Electronic Colloquium on Computational Complexity, Report No. 142 (2014)

Candidate Lasserre Integrality Gap For Unique Games

Subhash Khot * Dana Moshkovitz †

November 1, 2014

Abstract

We propose a candidate Lasserre integrality gap construction for the Unique Games problem. Our construction is based on a suggestion in [27] wherein the authors study the complexity of approximately solving a system of linear equations over reals and suggest it as an avenue towards a (positive) resolution of the Unique Games Conjecture. We use a new encoding scheme that we call the *real code*. The real code has two useful properties: like the long code, it has a unique local test, and like the Hadamard code, it has the so-called *sub-code covering* property.

This write-up represents a part of a work in progress.

1 Introduction

The Unique Games Conjecture [22] is currently one of the important questions in theoretical computer science. It is a perplexing question in the sense that researchers have no consensus regarding its correctness and a tantalizing question in the sense that its resolution might possibly be on the horizon. As shown in [24], the conjecture can be phrased equivalently in terms of solving a nearly-satisfiable system of *discrete* linear equations, where each equation depends on two variables:

Definition 1. 2Lin($\mathbb{F}$) Problem: *Given N variables $x_1, \ldots, x_N$ taking values over a finite field $\mathbb{F}$ and M equations $C_1, \ldots, C_M$ where each equation C_i is of the form $x_{i_1} - x_{i_2} = b_i$ and $b_i \in \mathbb{F}$. The goal is to find an assignment that maximizes the fraction of equations satisfied.*

Note that in each equation, for any value for either of the two variables, there is a *unique* value for the other variable that satisfies the equation. The Unique Games problem is a bit more general: each constraint is on two variables, the variables take values from an alphabet Σ, and for any value for either of the two variables in a constraint, there is a *unique* value for the other variable that satisfies the constraint (so a constraint corresponds to a permutation on Σ and different constraints may correspond to different permutations). As shown in [24], the essence of the Unique Games problem is captured even when the constraints are linear over a finite field and one may restrict to the 2Lin($\mathbb{F}$) problem. In particular, the Unique Games Conjecture can be stated as:

*khot@cims.nyu.edu. Computer Science Department, Courant Institute of Mathematical Sciences, New York University. Research supported by NSF grants CCF 1422159, 1061938, 0832795 and Simons Collaboration on Algorithms and Geometry grant.

†dmoshkov@csail.mit.edu. Department of Electrical Engineering and Computer Science, MIT. This material is based upon work supported by the National Science Foundation under Grant Number 1218547.

ISSN 1433-8092

Definition 2 (The Unique Games Conjecture)**.** *For any constant $\varepsilon > 0$, there is a finite field $\mathbb{F}$ of a constant size, such that given a* 2Lin($\mathbb{F}$) *instance that has an assignment that satisfies $(1-\varepsilon)$ fraction of all equations, it is NP-hard to find an assignment that satisfies even an ε fraction of all equations.*

If true, the Unique Games Conjecture implies optimal NP-hardness of approximation for a large number of optimization problems and in some cases, for entire classes of problems (see the surveys [23, 37] for more background on the Unique Games Conjecture). For example, Raghavendra [30] shows that, assuming the Unique Games Conjecture (and P $\neq$ NP), basic semidefinite programs (SDP) yield optimal approximation algorithms for constraint satisfaction problems and in particular, for the Unique Games problem itself. Indeed, researchers had already designed algorithms for the Unique Games problem based on semidefinite programs [16, 22, 12, 36], as well as constructed matching integrality gaps, showing that these algorithms do not disprove the Unique Games Conjecture [26]. After Raghavendra's work, the integrality gap constructions have been extended to SDPs that are more general, amounting to a combination of a basic SDP and a super-constant number of rounds of the so-called Sherali-Adams linear programming relaxation [28, 31].

In more recent years, researchers have been looking at semidefinite programs that are even more general. A promising approach is to consider the Lasserre hierarchy of semidefinite programs [6, 18]. In contrast to all the semidefinite programs considered before, current techniques seem inadequate to construct integrality gaps against the Lasserre hierarchy. Some of the limitations of current techniques towards constructing Lasserre integrality gaps have been formalized in [3]. The latter work shows that integrality gaps against weaker SDPs can be solved using only a *constant* number of rounds of the Lasserre hierarchy.

In the authors' opinion, a most pressing question at present is whether the Lasserre SDP yields a disproval of the Unique Games Conjecture, and if not, whether we can construct an integrality gap demonstrating that. This note is an approach towards constructing such integrality gaps. We consider the 2Lin($\mathbb{F}$) problem over boolean field:

Definition 3. Boolean 2Lin: *Given N variables $x_1, \ldots, x_N$ taking $\{-1, 1\}$-values and M equations $C_1, \ldots, C_M$ where each equation C_i is of the form $x_{i_1} \cdot x_{i_2} = b_i$ and $b_i \in \{-1, 1\}$. The goal is to find an assignment that maximizes the fraction of equations satisfied.*

Assuming the Unique Games Conjecture, given an instance of Boolean 2Lin where $1-\varepsilon$ fraction of the equations can be satisfied, it is NP-hard to satisfy $1-\Omega(\sqrt{\varepsilon})$ fraction of the equations [24, 14]. Such an inapproximability result for Boolean 2Lin could perhaps be equivalent to the Unique Games Conjecture. Even though an equivalence is not known formally (in fact, a promising direction for proving it was ruled out in [33, 5]), researchers tend to agree that the boolean case captures the main difficulty of general Unique Games. If one were able to prove NP-hardness of $(1-\varepsilon, 1-\Omega(\sqrt{\varepsilon}))$ gap in the boolean case, the proof might likely extend to NP-hardness of $(1-\varepsilon, 1-K(\mathbb{F})\cdot\sqrt{\varepsilon})$ gap for the general finite field case, where $K(\mathbb{F})$ is a constant with $K(\mathbb{F}) \to \infty$ as $|\mathbb{F}| \to \infty$. The latter result would then be enough, via parallel repetition, to amplify the gap to $(1-o(1), o(1))$ and prove the Unique Games Conjecture [32]!

At present however, we do not even know a $(1-\varepsilon, 1-C\cdot\varepsilon)$ gap with $C \to \infty$, even for general Unique Games, and even as Lasserre integrality gap (as opposed to *more ambitious* NP-hardness result). In light of these considerations, we set ourselves the following, seemingly modest but still challenging in our opinion, goal: show a $(1-\varepsilon, 1-C\cdot\varepsilon)$ integrality gap for C rounds of the Lasserre SDP for the Boolean 2Lin($\mathbb{F}$) problem where $C \to \infty$ (and of course $\varepsilon \to 0$ necessarily).

Towards this goal, we indeed construct an instance of the Boolean 2Lin problem that has a Lasserre SDP solution with value $1 - \varepsilon$ after a super-constant number of rounds. We believe that the instance has no integral solution with value $1 - C \cdot \varepsilon$ where $C \to \infty$ as $\varepsilon \to 0$. Though we are unable so far to prove this *soundness* property formally, we consider several examples and sketch an argument as to why it might be true. In the remainder of the introduction we discuss some of the ideas that go into the construction.

1.1 The Real Code

To prove NP-hardness of approximation results, one typically composes a known "PCP" such as Projection Games, e.g. [15, 2, 1, 33], with a gadget based on either the long code, e.g. [7, 19, 20], or the Hadamard code, e.g. [21]. The Projection Games problem is more general than the Unique Games problem. For projection games there is a partition of the set of variables into two, and each constraint depends on a variable from each part. For every constraint, a value to the variable from the first part determines at most one value to the variable from the second part, but not (necessarily) vice versa. Unlike Unique Games, Projection Games are known to be NP-hard to approximate, and serve as a typical PCP to compose with the long code or the Hadamard code.

The long code encodes an index $i \in [n]$, or equivalently a $\log n$ bit string, as the dictator function $f(x_1, \ldots, x_n) = x_i$ defined over $(x_1, \ldots, x_n) \in \{-1, 1\}^n$. Hadamard code encodes a string $(\sigma_1, \ldots, \sigma_n) \in \{-1, 1\}^n$ as the function $f(x_1, \ldots, x_n) = \prod_{j:\sigma_j=-1} x_j$ defined over $(x_1, \ldots, x_n) \in \{-1, 1\}^n$. Evidently, the long code has a much worse rate than the Hadamard code; that is, one encodes much less information using the long code compared to the Hadamard code.

A big advantage of the long code is that it has a "unique local test", i.e. a 2-query test whose predicate is unique, meaning an answer to either of the two queries determines the answer to the other query, if the tester is to accept. In this test, also known as the noise test, one picks $x = (x_1, \ldots, x_n) \in \{-1, 1\}^n$ at random, then slightly perturbs x by re-sampling each coordinate with probability δ to get a new vector $x' \in \{-1, 1\}^n$. The test checks that $f(x) = f(x')$. The test passes with probability $1 - \delta$ for dictators, and passes with much less probability, i.e. $1 - \Omega(\sqrt{\delta})$, for balanced functions that are far from dictators (for an appropriate definition of "far from dictator") [10, 24, 14]. Towards proving the Unique Games Conjecture (possibly in a weaker form), one may now attempt to compose Projection Games with the noise test, however, this fails. We won't elaborate this point further, but in short, the reason is that the long code is too long and does not have the sub-code covering property (described next) which, at least in the authors' opinion, seems essential to construct a 2-query PCP with unique local test, i.e. to prove the Unique Games Conjecture or its weaker form.

