

MIT Open Access Articles

Tight SoS-Degree Bounds for Approximate Nash Equilibria

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Harrow, Aram et al. "Tight SoS-Degree Bounds for Approximate Nash Equilibria." Leibniz International Proceedings in Informatics 22 (2016): 22:2-22:25.

As Published: <http://dx.doi.org/10.4230/LIPIcs.CCC.2016.22>

Publisher: Schloss Dagstuhl--Leibniz-Zentrum fuer Informatik

Persistent URL: <http://hdl.handle.net/1721.1/109803>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: Creative Commons Attribution 4.0 International License



Tight SoS-Degree Bounds for Approximate Nash Equilibria

Aram W. Harrow¹, Anand V. Natarajan², and Xiaodi Wu³

1 Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, USA

2 Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, USA

3 Department of Computer and Information Science, University of Oregon, Eugene, USA

Abstract

Nash equilibria always exist, but are widely conjectured to require time to find that is exponential in the number of strategies, even for two-player games. By contrast, a simple quasi-polynomial time algorithm, due to Lipton, Markakis and Mehta (LMM), can find *approximate* Nash equilibria, in which no player can improve their utility by more than ϵ by changing their strategy. The LMM algorithm can also be used to find an approximate Nash equilibrium with near-maximal total welfare. Matching hardness results for this optimization problem were found assuming the hardness of the planted-clique problem (by Hazan and Krauthgamer) and assuming the Exponential Time Hypothesis (by Braverman, Ko and Weinstein).

In this paper we consider the application of the sum-squares (SoS) algorithm from convex optimization to the problem of optimizing over Nash equilibria. We show the first unconditional lower bounds on the number of levels of SoS needed to achieve a constant factor approximation to this problem. While it may seem that Nash equilibria do not naturally lend themselves to convex optimization, we also describe a simple LP (linear programming) hierarchy that can find an approximate Nash equilibrium in time comparable to that of the LMM algorithm, although neither algorithm is obviously a generalization of the other. This LP can be viewed as arising from the SoS algorithm at $\log n$ levels – matching our lower bounds. The lower bounds involve a modification of the Braverman-Ko-Weinstein embedding of CSPs into strategic games and techniques from sum-of-squares proof systems. The upper bound (i.e. analysis of the LP) uses information-theory techniques that have been recently applied to other linear- and semidefinite-programming hierarchies.

1998 ACM Subject Classification F.2 [Theory of Computation] Analysis of Algorithms and Problem Complexity

Keywords and phrases Approximate Nash Equilibrium, Sum of Squares, LP, SDP

Digital Object Identifier 10.4230/LIPIcs.CCC.2016.22

1 Introduction

1.1 Motivation and background

Game theory, broadly speaking, seeks to explain the decision-making of interacting, self-interested agents. Mathematically it can be seen as a generalization of optimization problems in which the goal is to minimize or maximize some function. Instead each player wishes to maximize its own payoff, which in general will depend on the actions of the other players as well. The standard solution concept here is a Nash equilibrium, meaning a set



© Aram W. Harrow, Anand V. Natarajan, and Xiaodi Wu;
licensed under Creative Commons License CC-BY

31st Conference on Computational Complexity (CCC 2016).

Editor: Ran Raz; Article No. 22; pp. 22:1–22:25



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



of (uncoordinated) mixed strategies for which no player can unilaterally increase their own payoff by changing strategy. Here “mixed strategy” means a probability distribution over the basic “pure strategies” of the game and “uncoordinated” means that these distributions are uncorrelated. A related notion is a “correlated equilibrium” where again players cannot unilaterally improve their payoffs, but this time a “signal” random variable is broadcast to all the players who are free to choose their strategy based on the common signal; e.g. consider the role of a traffic light in suggesting that one car stop and another car drive.

