straightforward consequence of [JNV$^+$20b] and the analysis given in [Vid16] we do not show nor claim the results in this erratum.

3 Consequences for subsequent work

Since the publication of [Vid16] a number of works have built on the results announced therein. We discuss consequences of the error on the three principal such works that we are aware of.

Consequences on [NV18b]. The main result of [NV18b] is a two-prover analogue of [Vid16, Theorem 3.1]. The proof re-uses most steps of the proof of Theorem 3.1, and in particular the consolidation procedure, as a black box. As a result, the main result of [NV18b] is invalidated and the paper has been withdrawn. A weaker version of that result, keeping the two provers but weakening the soundness guarantees, follows from [JNV$^+$20b] as discussed in the previous section.

Consequences on [NV18a]. The main result of [NV18a] is a quantum analogue of [Vid16, Theorem 3.1], and uses [Vid16, Theorem 3.1] as a black-box in its proof. Replacing the use of [Vid16, Theorem 3.1] by [JNV$^+$20b] leads to a weakening of the main result of [NV18a]. A precise formulation of this weakening is stated as [JNV$^+$20a, Theorem 7.16] and a full proof is given in [JNV$^+$20a, Appendix A]. Analogously to the weakened hardness results described in the previous section, the weakened version of the quantum low-degree test from [NV18a] (i.e. [JNV$^+$20a, Theorem 7.16]) leads to a weakened “quantum games PCP” where the size of the games is quasi-polynomial in the input size, as opposed to polynomial as claimed in [NV18a].

Consequences on [JNV$^+$20a]. An initial version of [JNV$^+$20a] posted on the arXiv (version 1) made use of [Vid16, Theorem 3.1] in an essential way. Subsequently to the discovery of the error in that result reported here, the authors established the weaker version given in [JNV$^+$20b] and showed that it is sufficient to establish the main results from [JNV$^+$20a], and in particular the equality $\mathrm{MIP}^* = \mathrm{RE}$, in their original form. (The main result of [NW19], the inclusion $\mathrm{NEEXP} \subseteq \mathrm{MIP}^*$, also holds in its original form.)

The characterization $\mathrm{MIP}^* = \mathrm{RE}$ supersedes the main hardness results from [Vid16] (and also [NV18a]) in many but not all ways. In particular, we believe (but do not prove here) that combining $\mathrm{MIP}^* = \mathrm{RE}$ with the standard answer-reduction techniques described in [Vid16] it is possible to show that the class of entangled three-prover interactive proof systems based on XOR games equals RE. From a purely complexity-theoretic point of view this result is stronger than the NP-hardness claimed in [Vid16]. However, there are ways in which such a result would be weaker than e.g. the version of the hardness from [Vid16, NV18a] described in Section 1 of this erratum and in the previous paragraph respectively. The main way in which the result would be weaker is that, at least in a black-box way, it would not suffice to establish that for languages in NP (resp. QMA) it suffices for the provers to have the power of BPP (resp. BQP) together with suitable access to a witness in order to succeed in the protocol, for the case of a YES-instance.

References

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

[JNV$^+$20a] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. MIP*= RE. *arXiv preprint arXiv:2001.04383v2*, 2020.

[JNV$^+$20b] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. Quantum soundness of the classical low individual degree test, 2020. Manuscript.

[NV18a] Anand Natarajan and Thomas Vidick. Low-degree testing for quantum states, and a quantum entangled games PCP for QMA. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 731–742. IEEE, 2018.

[NV18b] Anand Natarajan and Thomas Vidick. Two-player entangled games are NP-hard. In *33rd Computational Complexity Conference (CCC 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.

[NW19] Anand Natarajan and John Wright. NEEXP is contained in MIP*. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 510–518. IEEE, 2019.

[RS97] Ran Raz and Shmuel Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In *Proceedings of the Twenty-ninth Annual ACM Symposium on Theory of Computing*, STOC '97, pages 475–484, New York, NY, USA, 1997. ACM.

[Vid16] Thomas Vidick. Three-player entangled XOR games are NP-hard to approximate. *SIAM Journal on Computing*, 45(3):1007–1063, 2016.