On the other hand, the Hadamard code does not have a unique local test (it is nevertheless very useful in other applications, thanks to tests with three or more queries [8]). The reason is simple: for any two distinct locations $x, x' \in \{-1, 1\}^n$, half of the legitimate Hadamard codes satisfy $f(x) = f(x')$ and the remaining half satisfy $f(x) \neq f(x')$, and thus any unique local test fails with probability $\frac{1}{2}$ on some legitimate Hadamard code. However, the Hadamard code has the following sub-code covering property: the Hadamard code of a string $(\sigma_1, \ldots, \sigma_n)$ is nearly uniformly covered by the Hadamard codes of its proper substrings.[1] This property is

[1]Here is a more precise statement. Pick a random subset $S \subseteq [n]$ of size $(1 - \delta)n$ and a random string $x' \in \{-1, 1\}^S$; define a string $x \in \{-1, 1\}^n$ by letting $x_j = x'_j$ if $j \in S$ and $x_j = 1$ otherwise; then the distribution of x is statistically close to uniform over $\{-1, 1\}^n$ (provided $\delta \ll \frac{1}{\sqrt{n}}$, see [25]). Note that x denotes a typical

potentially useful in a PCP composition as follows: while composing a Projection Game with an encoding scheme with the sub-code covering property, an encoding on the "larger side" of the game is nearly uniformly covered by the encodings of its "neighbors" on the "smaller side" of the game. Since the encodings on the smaller side are already "contained" in the encodings on the larger side, only the latter explicitly appear in the "PCP proof", the former being "present implicitly". Now, one may simply run a test (a unique local test if one intends to show hardness of the Unique Games problem) on the encoding on the larger side. In the soundness analysis, one is able to "list-decode" this encoding. Since the encoding is nearly uniformly covered by the encodings on the smaller side, essentially the same list-decoding also serves as the list-decoding on the smaller side, leading to "consistent decodings" on both the sides, completing the soundness analysis. In [25], this recipe is demonstrated using the Hadamard code. Therein the application is different (and not to the Unique Games problem, since the Hadamard code does not have a unique local test).

Having explained the limitation of using either the long code or the Hadamard code towards proving hardness of the Unique Games problem, we now sketch a key idea in this note. Our construction is based on a new encoding scheme over reals, which has a property analogous to the sub-code covering property of the Hadamard code, while also having a unique local test analogous to the long code (at least in a loose sense). We call the new code the *real code*. Unlike Hadamard or the long code, the real code no longer uses functions over the boolean hypercube $\{-1,1\}^n$, but over $\mathbb{R}^n$ with the underlying space equipped with the standard Gaussian measure. The range of the functions is $\{-1,1\}$ still. Specifically, we use half-spaces (or rather a "periodized" version of half-spaces). For $z \in \mathbb{R}$, define:

$$\mathsf{interval}(z) = \begin{cases} +1 & \text{if} \quad z \in [2k, 2k+1) \text{ for some } k \in \mathbb{Z} \\ \\ -1 & \text{if} \quad z \in [2k-1, 2k) \text{ for some } k \in \mathbb{Z}. \end{cases}$$

The *real encoding* of a string $\sigma = (\sigma_1, \ldots, \sigma_n) \in \{-1,1\}^n$ is now defined as the function $f_\sigma : \mathbb{R}^n \to \{-1,1\}$:

$$f_\sigma(x) = \mathsf{interval}\left(\sum_{i=1}^{n} \sigma_i \cdot x_i\right).$$

Just like the noise test for functions on the boolean hypercube, one can define a natural noise test for functions $f : \mathbb{R}^n \to \{-1,1\}$ on the Gaussian space. This is precisely the unique local test we use to test the real code. By a result of Borell [9], half-spaces maximize the acceptance probability of a Gaussian noise test and since the legitimate real encodings are indeed half-spaces (leaving aside the issue of "periodizing"), they maximize the acceptance probability of the unique local test. More details appear in Section 4. Moreover, the real code has a property analogous to the sub-code covering property (see footnote on the previous page for a comparison with the Hadamard code). Pick a random subset $S \subseteq [n]$ of size $(1-\delta)n$ and a random input $x' \in \mathbb{R}^S$ from the standard Gaussian measure; define an input $x \in \mathbb{R}^n$ by letting $x_j = x'_j$ if $j \in S$

location in the Hadamard code of a string $\sigma \in \{-1,1\}^n$. The bit of the code at this location is $\prod_{j:\sigma_j=-1} x_j$ and since the coordinates of x outside S are set to 1, this bit depends only on the coordinates in S, i.e. on x'. On the other hand, x' denotes a typical location in the Hadamard code of the substring $\sigma|_S$, i.e. σ restricted to S. In this sense, the Hamadard codes of substrings of σ of length $(1-\delta)n$ nearly uniformly cover the Hadamard code of σ. The specific manner in which x is chosen can be restated as follows. Pick a random subset $S \subseteq [n]$ of size $(1-\delta)n$; pick the coordinates in S uniformly at random from $\{-1,1\}$; pick the coordinates outside S to be "small" in value. In the boolean field $\{-1,1\}$, "small" amounts to the value 1, i.e. a value that has no effect on the bit of the code at location x.

and a uniformly random number in $[-\delta, \delta]$ otherwise; then x "looks like" an input chosen from $\mathbb{R}^n$ with the standard Gaussian measure. The reason is that a typical input chosen from $\mathbb{R}^n$ with the standard Gaussian measure does have a fraction δ of the coordinates with magnitude $O(\delta)$ and so "looks like" the input x. Akin to the Hadamard code, the coordinates of x outside S are small in value and do not much influence the bit of the real encoding at location x. More details appear in Section 6. Our Lasserre integrality gap candidate for the Boolean 2Lin problem uses the real code and appeals to both these properties.

Remark 1.1. *In recent years, researchers suggested the "short code" (aka the "low degree long code") as a more efficient alternative to the long code [4, 13]. The short code has a unique test, but does not have the sub-code covering property.*

1.2 Approximate Real Linear Equations Problem and Integrality Gaps

The starting point of our construction is a Lasserre integrality gap for the approximate real linear equations problem. In this problem, one is given a system of linear equations over reals and each equation is of the form:

$$\sum_{j=1}^{k} b_j y_j = 0,$$

where $b_j \in \{-1, 1\}$ and y_j are real variables. One wishes to satisfy the equations approximately, and not necessarily exactly. Also, one wishes to assign at least a constant fraction of the variables, values that are at least a constant in magnitude (and so, one cannot "cheat" by assigning the zero value to all the variables). Given a real valued assignment to the variables, the margin (or the error) on a typical equation as above is $\left|\sum_{j=1}^{k} b_j y_j\right|$ and the goal is to find an assignment that (approximately) minimizes the average margin over all the equations. Note that, assigning the variables y_j random $\{-1, 1\}$ values, one can always achieve a margin of $O(\sqrt{k})$ on average, so the question is whether one can do better (and the answer is negative as explained next).

In analogy to the discrete case, we refer to the approximate real linear equations problem as kLin($\mathbb{R}$). In the paper [27], the authors prove an optimal NP-hardness result for the problem 3Lin($\mathbb{R}$), i.e. even when each equation has only three variables. The authors show that when an assignment with average margin ε exists, it is NP-hard to find an assignment with average margin $O(\sqrt{\varepsilon})$, i.e. there is a quadratic gap. This is shown to be optimal in the sense that there is a matching SDP algorithm (effectively a *least square fit* algorithm). In [27], the coefficients b_j in the equations are allowed to be in a bounded interval, as opposed to being $\{-1, 1\}$, but this is a minor point.

In the current note, we start with a similar result for larger k, but where instead of NP-hardness we have an integrality gap against a linear number of rounds of the Lasserre hierarchy. The integrality gap instance amounts to saying that, on that instance, the Lasserre SDP pretends that there is an assignment with average margin $O(1)$ (in the sense that the SDP has a feasible solution with objective $O(1)$) whereas, actually, every assignment has average margin $\Omega(\sqrt{k})$. Dividing by a normalizing factor of k, the gap is same as $O(\frac{1}{k})$ versus $\Omega(\frac{1}{\sqrt{k}})$, i.e. a quadratic gap as in [27]. We derive this integrality gap by essentially re-interpreting Tulsiani's [38] integrality gap for constraint satisfaction problems[2]. Working with integrality gaps instead of NP-hardness results has several advantages that we employ in our construction: (1) Tulsiani's result is for

[2]The derivation was observed during discussions with Tulsiani and Worah.

random instances, and hence the integrality gap we start with is for random instances. (2) Tulsiani has a guarantee on *all linear-sized* sets of variables; one does not have this luxury in the NP-hardness setting as it would amount to an exponential time, hence a meaningless, reduction.

1.3 Organization

We discuss constraint satisfaction problems, their Lasserre semidefinite programs, and Tulsiani's result in Section 2. We obtain the integrality gap for approximate real linear equations problem in Section 3. We discuss the real code in Section 4, and how to incorporate constraint test in the real code in Section 5. We show how to check consistency between real codes in Section 6. The overall candidate integrality gap instance for the Boolean 2Lin problem is in Section 7, and in Section 8 we sketch why it might work through examples.