As with optimization problems, the practical applicability of Nash equilibria and correlated equilibria depend on their computational complexity. Nash’s 1950 existence theorem proved that Nash equilibria exist under very general conditions [24, 25] but turning the proof into an algorithm results in an exponential runtime. Indeed, finding a Nash equilibrium is PPAD-complete, meaning that it is as hard as solving an abstract fixed-point problem [10]. If we instead consider the problem of maximizing a linear function (e.g. total payoff) over the space of Nash equilibria then the problem becomes NP-complete [10]. The difference in complexity (PPAD vs NP) reflects the fact that the former is a problem of searching for a solution that is known to exist whereas the latter problem is to determine whether a solution exists. Finding a correlated equilibrium, by contrast, can be achieved in poly time with linear programming.

One could reasonably argue that exact Nash equilibria are implausible models of rational behavior and that deviating from an equilibrium strategy might only happen in practice when the benefit is greater than zero by some non-negligible amount. An ϵ -approximate Nash equilibrium (aka. ϵ -ANE) is thus defined to be a set of uncorrelated strategies for which no player can improve their payoff by more than ϵ by changing strategies. It turns out that the complexity of finding ϵ -ANEs is significantly lower than that of exact Nash equilibria. In 2003, Lipton, Markakis and Mehta [21] gave an algorithm for finding an ϵ -ANE in quasipolynomial time, e.g. $n^{O(\log(n)/\epsilon^2)}$ for two-player games where both strategy sets have size n . Their algorithm was based on enumeration over a suitably chosen net of strategies. This net-based framework has been refined in [2, 6] to yield PTASs in some special cases. On the hardness side, Braverman, Ko and Weinstein [9] recently showed that finding the best (i.e. highest total payoff) ϵ -ANE in $n^{o(\log n)}$ time would violate the Exponential Time Hypothesis (ETH). (The ETH posits that 3-SAT instances on n variables require $\exp(\Omega(n))$ time.)

1.2 Main results

Our paper investigates ϵ -ANE from the perspective of convex optimization. Since the sets of Nash equilibria and ϵ -ANE are not convex, it is not immediately obvious how to relate the problem of finding an ANE to a convex optimization problem. To this end, we can consider the convex hull of all ϵ -ANE, for a given set of payoff functions. Optimizing a linear function over one of these sets is equivalent to optimizing a linear function over the set of ϵ -ANE. Note that these sets can be far from the much-more-tractable set of correlated equilibria; additionally, even though it is easy to test whether a strategy is an ϵ -NE, this does not extend to a test for whether a strategy is in the convex hull of ϵ -ANE (and likewise for Nash equilibria). Indeed, standard arguments mean that optimizing a linear function over or testing membership in these sets have approximately the same complexity [13].

Our first main result is a no-go theorem for a family of approximation algorithms based on semidefinite programming (SDP), called the sum of squares (SoS) hierarchy. In particular we show that in order to achieve a constant-factor approximation for 0.1-ANE, one must go to a level of at least $\Omega(\log n / \text{poly log log } n)$ in the SoS hierarchy (where n is the size of the game).

This translates to an SDP of size $\Omega(n^{\log n})$. Unlike all previous results on the hardness of NE and ANE, our result is unconditional; i.e. does not depend on any assumptions about the hardness of 3-SAT, planted clique or other problems. Our result is also surprising in part because the best results for planted clique, an apparently comparable problem [16], only extend to ruling out SDPs arising from the SoS hierarchy of size $O(n^4)$ [17].

► **Theorem 1.1** (informal). *Given a game of size N with payoffs bounded by a constant, deciding whether either (1) there exists a Nash equilibrium with average payoff ≥ 1 or (2) all 0.1-approximate Nash equilibria have average payoff at most δ requires at least $\Omega(\log N / \text{poly log log } N)$ levels in the SoS hierarchy.*

Our proof makes use of a classic result of Grigoriev [12], showing hardness for the problem 3XOR in the SoS model. We use reductions (with some properties as discussed below) to extend the hardness of 3XOR to ANE. There have been quite a few examples using reductions to prove the hardness in the SoS model (e.g., [34, 27]). However, each proof requires slightly different properties about reductions and there is no explicit unified framework for doing so. We follow a recent result in [15], which aims to serve as one such framework to facilitate the proof of hardness in the SoS model.