2 Constraint Satisfaction Problems and their Lasserre Semidefinite Programs

A predicate $P : \{-1,1\}^k \mapsto \{0,1\}$ leads to a constraint satisfaction problem (CSP) as follows. There are N variables taking values in $\{-1,1\}$ and M constraints, each defined on some (ordered) tuple of k variables. In a constraint, each variable first gets a $\{-1,1\}$ sign ("polarity"), and then the predicate P is applied on the tuple of polarized variables:

Definition 4. *For a predicate $P : \{-1,1\}^k \mapsto \{0,1\}$, an instance of CSP($P$) consists of N variables $x_1,\ldots,x_N$ and M constraints $C_1,\ldots,C_M$, where each constraint C is over a k-tuple of variables $\{x_{i_1},\ldots,x_{i_k}\}$ and is of the form $P(b_1x_{i_1},\ldots,b_kx_{i_k})$ where $b_1,\ldots,b_k \in \{-1,1\}$.*

We overload notation by using C to denote a typical constraint, as well as the tuple of variables appearing in it. For $j \in [k]$, we let $C[j] \in [N]$ denote the index of the j^{th} variable in C and $b_C \in \{-1,1\}^k$ be a vector such that its j^{th} coordinate $b_C[j]$ indicates the polarity of the j^{th} variable in C. For $u,v \in \{-1,1\}^k$, let $u \circ v \in \{-1,1\}^k$ denote the coordinate-wise product of u,v. A random instance of CSP(P) is one where the constraints are on randomly chosen k-tuples of variables, and the polarities of variables are randomly chosen as well (independently for occurrences in different constraints).

The optimization problem associated with CSP(P) is to find an assignment to the variables that satisfies the largest number of constraints. If Φ is an instance of CSP(P), then we denote the maximum number of satisfied constraints by $OPT(\Phi)$.

The t-round Lasserre semidefinite program for the CSP(P) problem has a vector variable $V_{S,\alpha}$ for every set of variables $S \subseteq [N]$, $|S| \leq t$, and an assignment $\alpha \in \{-1,1\}^{|S|}$ to the variables in S. In the intended solution, $V_{S,\alpha}$ is some (globally fixed) unit vector if S is assigned α, and is the zero vector otherwise. If $\alpha_1 \in \{-1,1\}^{t_1}$ is an assignment to a set $S_1 \subseteq [N]$ of variables, and $\alpha_2 \in \{-1,1\}^{t_2}$ is an assignment to a set $S_2 \subseteq [N]$ of variables, then we say that α_1 and α_2 agree if they assign the same values to variables in $S_1 \cap S_2$; otherwise, we say that they disagree. The Lasserre program attempts to maximize the number of satisfied constraints as reflected by the vector variables, subject to consistency constraints on the variables:

Lasserre semidefinite program

$$
\begin{aligned}
&\max \quad \sum_{i=1}^{M} \sum_{\alpha \in \{-1,1\}^k} P(\alpha \circ b_C) \|V_{C_i,\alpha}\|^2 \\
&\text{s.t.}
\end{aligned}
$$

$$
\begin{aligned}
\textbf{Orthogonality}: &\quad \langle V_{S,\alpha}, V_{S,\beta} \rangle = 0 && \forall S, \alpha \neq \beta && (1)\\
\textbf{Consistency}: &\quad V_{S,\alpha} = V_{S\cup\{x\},\alpha\cup\{+1\}} + V_{S\cup\{x\},\alpha\cup\{-1\}} && \forall S, \alpha, x \notin S && (2)\\
\textbf{Non-negativity}: &\quad \langle V_{S,\alpha}, V_{T,\beta} \rangle \geq 0 && \forall S, T, \alpha, \beta && (3)\\
\textbf{Normalization}: &\quad \|V_{\phi,\phi}\|^2 = 1. && && (4)
\end{aligned}
$$

Local distributions: In any feasible solution to the semidefinite program, for every S, we have

$$\sum_{\beta \in \{-1,1\}^{|S|}} \|V_{S,\beta}\|^2 = 1.$$

Thinking of $\|V_{S,\beta}\|^2$ as probabilities, a vector solution to the Lasserre semidefinite program induces a distribution over assignments to S, referred to as the local distribution on S. In particular, for each constraint C, one has a local distribution on assignments to that constraint. Any set $S \supseteq C$ induces a distribution over assignments to C by picking $\beta \in \{-1,1\}^{|S|}$ with probability $\|V_{S,\beta}\|^2$ and restricting β to the variables in C. The feasibility of the solution ensures that this induced distribution coincides with the local distribution on C.

Note that the objective function of the program, measuring the quality of the solution, depends only on the local distributions. If one only has local distributions and consistency among them, rather than vector solutions and consistency among them, one gets the so-called Sherali-Adams linear programming relaxation. In the first reading, the reader might want to focus only on the local distributions. We rely on a result of Tulsiani concerning CSPs with the following linear predicate:

Definition 5. The Hypergraph Linearity Test Predicate: *For $k = 2^s - 1$, the hypergraph linearity test predicate $P_{H\mathrm{LIN}} : \{-1,1\}^k \mapsto \{0,1\}$ is defined as follows. Index the k coordinates by non-empty subsets $A \subseteq [s]$ and assume w.l.o.g. that the first s coordinates correspond to the singleton sets. Then $x \in \{-1,1\}^k$ is a satisfying assignment of the predicate $P_{H\mathrm{LIN}}$ (i.e., $P_{H\mathrm{LIN}}(x) = 1$) if and only if*

$$x_A = \prod_{i \in A} x_{\{i\}} \quad \forall\ 2 \leq |A| \leq s.$$

In other words, the satisfying assignments of the predicate are precisely the Hadamard codewords and $2^s = k+1$ in number.

Samorodnitsky and Trevisan [34] constructed a UGC-based PCP using the Hypergraph Linearity Test Predicate, and Chan [11], in a recent remarkable result, constructed a similar PCP without relying on the UGC. Regarding the Lasserre integrality gap, Tulsiani, building on the works of Grigoriev and Schoenebeck [17, 35], shows the following (the statement is tailored to our needs):

Theorem 6 (Tulsiani [38])**.** *Let T be an arbitrarily large constant. Let Φ be a randomly chosen instance of $\mathsf{CSP}(P_{H\mathrm{LIN}})$ with N variables and $M = TN$ constraints for large enough (growing) N. Let $n_0 = \lfloor \eta N \rfloor$ where $\eta = \frac{1}{T^{25}}$. Then with high probability over the choice of Φ:*

1. Completeness: Φ *has an* n_0*-round Lasserre SDP solution with objective value* M*, where for every constraint* C*, the local distribution on* C *is uniform over the satisfying assignments to* C *(i.e. uniform over* $P_{H\text{LIN}}^{-1}(1) \circ b_C$*).*

2. Soundness: $OPT(\Phi) \leq (1 + o(1))\frac{k+1}{2^k} \cdot M$.

In the soundness case, note that since the hypergraph linearity test predicate has $k+1$ satisfying assignments, the expected number of constraints satisfied by a random $\{-1, 1\}$ assignment to the variables is $((k+1)/2^k) \cdot M$. A standard argument shows that, with high probability over the choice of the instance, no assignment satisfies a slightly larger fraction of constraints (here the $o(1)$ term becomes arbitrarily small as T increases). What is remarkable is that in the completeness case, there is a SDP solution, up to a linear number of rounds of Lasserre, which "pretends" that there is an assignment satisfying all the constraints.

3 Lasserre Integrality Gap for Approximate Real Linear Equations Problem

In this section, we construct a Lasserre integrality gap for approximate real linear equations problem. The construction is essentially a re-interpretation of Tulsiani's Lasserre integrality gap for the hypergraph linearity test predicate, where we re-interpret a predicate over boolean domain as an equation over reals with carefully chosen coefficients. As Tulsiani's instance is a random instance of CSP(P_{HLIN}), our integrality gap can be thought of as a random, or average-case, analog of our NP-hardness result in [27].

Recall that the predicate $P_{H\text{LIN}} : \{-1, 1\}^k \mapsto \{0, 1\}$ has exactly $k+1$ satisfying assignments and $k = 2^s - 1$. It is easily verified that for any two distinct satisfying assignments $a, b \in P_{H\text{LIN}}^{-1}(1)$, we have $\sum_{j=1}^{k} a_j b_j = -1$. Indeed, since the k co-ordinates are indexed by non-empty subsets $A \subseteq [s]$, for some distinct $x, y \in \{-1, 1\}^s$, we have

$$\sum_{j=1}^{k} a_j b_j = \sum_{A \subseteq [s],\ A \neq \phi} \prod_{i \in A} x_i \cdot \prod_{i \in A} y_i = -1.$$

Let Φ be an instance of CSP($P_{H\text{LIN}}$) with N variables and M constraints. A typical constraint is denoted as C along with the vector b_C of polarities. We construct a system of linear equations over reals by replacing each constraint C by a set of $k+1$ linear equations over reals, one equation, as below, for each *sign vector* $\epsilon \in P_{H\text{LIN}}^{-1}(1)$:

$$\sum_{j=1}^{k} \epsilon_j \cdot b_C[j] \cdot y_{C[j]} = 0,$$

where $y_1, \ldots, y_N$ are real-valued variables. In the following, a constraint C will refer to the constraint as in the CSP($P_{H\text{LIN}}$) instance, and also to any of the $k+1$ real linear equations constructed from it. It should be clear from the context which is being referred to. We observe that a uniformly random satisfying assignment to the constraint C (i.e. uniform in $P_{H\text{LIN}}^{-1}(1) \circ b_C$) is, on average, a good assignment to each of the linear equations constructed from it, in terms of the average ℓ_1 error (i.e. margin).

Fact 3.1. *Let $\epsilon \in P^{-1}_{H\textsc{Lin}}(1)$ be any fixed sign vector. Then*

$$\mathop{\mathbf{E}}_{\alpha \in P^{-1}_{H\textsc{Lin}}(1) \circ b_C} \left[\left| \sum_{j=1}^{k} \epsilon_j \cdot b_C[j] \cdot \alpha_j \right| \right] = \frac{2k}{k+1}.$$

Proof. Substituting $\alpha = \beta \circ b_C$ above and canceling out the polarities, the expectation is

$$\mathop{\mathbf{E}}_{\beta \in P^{-1}_{H\textsc{Lin}}(1)} \left[\left| \sum_{j=1}^{k} \epsilon_j \cdot \beta_j \right| \right].$$

Note that $P_{H\textsc{Lin}}$ has $k+1$ satisfying assignments. Out of these, there is one assignment that equals ϵ and the inner sum equals k. For the remaining k assignments $\beta \neq \epsilon$ and the inner sum equals -1 as observed before. □

Motivated by the observation that $\{-1,1\}$-valued assignments to the variables suffice for approximate satisfaction of the real linear equations we defined, we continue to refer to the Lasserre semidefinite program we described before, which has a variable $V_{S,\alpha}$ per set S of at most t variables and per $\{-1,1\}$ assignment α to the variables in S. We keep the feasibility conditions of this program, but drop the objective function (which talks about satisfying the predicate $P_{H\textsc{Lin}}$).

Theorem 6 now directly implies our Lasserre integrality gap for approximate real linear equations problem, stated as Theorem 7 below. In the completeness part, we have a feasible vector solution inducing local distributions that on average approximate each equation up to a margin $O(1)$. In the soundness part, we have that every boolean assignment to the variables has average margin $\Omega(\sqrt{k})$, averaged over all the equations.

Theorem 7. *Let T be an arbitrarily large constant. Let Φ be a randomly chosen instance of $k\textsc{Lin}(\mathbb{R})$ with N variables and $M = TN$ constraints for large enough (growing) N, as described above. Note that every equation is of the form $\sum_{j=1}^{k} \epsilon_j \cdot b_C[j] \cdot y_{C[j]} = 0$. Let $n_0 = \lfloor \eta N \rfloor$ where $\eta = \frac{1}{T^{25}}$. Then w.h.p. over the choice of Φ we have:*

1. **Completeness:** *Φ has an n_0-round Lasserre SDP feasible solution, where for every constraint C, the local distribution is uniform on its satisfying assignments, i.e. $P^{-1}_{H\textsc{Lin}}(1) \circ b_C$. In particular, for every sign vector $\epsilon \in P^{-1}_{H\textsc{Lin}}(1)$, when taking expectation over the local distribution on C,*

$$\mathop{\mathbf{E}}_{\sigma \in P^{-1}_{H\textsc{Lin}}(1) \circ b_C} \left[\left| \sum_{j=1}^{k} \epsilon_j \cdot b_C[j] \cdot \sigma(C[j]) \right| \right] = \frac{2k}{k+1} \ \leq \ 2. \tag{5}$$

2. **Soundness:** *For some absolute constant $c > 0$, the following holds: for any global assignment $\sigma : [N] \mapsto \{-1,1\}$, for a c fraction of the equations (where an equation is specified by constraint C and a sign vector ϵ), the margin is at least $c \cdot \sqrt{k}$, i.e.*

$$\left| \sum_{j=1}^{k} \epsilon_j \cdot b_C[j] \cdot \sigma(C[j]) \right| \ \geq \ c \cdot \sqrt{k}. \tag{6}$$

Proof. The completeness part follows from Theorem 6 and Fact 3.1. The soundness part is a standard probabilistic argument over the choice of the instance Φ: fix a global assignment σ, fix the tuples of variables that appear in all the constraints, and consider the choice of the polarities for all the constraints. Given a constraint C and a sign vector ϵ, over the choice of random polarities $b_C[j]$, the sum $\sum_{j=1}^{k} \epsilon_j \cdot b_C[j] \cdot \sigma(C[j])$ is at least $c \cdot \sqrt{k}$ in magnitude with probability c for some absolute constant $c > 0$. Since the polarities are chosen independently for different constraints and the number of constraints TN is large relative to the number of variables N, one can apply Chernoff bound and a union bound. □

Remark 3.1. *(1) In the soundness case above, it is easy to extend the conclusion to all assignments $\sigma : [N] \mapsto \mathbb{R}$ (as opposed to only boolean assignments $\sigma : [N] \mapsto \{-1, 1\}$) as long as σ assigns, to a constant fraction of the variables, values that are at least a constant in magnitude. (2) The gap $(O(1),\ \Omega(\sqrt{k}))$ in the completeness versus the soundness case in Theorem 7 is our starting gap. We re-emphasize some of the points mentioned before. Dividing by a normalization factor of k, the gap here is $(O(\frac{1}{k}), \Omega(\frac{1}{\sqrt{k}}))$, i.e., a quadratic gap. In [27], the authors indeed prove that it is NP-hard to distinguish such a quadratic gap even when equations involve three variables. Thus Theorem 7 is the integrality gap analogue of NP-hardness result in [27]. The authors therein propose that there may be a further reduction from the system of equations with three variables (3*LIN*($\mathbb{R}$)) to a system of equations with two variables (2*LIN*($\mathbb{R}$)) and/or the closely related Boolean 2*LIN *problem. This note may substantiate their proposal, albeit in the context of Lasserre integrality gaps. We are at present unable to show soundness of our construction. If the construction is sound and the techniques to analyze soundness are developed, it may be possible to extend the integrality gap construction to a NP-hardness reduction.*

4 The Real Code and Gaussian Noise Test

The encoding scheme in our construction is the real code as explained in the introduction. Recall that for some assignment $\sigma : [n] \mapsto \{-1, 1\}$, its real code encoding $f_\sigma : \mathbb{R}^n \mapsto \{-1, 1\}$ is supposed to be the interval function of σ. Specifically, defining for $z \in \mathbb{R}$,

$$\mathsf{interval}(z) = \begin{cases} +1 & \text{if} \quad z \in [2k, 2k+1) \text{ for some } k \in \mathbb{Z} \\ \\ -1 & \text{if} \quad z \in [2k-1, 2k) \text{ for some } k \in \mathbb{Z}, \end{cases} \tag{7}$$

we let, with underlying standard Gaussian measure on $\mathbb{R}^n$ (denoted $\mathcal{N}^n$),

$$f_\sigma(x) = \mathsf{interval}\left(\sum_{i=1}^{n} \sigma(i) \cdot x_i\right). \tag{8}$$

Let Φ be an instance of kLIN($\mathbb{R}$) as in Theorem 7 along with the Lasserre SDP solution. For every set $S \subseteq [N]$, $|S| = n$, our final construction has a copy of the "gadget" $f_S : \mathbb{R}^n \mapsto \{-1, 1\}$ which is supposed to be the real encoding of the assignment to S. We use a standard PCP trick called *folding* to enforce certain basic properties of f_S:

Folding: The encoding $f = f_\sigma$ in Equation (8) satisfies

$$f(x_1, \ldots, x_{i-1}, x_i, x_{i+1}, \ldots, x_n) = -f(x_1, \ldots, x_{i-1}, x_i + 1, x_{i+1}, \ldots, x_n),$$

for any $i \in [n]$ and $x \in \mathbb{R}^n$. Moreover f is odd, i.e., $f(-x) = -f(x)$ for any $x \in \mathbb{R}^n$. By using these identities for evaluating f, we can assume that the functions f in our Boolean 2Lin instance always satisfy these identities.

Now we describe a test that checks that a given function $f : \mathbb{R}^n \mapsto \{-1, 1\}$, at least loosely speaking, resembles an encoding of some assignment $\sigma : [n] \mapsto \{-1, 1\}$, or more generally of some *reasonable* assignment $\sigma : [n] \mapsto \mathbb{R}$. We call it the *low boundary test.* It fails with only a *small* probability whenever f is indeed a correct encoding of a boolean assignment (there will be more tests, namely a *constraints test* and a *consistency test* that will be added later). The test below is applied on a given function $f : \mathbb{R}^n \mapsto \{-1, 1\}$. We think of the parameter α as infinitesimally small.

Low Boundary Test with Parameter α

- Pick $x, w \in \mathcal{N}^n$ independently and let $y = (1-\alpha)x + \sqrt{2\alpha - \alpha^2}w$ (thus y is a α-perturbation of x).
- Reject if and only if $f(x) \neq f(y)$.

Lemma 4.1. *If $f : \mathbb{R}^n \mapsto \{-1, 1\}$ is an encoding of a boolean assignment $\sigma : [n] \mapsto \{-1, 1\}$ as in Equation* (8)*, then f rejects the low boundary test with probability $O(\sqrt{\alpha n})$.*

Proof. This is because f, viewed as a partition of $\mathbb{R}^n$, has a Gaussian boundary/surface-area $O(\sqrt{n})$ and then one uses Corollary 14 in [29]. Alternately, it is easily seen that the two sums $\sum_{i=1}^{n} \sigma(i) \cdot x_i$ and $\sum_{i=1}^{n} \sigma(i) \cdot y_i$ are typically spread over a band of width $\Theta(\sqrt{n})$ around the origin and typically differ by $O(\sqrt{\alpha n})$. Thus the probability that the two sums lie in adjacent odd/even intervals is $O(\sqrt{\alpha n})$. $\square$

5 The Real Code Augmented with Constraint Test

We now augment the basic gadget with an additional test that allows for checking kLin($\mathbb{R}$) constraints. Suppose that there is a constraint $C \subseteq [n]$ of the kLin($\mathbb{R}$) instance of the form:

$$\sum_{j=1}^{k} \epsilon_j \cdot b_C[j] \cdot y_{C[j]} = 0.$$

For any such constraint C, let $\mathbf{v}_C$ denote the unit vector in $\mathbb{R}^n$ that has $\frac{\epsilon_j \cdot b_C[j]}{\sqrt{k}}$ in the position $C[j]$ and zero elsewhere. Let β be a parameter thought of as infinitesimally small.