To obtain integrality gaps in the SoS model, one needs to show that (a) the SoS solution believes the value is large up to some high level (degree), and (b) the true value is actually small. To achieve (a), we follow the notion of *low-degree reductions* [15], in which one requires the reductions preserve a SoS solution for the reduced problems with almost the same value and a small amount of loss of the degree¹. To achieve (b), we need the reductions to have some kind of *soundness*.

Our specific reduction is a variant of the one used in proving the ETH-hardness of ANE by Braverman, Ko and Weinstein [9], with tweaks to ensure it has *low-degree*, *soundness*, and *embedding* properties. Our hardness analysis is also inspired by a recent result of ours [15] that extends the SoS lower bounds to quantum information problems. This connection is natural given the intimate relationship between [1] and quantum information, which serves as the first step in the reduction of [9].

We also make use of the explicit construction of a two-player strategic game for which the optimal payoff of a Nash equilibrium is related to the value of a constraint satisfaction problem, in the second step of [9]. In their game, the two players each specify an assignment to a subset of \sqrt{n} variables from the CSP, and receive a payoff if they are consistent with each other and satisfy the clauses. There are further penalties that ensure that each player must choose their subset close to uniformly at random. It is then proven that if the underlying CSP is satisfiable, the optimal Nash equilibrium is an “honest” strategy, where both players answer according to a fixed assignment to the variables. This establishes reductions with good completeness and soundness from CSPs to HONESTNASH of optimizing over honest strategies to this game. We prove that their reduction is also low-degree and pseudosolution-preserving, which allows us to obtain an SoS hardness result for HONESTNASH.

However, this game is not convenient for obtaining hardness for APPROXIMATENASH, since not all honest strategies are in fact Nash equilibria. The problem is that Alice and Bob are punished for honest strategies that do not satisfy clauses. Additionally, the game depends on underlying CSP. We fix both problems at once by giving Alice and Bob a payoff

¹ This roughly refers to the “Vector Completeness” in [34] and “SoS Completeness” in [27]. It explicitly requires the existence of a mapping that is a polynomial of low-degree, which maps a SoS solution of the original problem to a SoS solution of the reduced one.

simply for answering consistently, independent of the clauses. We then offload all dependence on clause structure into our objective function.² Our objective function is 0 whenever Alice or Bob output sets of variables that do not contain enough clauses and otherwise equals the fraction of satisfied clauses. Maximizing this objective function over ϵ -ANE is then roughly equivalent to maximizing the number of satisfied clauses, which completes our reduction.

Our second main result is a linear programming (LP) lift of NE with dimension $n^{O(\log(n)/\epsilon^2)}$, matching our above lower bound. We describe lifts formally in Section 4, but intuitively an approximate lift of dimension D is a polytope in \mathbb{R}^D whose projection onto some lower-dimensional space includes all Nash equilibria, but ideally not too many additional points. In Section 4 we prove the following theorem (more formally stated in Theorem 4.1).

► **Theorem 1.2 (informal).** *Consider a two-player game with strategy sets of size n_1, n_2 and payoffs in $[-1, 1]$. There exists a polytope in \mathbb{R}^D with $D = \exp(O(\log(n_1) \log(n_2)/\epsilon^2))$ such that its projection onto $\mathbb{R}^{n_1 n_2}$ contains all Nash equilibria and is contained in the ϵ -neighborhood of the convex hull of all ϵ -ANE. This polytope has an explicit efficient description, so that we can find an ϵ -ANE in time $\text{poly}(D)$; or if f is an efficiently computable concave function, we can estimate its maximum value over the NE efficiently.*

The same result holds for m -player games where each player has a strategy set of size n if we set $D = \exp(O(m^3 \ln^2(n)/\epsilon^2))$.