The Constraint Test for a Given Constraint C and Parameter β

- Pick $x, y \in \mathbb{R}^n$ such that both x, y are distributed as $\mathcal{N}^n$ and $y = x + \beta\ell\mathbf{v}_C$ and $\ell \sim \mathcal{N}$. Specifically, x, y are picked by first selecting their common component orthogonal to $\mathbf{v}_C$ and then selecting their components along $\mathbf{v}_C$ in a $(1 - \frac{\beta^2}{2})$-correlated manner.[3]

[3] If y^* and x^* denote the components along $\mathbf{v}_C$, then this amounts to saying $y^* = (1 - \frac{\beta^2}{2})x^* + \sqrt{\beta^2 - \frac{\beta^4}{4}}w^*$ for $x^*, w^* \sim \mathcal{N}$. Thus $y^* - x^* \sim \beta\mathcal{N}$.

- Reject if and only if

$$f(x) \neq f(y). \tag{9}$$

Note that x, y differ only on coordinates in C. Next we analyze the behavior of the test on local distributions induced by the Lasserre SDP solution:

Lemma 5.1. *Suppose an assignment $\sigma : [n] \mapsto \{-1, 1\}$ is sampled from the local distribution on set S as given by the Lasserre SDP solution. Let $C \subseteq [n]$ be a constraint with a corresponding linear equation of the form $\sum_{j=1}^{k} \epsilon_j \cdot b_C[j] \cdot y_{C[j]} = 0$. Then the average rejection probability of the Constraint Test (w.r.t. constraint C) over the choice of σ is at most $O(\frac{\beta}{\sqrt{k}})$.*

Proof. Note that in the Lasserre solution, the restriction of σ to the constraint C is uniformly distributed over the satisfying assignments to C. Hence, over the choice of σ,

$$\sum_{j=1}^{k} \epsilon_j \cdot b_C[j] \cdot \sigma(C[j]) = \begin{cases} -1 & \text{with probability} \quad \frac{k}{k+1} \\ \\ k & \text{with probability} \quad \frac{1}{k+1}. \end{cases} \tag{10}$$

In the first case, $\langle \sigma, \mathbf{v}_C \rangle = -\frac{1}{\sqrt{k}}$, hence

$$\langle \sigma, y - x \rangle \sim \frac{\beta}{\sqrt{k}} \, \mathcal{N},$$

and the Constraint Test (9) rejects with probability $O(\frac{\beta}{\sqrt{k}})$. This is because the sums $\sum_{i=1}^{n} \sigma(i) \cdot x_i$ and $\sum_{i=1}^{n} \sigma(i) \cdot y_i$ are spread over a band of width $\Theta(\sqrt{n})$ around the origin, and their difference is distributed as $\frac{\beta}{\sqrt{k}} \, \mathcal{N}$ as shown. Thus the probability that the two sums lie in adjacent odd/even intervals is $O(\frac{\beta}{\sqrt{k}})$.

Similarly, in the second case, $\langle \sigma, \mathbf{v}_C \rangle = \sqrt{k}$, hence

$$\langle \sigma, y - x \rangle \sim \beta \sqrt{k} \, \mathcal{N},$$

and the Constraint Test (9) rejects with probability $O(\beta\sqrt{k})$. Overall, the Constraint Test rejects with probability $O(\frac{\beta}{\sqrt{k}})$. □

The constraint test examines the behavior of the function along a small number k of coordinates among the n coordinates, whereas the low boundary test might be insensitive to changes in such a small number of coordinates. This motivates a generalization of the low boundary test which focuses on any given subset K of the coordinates (this generalization will also be a part of our final construction). For $x \in \mathbb{R}^n$, let x_K denote the restriction of x to the coordinates in $K \subseteq [n]$, and let $x_{\overline{K}}$ denote the restriction of x to coordinates in $\{1, \ldots, n\} \setminus K$. As before, we think of the parameter α as infinitesimally small.

General Low Boundary Test with Parameter α on Subset $K \subseteq \{1, \ldots, n\}$

- Pick $x, w \in \mathcal{N}^n$ independently. Let $y_{\overline{K}} = x_{\overline{K}}$ and $y_K = (1 - \alpha)x_K + \sqrt{2\alpha - \alpha^2} w_K$.
- Reject if and only if $f(x) \neq f(y)$.

When $K = \{1, \ldots, n\}$, this test is same as the low boundary test we defined before. Lemma 4.1, analyzing the boundary test on the real code, continues to hold for any $K \subseteq \{1, \ldots, n\}$ with the appropriate scaling:

Lemma 5.2. *If $f : \mathbb{R}^n \mapsto \{-1, 1\}$ is an encoding of a boolean assignment $\sigma : [n] \mapsto \{-1, 1\}$ as in Equation* (8)*, then f rejects the general low boundary test on subset K with probability $O(\sqrt{\alpha\, |K|})$.*

6 The Consistency Test

Our proposed integrality gap instance for the Boolean 2Lin problem (in Section 7) has a block of variables for every subset $S \subseteq [N], |S| = n$, for an appropriate setting of parameter n. The variables correspond to points in space $\mathbb{R}^n$ (discretized appropriately and weighed according to the standard Gaussian measure). The variables are boolean and an assignment to the variables in a block corresponds to a function $f_S : \mathbb{R}^n \mapsto \{-1, 1\}$. For some global assignment $\tau : [N] \mapsto \{-1, 1\}$, the function f_S is intended to be the encoding of $\sigma = \tau|_S$ as in Equation (8). On each block, we are going to perform the two tests described so far. In addition, we need a test to check consistency between different blocks (i.e., to check, in at least some loose sense, that the functions f_S for different blocks are encodings of block assignments $\sigma(S)$ that are consistent across blocks, giving rise to consistent global assignment). We describe the consistency test next.

Roughly speaking if there are two blocks S and R such that $|S \cap R| \approx (1 - \delta)n$, then the (intended) linear interval functions f_S and f_R are nearly the same and thus we may test that this is indeed the case. We describe the test formally below. An absolutely crucial aspect of the test is that the coordinates in $S \setminus R$ and $R \setminus S$ are *very small* compared to the coordinates in $S \cap R$.

The test has a parameter δ satisfying $\frac{1}{n^{1/2}} \ll \delta \ll \frac{1}{n^{1/3}}$. It is not clear what the correct setting should be. For now think of $\delta = \frac{1}{\sqrt{n}}$. Let $I = [-s, s]$ be an interval whose measure w.r.t. the standard Gaussian is δ (and thus $s \approx \sqrt{2\pi}\, \frac{\delta}{2}$). Let $\mathcal{D}_I$ and $\mathcal{D}_{\bar{I}}$ denote the distribution of $x \sim \mathcal{N}$ conditional on being $x \in I$ and $x \notin I$ respectively.

Consistency Test with Parameter δ

Given functions $\{f_S : \mathbb{R}^n \mapsto \{-1, 1\} \mid S \subseteq [N], |S| = n\}$.

- Pick a set $S \subseteq [N], |S| = n$ at random.
- Pick $U \subseteq S$ by including each element of S with probability $1 - \delta$ and let $|U| = m$. Pick $R \subseteq [N], |R| = n$ such that $S \cap R = U$.
- Pick $x \sim \mathcal{D}_{\bar{I}}^m$. Pick $y^S, y^R \sim \mathcal{D}_I^{n-m}$ independently. We think of the coordinates of x, y^S, y^R as indexed by elements of $U, S \setminus U$ and $R \setminus U$ respectively.
- Reject if and only if
$$f_S(x, y^S) \neq f_R(x, y^R).$$

Note that the distribution of both the queries (x, y^S) and (x, y^R) is precisely $\mathcal{N}^n$. We have the following lemma regarding the rejection probability of the test when the functions f_S and f_R are indeed encodings of consistent assignments.

Lemma 6.1. *Suppose the functions f_S and f_R are encodings of assignments $\sigma(S) : S \mapsto \{-1, 1\}$ and $\sigma(R) : R \mapsto \{-1, 1\}$ respectively such that $\sigma(S)|_{S\cap R} = \sigma(R)|_{S\cap R} =: \pi$. Then the failure probability of the Consistency Test above is at most $O(\delta\sqrt{\delta n})$ (which is $\ll 1$ by our choice of $\delta \ll \frac{1}{n^{1/3}}$).*

Proof. We have:

$$f_S(x, y^S) = \mathsf{interval}\left(\sum_{i\in U} \pi(i) \cdot x_i + \sum_{\ell\in S\setminus U} \sigma(S)(\ell) \cdot y^S_\ell\right).$$

$$f_R(x, y^R) = \mathsf{interval}\left(\sum_{i\in U} \pi(i) \cdot x_i + \sum_{\ell\in R\setminus U} \sigma(R)(\ell) \cdot y^R_\ell\right).$$

Note that the sums are spread over a band of width $\Theta(\sqrt{n})$ around the origin whereas the difference between the two sums is attributed to y^S and y^R and is typically $O(\delta\sqrt{n-m}) = O(\delta\sqrt{\delta n})$ in magnitude. Therefore the two interval functions differ with probability $O(\delta\sqrt{\delta n})$. □

7 The Overall Construction

We are now ready to describe our proposed construction of the Boolean 2Lin integrality gap instance. Let Φ be the instance of kLin($\mathbb{R}$) with N variables and $M = TN$ constraints as in Theorem 7. We think of k as a large constant, T as a constant large enough after choosing k and N as a growing parameter. As we noted, every constraint C of Φ is a homogeneous linear equation over reals:

$$\sum_{j=1}^{k} \epsilon_j \cdot b_C[j] \cdot y_{C[j]} = 0.$$

Let $n = n_0/t$ for n_0 as in Theorem 7 and t is the number of Lasserre rounds for the instance we construct. As mentioned before, in our Boolean 2Lin instance, there is a block of variables for every subset $S \subseteq [N], |S| = n$. The variables correspond to points in the space $\mathbb{R}^n$ and an assignment to this block corresponds to a function $f_S : \mathbb{R}^n \mapsto \{-1, 1\}$.

Choice of Parameters: Let α, β be infinitesimally small, $\frac{1}{n^{1/2}} \ll \delta \ll \frac{1}{n^{1/3}}$.

Test: Run the following three tests with appropriate probabilities:

(1a) *Low Boundary Test with parameter α* is carried out with probability proportional to $\frac{1}{\sqrt{\alpha n}}$.

Pick a set $S \subseteq [N], |S| = n$ at random. Run the Low Boundary Test with parameter α on f_S.

(1b) *General Low Boundary Test with parameter α* is carried out with probability proportional to $\frac{1}{\sqrt{\alpha k}}$.

Pick a set $S \subseteq [N], |S| = n$ at random. Pick a constraint $C \subseteq S$ at random (note that $|C| = k$). Run the General Low Boundary Test with parameter α on subset C and f_S.

(2) *Constraint Test with Parameter β* is carried out with probability proportional to $\frac{1}{\beta/\sqrt{k}}$.

Pick a set $S \subseteq [N], |S| = n$ at random. Pick a constraint $C \subseteq S$ at random. Run the Constraint Test for constraint C and parameter β on f_S.

(3) *Consistency Test with Parameter δ* is carried out with probability proportional to $\frac{1}{\delta\sqrt{\delta n}}$.

Remark 7.1. *Note that the probability with which each test is performed is inversely proportional to the rejection probability of the test in Lemmas [4.1, 5.2], 5.1 and 6.1 respectively. These are the rejection probabilities in the "completeness case" (see below). Thus, in the completeness case, the different tests contribute equally towards the overall rejection probability. In the soundness case, it is enough to show that for any integral solution, for at least one of the tests, the rejection probability is significantly larger than that in the completeness case.*

Next we describe how the n_0-round Lasserre integrality gap for approximate real linear equations instance in Section 3 can be transformed into a feasible solution for t-rounds of Lasserre SDP for our Boolean 2Lin instance. We also analyze the objective value achieved by the Lasserre solution. While we do not analyze the soundness of our construction, in Section 8 we argue soundness against select global strategies.

Let $\{V_{S,\alpha}\}$ be the vector solution for approximate real linear equations instance. Let us denote the vector solution for our Boolean 2Lin instance by $\{U_{T,\beta}\}$. Let T be a set of at most t variables from the Boolean 2Lin instance. Let the i^{th} variable in T correspond to a pair (S_i, x_i) where $S_i \subseteq [N], |S| = n$ and $x_i \in \mathbb{R}^n$, i.e. the variable appears in the block S_i and corresponds to the point $x_i \in \mathbb{R}^n$ in that block. Let $S = \bigcup_{i=1}^{|T|} S_i$, so that $|S| \leq t \cdot n = n_0$. The set S will be referred to as the super-block corresponding to set T. The main observation is that every assignment $\alpha \in \{-1,1\}^{|S|}$ to the variables in S induces an assignment $\alpha(T) \in \{-1,1\}^{|T|}$ for the variables in T as follows: an assignment $\alpha \in \{-1,1\}^{|S|}$ induces, by restriction, assignments $\sigma_i : S_i \mapsto \{-1,1\}$ to the blocks, which in turn induce assignments $f_{\sigma_i}(x_i)$ to the points $x_i \in \mathbb{R}^n$ via the encodings $f_{\sigma_i} : \mathbb{R}^n \mapsto \{-1,1\}$. This yields the induced assignment $\alpha(T)$ is claimed. With this observation in mind, for every $\beta \in \{-1,1\}^{|T|}$, let

$$U_{T,\beta} = \sum_{\alpha:\alpha(T)=\beta} V_{S,\alpha}.$$

For every T, by orthogonality,

$$\sum_{\beta\in\{-1,1\}^{|T|}} \|U_{T,\beta}\|_2^2 = \sum_{\beta\in\{-1,1\}^{|T|}} \sum_{\alpha:\alpha(T)=\beta} \|V_{S,\alpha}\|_2^2 = \sum_{\alpha} \|V_{S,\alpha}\|_2^2 = 1.$$

The local distributions associated with $\{U_{T,\beta}\}$ assign β to T with probability $\|U_{T,\beta}\|_2^2$. As we show below, the consistency conditions of the solution $\{U_{T,\beta}\}$ follow from the consistency conditions of the solution $\{V_{S,\alpha}\}$.

Orthogonality: Let T be a set of at most t variables from our Boolean 2Lin instance. Let $\beta_1 \in \{-1,1\}^{|T|}$, $\beta_2 \in \{-1,1\}^{|T|}$ be distinct assignments to the variables of T. Let $S \subseteq [N]$ be the super-block associated with T. For every assignments α_1, α_2 to S such that $\alpha_1(T_1) = \beta_1$ and $\alpha_2(T_2) = \beta_2$, it holds that α_1, α_2 disagree. Hence, from the feasibility of the original solution,

$\langle V_{S,\alpha_1}, V_{S,\alpha_2}\rangle = 0$. Therefore,

$$\begin{aligned}
\langle U_{T,\beta_1}, U_{T,\beta_2}\rangle &= \langle \sum_{\alpha_1:\alpha_1(T)=\beta_1} V_{S,\alpha_1}, \sum_{\alpha_2:\alpha_2(T)=\beta_2} V_{S,\alpha_2}\rangle \\
&= \sum_{\alpha_1:\alpha_1(T)=\beta_1} \sum_{\alpha_2:\alpha_2(T)=\beta_2} \langle V_{S,\alpha_1}, V_{S,\alpha_2}\rangle \\
&= 0.
\end{aligned}$$

Consistency: Let T be a set of at most $t-1$ variables from the Boolean 2Lin instance. Let $p \notin T$ be an additional variable. Let $S, S^+ \subseteq [N]$ be the super-blocks associated with T and $T \cup \{p\}$ respectively. Note that either $S^+ = S$ or $S^+ = S \cup S'$ for some $S' \subseteq [N]$, $|S'| = n$. If $S^+ = S$, then

$$\begin{aligned}
U_{T,\beta} &= \sum_{\alpha:\alpha(T)=\beta} V_{S,\alpha} \\
&= \sum_{\alpha:\alpha(T\cup\{p\})=\beta\cup\{+1\}} V_{S,\alpha} + \sum_{\alpha:\alpha(T\cup\{p\})=\beta\cup\{-1\}} V_{S,\alpha} \\
&= U_{T\cup\{p\},\beta\cup\{+1\}} + U_{T\cup\{p\},\beta\cup\{-1\}}.
\end{aligned}$$

If $S^+ = S \cup S'$, then

$$\begin{aligned}
U_{T,\beta} &= \sum_{\alpha:\alpha(T)=\beta} V_{S,\alpha} \\
&= \sum_{\alpha:\alpha(T)=\beta} \sum_{\alpha'} V_{S\cup S',\alpha\cup\alpha'} \\
&= \sum_{\alpha:\alpha(T)=\beta} \sum_{\alpha':\alpha'(p)=+1} V_{S\cup S',\alpha\cup\alpha'} + \sum_{\alpha:\alpha(T)=\beta} \sum_{\alpha':\alpha'(p)=-1} V_{S\cup S',\alpha\cup\alpha'} \\
&= U_{T\cup\{p\},\beta\cup\{+1\}} + U_{T\cup\{p\},\beta\cup\{-1\}}.
\end{aligned}$$

Non-negativity: For any sets T_1, T_2 of up to t variables and assignments β_1, β_2 to them, letting S_1, S_2 be the super-blocks associated with them, we have,

$$\begin{aligned}
\langle U_{T_1,\beta_1}, U_{T_2,\beta_2}\rangle &= \langle \sum_{\alpha_1:\alpha_1(T_1)=\beta_1} V_{S_1,\alpha_1}, \sum_{\alpha_2:\alpha_2(T_2)=\beta_2} V_{S_2,\alpha_2}\rangle \\
&= \sum_{\alpha_1:\alpha_1(T_1)=\beta_1} \sum_{\alpha_2:\alpha_2(T_2)=\beta_2} \langle V_{S_1,\alpha_1}, V_{S_2,\alpha_2}\rangle \geq 0.
\end{aligned}$$

Completeness: The completeness of our construction follows from the completeness of the approximate real linear equations instance in Theorem 7, and Lemmas [4.1, 5.2], 5.1 and 6.1 analyzing the boundary, constraint and consistency tests, respectively. We elaborate a bit more.

Consider a hypothetical scenario that the instance in Theorem 6 and Theorem 7 has a perfectly satisfying (global) assignment. Considering its restrictions to blocks, we have: (1) for each block S, an assignment $\sigma(S)$ such that (2) $\sigma(S)$ satisfies all the constraints C that appear inside S and (3) the assignments $\sigma(S)$ and $\sigma(R)$ for any two blocks are consistent, i.e. they agree on $S \cap R$. In this scenario, letting each function f_S to be the correct encoding $f_{\sigma(S)}$, the failure probability of all the tests is bounded as in Lemmas [4.1, 5.2], 5.1 and 6.1.

Of course, the instance in Theorem 6 and Theorem 7 is highly unsatisfiable and the scenario is impossible. Still, the main point is that the Lasserre SDP solution to the instance effectively pretends that the hypothetical scenario holds. Namely, we have (1) for each block S, a set of assignments $\sigma(S)$ (the "local" distribution is uniform on this set) such that (2) every assignment $\sigma(S)$ satisfies all the constraints C that appear inside S and (3) sampling a random assignment τ for the block $S \cup T$ and letting $\sigma(S) = \tau|_S$ and $\sigma(R) = \tau|_R$ yields assignments to blocks S and R that are consistent.

Moreover, there are vectors $V_{S,\sigma(S)}$ that satisfy all the Lasserre feasibility conditions. Now, to each block S, instead of assigning an encoding $f_{\sigma(S)} : \mathbb{R}^n \mapsto \{-1, 1\}$, we effectively assign a "vector super-position" of such encodings, informally written as

$$\sum_{\sigma(S)} f_{\sigma(S)} \cdot V_{S,\sigma(S)}.$$

To be precise, if a typical variable in the Boolean 2LIN instance is denoted by a pair (S, x), then

$$U_{(S,x),\{+1\}} = \sum_{\sigma(S):f_{\sigma(S)}(x)=+1} V_{S,\sigma(S)},$$

$$U_{(S,x),\{-1\}} = \sum_{\sigma(S):f_{\sigma(S)}(x)=-1} V_{S,\sigma(S)}.$$

We need to show that the SDP solution achieves an objective value that is same as the failure probability of the tests in Lemmas [4.1, 5.2], 5.1 and 6.1. We demonstrate this for the low boundary test (i.e. Lemmas [4.1) and the others are treated similarly.

The low boundary test (for some fixed block S) picks two points $x, y \in \mathbb{R}^n$ and rejects if $f_S(x) \neq f_S(y)$. The analogue of the rejection probability from the viewpoint of the SDP objective is

$$\|U_{\{(S,x),(S,y)\},\{+1,-1\}}\|^2 + \|U_{\{(S,x),(S,y)\},\{-1,+1\}}\|^2,$$

or more precisely, the expectation of this expression over the choice of x and y. Using the feasibility conditions, this is same as

$$\langle U_{(S,x),\{+1\}}, U_{(S,y),\{-1\}}\rangle + \langle U_{(S,x),\{-1\}}, U_{(S,y),\{+1\}}\rangle.$$

Using the expression for the vector $U_{(S,x),\{+1\}}$ and others as observed, this is same as

$$\sum_{\sigma(S):f_{\sigma(S)}(x)\neq f_{\sigma(S)}(y)} \|V_{S,\sigma(S)}\|^2.$$

Taking the expectation over the choice of x and y, the SDP objective is

$$\sum_{\sigma(S)} \|V_{S,\sigma(S)}\|^2 \cdot \Pr_{x,y}\left[f_{\sigma(S)}(x) \neq f_{\sigma(S)}(y)\right].$$

Now we observe that $\sum_{\sigma(S)} \|V_{S,\sigma(S)}\|^2 = 1$ and $\Pr_{x,y}\left[f_{\sigma(S)}(x) \neq f_{\sigma(S)}(y)\right]$ is the rejection probability of the test as in Lemma 4.1.

8 Soundness Against Potential Counterexamples

We now look at a few potential counterexamples to the soundness analysis and observe that these do not (as far as we see) pose danger to the soundness of the construction. The first three examples demonstrate that each of the three tests is necessary, i.e., for any two of the tests, there is an integral solution whose failure probability on these two tests is same as the completeness case. The fourth example is open-ended and less concrete. In the following, we focus only on the tests **(1a)**, **2**, and **3**, and ignore the test **(1b)**.[4]

Example 1

It is possible to have an integral solution whose failure probability of the Low Boundary Test and the Consistency Test is $O(\sqrt{\alpha n})$ and $O(\delta\sqrt{\delta n})$, i.e., same as in the completeness case. Simply fix an arbitrary global assignment $\tau : [N] \mapsto \{-1, 1\}$ and for every $S \subseteq [N], |S| = n$, let f_S be the linear interval function of the assignment $\tau|_S$. This solution looks like the intended solution in the completeness case to the Low Boundary Test and the Consistency Test, implying that their failure probabilities are the same as in the completeness case. However the Constraint Test fails with probability $\Omega(\beta)$. For a typical constraint C, according to Equation (6), $\sum_{j=1}^{k} \epsilon_j \cdot b_C[j] \cdot \tau(C[j]) \geq \Omega(\sqrt{k})$. Looking at the proof of Lemma 5.1 then, with $\sigma = \tau|_S$, we have $\langle \sigma, \mathbf{v}_C \rangle = \sum_{j=1}^{k} \frac{\epsilon_j \cdot b_C[j]}{\sqrt{k}} \cdot \sigma(C[j])$ is $\Omega(1)$ in magnitude, hence $\langle \sigma, y - x \rangle$ is typically $\Omega(\beta)$ in magnitude, and hence the Constraint Test rejects with probability $\Omega(\beta)$. This rejection probability is significantly higher (by a large multiplicative factor) than the failure probability $O(\frac{\beta}{\sqrt{k}})$ achieved in the completeness case.

Example 2

It is possible to have an integral solution whose failure probability of the Low Boundary Test is $O(\sqrt{\alpha n})$ and that of the Constraints Test is $O(\frac{\beta}{\sqrt{k}})$, i.e., the same as in the completeness case. Fix an arbitrary assignment $\tau : [N] \mapsto \{-1, 1\}$. For every set $S \subseteq [N], |S| = n$, write $S = A \cup (S \setminus A)$ where $A = \cup_{C \subseteq S} C$ is the set of all variables that occur in constraints that are contained in S. Note that for a typical S, the number of constraints inside S is close to the expected number of $m \approx \eta^k M = \eta^k TN$ and hence $|A| \approx k\eta^k TN \ll |S| = \eta N$ since $\eta \approx \frac{1}{T^{25}}$. Moreover, (for the typical S) these constraints are disjoint. Let $\sigma(S) : S \mapsto \{-1, 1\}$ be an assignment that equals the global assignment $\tau|_{S \setminus A}$ on $S \setminus A$ and on A, the assignment is chosen to be a satisfying assignment for every constraint. Now let f_S be the linear interval function of the assignment $\sigma(S)$. Again, this solution looks like the intended solution in the completeness case to the Low Boundary Test, its failure probability being $O(\sqrt{\alpha n})$, and also to the Constraints Test, its failure probability being $O(\frac{\beta}{\sqrt{k}})$. However since the local assignments $\sigma(S)$ are not globally consistent, the Consistency Test fails with constant probability. Consider typical sets S, R such that $|S| = n, |R| = n, |S \cap R| \approx (1 - \delta)n$. Then $\sigma(S)$ and $\sigma(R)$ are (likely) inconsistent on variables in $C \cap (S \cap R)$ where C is any constraint such that $C \subseteq S, |C \cap (S \setminus R)| = 1$ (and the symmetric case with the roles of S, R switched). The number of all such variables is $t \approx (k - 1) \cdot 2\delta km$. Since $\sqrt{t} = \Omega(\sqrt{\delta n \cdot k^2 \eta^{k-1} T}) = \Omega(1)$, the Consistency Test fails with probability $\Omega(1)$ (note that we think of η, k, T as constants and $\delta n \gg 1$). On the other hand,

[4] The question here is whether the test **(1b)** is needed towards the soundness of our construction. We believe that it is indeed needed, but the counter-example demonstrating it is complicated and less clear, and is omitted from this note.

the failure probability of the Consistency Test is $O(\delta\sqrt{\delta n}) \ll 1$ in the completeness case. This is where we need the fact that the coordinates in $S\Delta R$ are chosen to be small (i.e., at most $s = \Theta(\delta)$ in magnitude) instead of choosing them from the standard Gaussian.

Example 3

It is possible to have an integral solution whose failure probability of the Consistency Test is $O(\delta\sqrt{\delta n})$ and that of the Constraints Test is $O(\frac{\beta}{\sqrt{k}})$, i.e., the same as the completeness case. We note that in the example below, strictly speaking the functions f_S are not necessarily folded, but it does not seem justified to rule out the example on this ground. The example however can be ruled out on the grounds that the Low Boundary Test fails with much higher probability than the completeness case.

In order to define the function $f_S : \mathbb{R}^n \mapsto \{-1, 1\}$, for every set $A \subseteq [n]$, let $\mathcal{P}_A \subseteq \mathbb{R}^n$ be the region where the coordinates in $[n] \setminus A$ are at most $s = \Theta(\delta)$ in magnitude and the coordinates in A have magnitude larger than s. Clearly, $\{\mathcal{P}_A\}_{A\subseteq[n]}$ defines a partition of $\mathbb{R}^n$. Note that since a typical point $z \in \mathcal{N}^n$ has about δn coordinates that are at most s in magnitude, almost all the measure of $\mathcal{N}^n$ is contained in regions $\mathcal{P}_A$ with $|A| \geq (1 - 5\delta)n$.

For any region $\mathcal{P}_A$, let $\Gamma_A = \cup_{C\subseteq A} C$, i.e., subset of the variables in all constraints contained inside A. Fix an arbitrary global assignment $\tau : [N] \mapsto \{-1, 1\}$. Let $\sigma(A) : S \mapsto \{-1, 1\}$ be an assignment that equals $\tau|_{S\setminus\Gamma_A}$ on $S \setminus \Gamma_A$ and on Γ_A chosen so that all constraints therein are satisfied (with some globally fixed assignment to each constraint). Note that since $|A| \geq (1-5\delta)n$, essentially all the constraints contained in S are in fact contained in A (we will ignore those which aren't) and are disjoint. Finally, define the function $f_S(x)$ to be the function that equals the interval function of the assignment $\sigma(A)$ on region $\mathcal{P}_A$.

Observe first that the failure probability of the Constraints Test is $O(\frac{\beta}{\sqrt{k}})$ since (almost) all constraints are satisfied. Now consider a Consistency Test between input z in block S and input w in block R. Note that $z|_{S\cap R} = w|_{S\cap R}$ and that the coordinates of z in $S \setminus R$ and those of w in $R \setminus S$ are at most s in magnitude. Thus if $\mathcal{P}_A$ denotes the region containing z, then $A = S \cap R$ and A is precisely the set of coordinates of z that have magnitude larger than s. Moreover Γ_A is then obtained by including all constraints contained in A. Similarly if $\mathcal{P}_B$ denotes the region containing w, then $B = S\cap R$ and B is precisely the set of coordinates of w that have magnitude larger than s. Moreover Γ_B is then obtained by including all constraints contained in B. Since $z|_{S\cap R} = w|_{S\cap R} = y$ say, it follows that $A = B$ and $\Gamma_A = \Gamma_B$. The assignments $\sigma(A)$ and $\sigma(B)$ then agree on $S\cap R$ and since the coordinates of z, w outside $S\cap R$ are at most s in magnitude, the Consistency Test accepts with probability $O(\delta\sqrt{\delta n})$.

Thankfully, we observe that for the function f_S defined as above, the Low Boundary Test fails with probability $\Omega(\eta^{k-1}T\sqrt{\alpha}\ n)$ which is much higher than $O(\sqrt{\alpha n})$ in the completeness case (noting that η, k, T are thought of as constants). Consider the Low Boundary Test on the function $f_S : \mathbb{R}^n \mapsto \{-1, 1\}$ with $(1 - \alpha)$-correlated points $x, y \in \mathcal{N}^n$. Since for each coordinate i, x_i and y_i differ typically by $\sqrt{\alpha}$ and α is infinitesimally small, with probability $\Omega(\sqrt{\alpha}\ n)$, there is precisely one coordinate j such that $|x_j| > s$ and $|y_j| \leq s$ or vice versa (let's say the former holds). Hence $x \in \mathcal{P}_A$ and $y \in \mathcal{P}_B$ with $B \subseteq A, A \setminus B = \{j\}$. With probability $\approx \eta^{k-1}T$, the index j is contained in a constraint $C \subseteq A$. The constraint C is included in Γ_A but not in Γ_B. Thus the assignments $\sigma(A), \sigma(B)$ (likely) disagree on variables in $C \setminus \{j\}$ and the boundary test fails with constant probability. The overall failure probability is $\Omega(\eta^{k-1}T\sqrt{\alpha}\ n)$ as claimed.

Example 4

In the examples above, the functions f_S are (or composed of several) linear interval functions. Towards a counter-example, one may use functions of the type $\mathsf{interval}(g)$ where g is a degree d polynomial. We do not know if this really helps towards a counter-example. If g is a degree d polynomial, then $\mathsf{interval}(g)$ may have Gaussian boundary/surface-area of $\Omega(d\sqrt{n})$ and the Low Boundary Test may fail with probability $\Omega(d\sqrt{\alpha n})$ (as opposed to the failure probability $O(\sqrt{\alpha n})$ in the completeness case), rendering the example void when d is large.

Acknowledgements

Many thanks to Madhur Tulsiani and Pratik Worah for discussion that led to the Lasserre integrality gap for approximate real linear equations problem. Thanks to everyone with whom we discussed related things over the years.

References

[1] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, 1998.

[2] S. Arora and S. Safra. Probabilistic checking of proofs: a new characterization of NP. *Journal of the ACM*, 45(1):70–122, 1998.

[3] B. Barak, F. G. S. L. Brandão, A. W. Harrow, J. A. Kelner, D. Steurer, and Y. Zhou. Hypercontractivity, sum-of-squares proofs, and their applications. In *Proc. 44th ACM Symp. on Theory of Computing*, pages 307–326, 2012.

[4] B. Barak, P. Gopalan, J. Håstad, R. Meka, P. Raghavendra, and D. Steurer. Making the long code shorter. In *Proc. 53rd IEEE Symp. on Foundations of Computer Science*, pages 370–379, 2012.

[5] B. Barak, M. Hardt, I. Haviv, A. Rao, O. Regev, and D. Steurer. Rounding parallel repetitions of unique games. In *Proc. 49th IEEE Symp. on Foundations of Computer Science*, pages 374–383, 2008.

[6] B. Barak, P. Raghavendra, and D. Steurer. Rounding semidefinite programming hierarchies via global correlation. In *Proc. 52nd IEEE Symp. on Foundations of Computer Science*, pages 472–481, 2011.

[7] M. Bellare, O. Goldreich, and M. Sudan. Free bits, PCPs, and nonapproximability—towards tight results. *SIAM Journal on Computing*, 27(3):804–915, 1998.

[8] M. Blum, M. Luby, and R. Rubinfeld. Self-testing/correcting with applications to numerical problems. *Journal of Computer and System Sciences*, 47(3):549–595, 1993.

[9] C. Borell. Geometric bounds on the ornstein-uhlenbeck velocity process. *Z. Wahrsch. Verw. Gebiete*, 70(1):1–13, 1985.

[10] J. Bourgain. On the distributions of the Fourier spectrum of Boolean functions. *Israel J. Math.*, 131:269–276, 2002.

[11] S. O. Chan. Approximation resistance from pairwise independent subgroups. In *Proc. 45th ACM Symp. on Theory of Computing*, pages 447–456, 2013.

[12] M. Charikar, K. Makarychev, and Y. Makarychev. Near-optimal algorithms for maximum constraint satisfaction problems. *ACM Transactions on Algorithms*, 5(3), 2009.

[13] I. Dinur and V. Guruswami. PCPs via low-degree long code and hardness for constrained hypergraph coloring. In *Proc. 54th IEEE Symp. on Foundations of Computer Science*, pages 340–349, 2013.

[14] R. O'Donnell E. Mossel and K. Oleszkiewicz. Noise stability of functions with low influences invariance and optimality. In *Proc. 46th IEEE Symp. on Foundations of Computer Science*, pages 21–30, 2005.

[15] U. Feige, S. Goldwasser, L. Lovasz, S. Safra, and M. Szegedy. Interactive proofs and the hardness of approximating cliques. *Journal of the ACM*, 43(2):268–292, 1996.

[16] Uriel Feige and László Lovász. Two-prover one-round proof systems: Their power and their problems (extended abstract). In *STOC*, pages 733–744, 1992.

[17] Dima Grigoriev. Linear lower bound on degrees of Positivstellensatz calculus proofs for the parity. *Theoret. Comput. Sci.*, 259(1-2):613–622, 2001.

[18] V. Guruswami and A.K. Sinop. Lasserre hierarchy, higher eigenvalues, and approximation schemes for quadratic integer programming with PSD objectives. In *Proc. 52nd IEEE Symp. on Foundations of Computer Science*, pages 482–491, 2011.

[19] J. Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, 182:105–142, 1999.

[20] J. Håstad. Some optimal inapproximability results. *Journal of the ACM*, 48(4):798–859, 2001.

[21] S. Khot. Improved inapproximability results for maxclique, chromatic number and approximate graph coloring. In *Proc. 42nd IEEE Symp. on Foundations of Computer Science*, pages 600–609, 2001.

[22] S. Khot. On the power of unique 2-prover 1-round games. In *Proc. 34th ACM Symp. on Theory of Computing*, pages 767–775, 2002.

[23] S. Khot. On the unique games conjecture (invited survey). In *IEEE Conference on Computational Complexity*, pages 99–121, 2010.

[24] S. Khot, G. Kindler, E. Mossel, and R. O'Donnell. Optimal inapproximability results for MAX-CUT and other two-variable CSPs? *SIAM Journal on Computing*, 37(1):319–357, 2007.

[25] S. Khot and S. Safra. A two prover one round game with strong soundness. In *Proc. 52nd IEEE Symp. on Foundations of Computer Science*, pages 648–657, 2011.

[26] S. Khot and N. K. Vishnoi. The unique games conjecture, integrality gap for cut problems and embeddability of negative type metrics into l_1. In *Proc. 46th IEEE Symp. on Foundations of Computer Science*, pages 53–62, 2005.

[27] Subhash Khot and Dana Moshkovitz. Np-hardness of approximately solving linear equations over reals. In *STOC*, pages 413–420, 2011.

[28] Subhash Khot and Rishi Saket. Sdp integrality gaps with local ℓ_1-embeddability. In *FOCS*, pages 565–574, 2009.

[29] A. Klivans, R. O'Donnell, and R. Servedio. Learning geometric concepts via gaussian surface area. In *Proc. 49th IEEE Symp. on Foundations of Computer Science*, pages 541–550, 2008.

[30] P. Raghavendra. Optimal algorithms and inapproximability results for every csp? In *Proc. 40th ACM Symp. on Theory of Computing*, pages 245–254, 2008.

[31] Prasad Raghavendra and David Steurer. Integrality gaps for strong sdp relaxations of unique games. In *FOCS*, pages 575–585, 2009.

[32] A. Rao. Parallel repetition in projection games and a concentration bound. *SIAM Journal on Computing*, 40(6):1871–1891, 2011.

[33] R. Raz. A parallel repetition theorem. In *SIAM Journal on Computing*, volume 27, pages 763–803, 1998.

[34] A. Samorodnitsky and L. Trevisan. Gowers uniformity, influence of variables, and PCPs. *SIAM Journal on Computing*, 39(1):323–360, 2009.

[35] Grant Schoenebeck. Linear level lasserre lower bounds for certain k-csps. In *FOCS*, pages 593–602, 2008.

[36] Luca Trevisan. Approximation algorithms for unique games. *Theory of Computing*, 4(1):111–128, 2008.

[37] Luca Trevisan. On Khot's unique games conjecture. *Bull. Amer. Math. Soc. (N.S.)*, 49(1):91–111, 2012.

[38] M. Tulsiani. CSP gaps and reductions in the lasserre hierarchy. In *Proc. 41st ACM Symp. on Theory of Computing*, pages 303–312, 2009.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de