In fact, the set of correlated equilibria can already be seen as an LP relaxation of Nash equilibria, since Nash equilibria can be alternately defined as correlated equilibria that are product distributions. This relaxation can be useful [28], but in general correlated equilibria can be far from Nash equilibria. Our lift can be thought of as a systematic hierarchy of successive refinements of the set of correlated equilibria. Inspired by [31, 7], our idea is to replace the player Bob with k replicas: Bob-1, Bob-2, ..., Bob- k . We will impose the constraint that the strategies of Alice and Bob- j form a correlated equilibrium, even when conditioned on the strategies of Bobs-1, ..., $j-1$. This approach prevents Alice from being simultaneously correlated with all of the Bobs. Indeed, if Alice were correlated with Bob-1, then conditioning on his strategy would reduce Alice's entropy. Continuing in this way we find that Alice must have low correlation with *most* of the Bobs, implying that if we choose j randomly from $\{1, \dots, k\}$ and condition on the random strategies of Bobs-1, ..., $j-1$, then the resulting distribution on Alice and Bob- j will be nearly product. This means the resulting correlated equilibrium can be easily rounded to an ϵ -ANE.

The LP we obtain can be viewed as arising from the SoS hierarchy at level $\log N$, with the omission of the positive semidefinite constraint. Thus, our analysis of the LP also implies that SoS is able to solve approximate APPROXIMATE NASH to constant accuracy at level $\log N$, matching our SoS lower bound.

Our LP relaxation is not the first approximation algorithm for ANE. In fact, a nearly identical runtime was achieved by Lipton, Markakis and Mehta in 2003 [21] using an algorithm that exhaustively searched over a set of sparse strategies. Using Chernoff bounds it is possible to show that any NE can be approximated in this way. By now we have seen a series of examples where net-based algorithms and LP/SDP hierarchies give very similar approximation guarantees, often in the regime of PTASs or quasipolynomial-time algorithms. These examples are summarized in Table 1 of [8] and include (1) optimizing polynomials over

² This has the advantage of making our results compatible with the framework of extension complexity, where one considers a polytope of feasible solutions (e.g. the matching polytope) that is independent of the function being maximized.

the simplex, (2) optimizing polynomials over the unit sphere, (3) free two-player games, (4) unique games, (5) small-set expansion and (6) optimizing linear functions over unentangled quantum states (see [8] for specific references). Despite this series of coincidences, it is still an open question to find a common explanation of the performance of both types of algorithms. Indeed our algorithm differs from that of LMM in ways that suggest there is no obvious way to map one onto the other. If there is a sparse NE then LMM will find it, while our algorithm may not. On the other hand, while both algorithms can freely add linear constraints (e.g. on the total payoff), only ours can add convex constraints, such as maximizing entropy. Indeed, if there exists an NE with entropy $\geq c \log(n)$ then our algorithm will always find a nearby ϵ -ANE with entropy $\geq (c - O(\epsilon)) \log(n)$, while by construction, LMM will only find ϵ -ANE with entropy $\leq \log \log(n/\epsilon) + O(1)$. We further compare the algorithms in Section 5.

1.3 Open problems

- *Extension complexity.* Our results show limitations on approximating Nash equilibria using the SoS hierarchy. But what about more general SDPs? Is it possible to find an $n^{\Omega(\log n)}$ lower bound on the extension complexity of ϵ -ANE? The approach of Lee, Raghavendra, and Steurer [20] does not apply directly here because our problems do not have the same self-embedding property that CSPs do. However, it seems likely that the [20] framework can be extended to cover ϵ -ANE.
- *Special cases.* While our results address the complexity of ϵ -ANE in the worst case, there are many special cases where it should be possible to find more efficient convex relaxations. Under various conditions on the payoff matrices, net-based algorithms can run more quickly [2, 6]. Without any further modification, our algorithm is already more effective when Alice has low entropy in all correlated equilibria. This condition can be checked quickly (since correlated equilibria form a polytope and entropy is a concave function) but appears incomparable to the conditions under which [21, 2, 6] outperform their worst-case guarantees. More generally we would like to know scenarios under which our algorithms or variants of them can perform significantly better.
- *Semidefinite and convex constraints.* A related question is whether SDP or other convex constraints give additional benefits not already captured by LP hierarchies. We could also add concave objective functions, such as entropy maximization. Do these have further application?
- *Search vs. optimization.* A major theme in work on the complexity of finding Nash equilibrium is the distinction between NP and PPAD. PPAD is an example of TFNP (“total function NP”) which is the class of search problems for which an answer is guaranteed to exist and can be efficiently verified. All known algorithms for finding Nash equilibria can also perform (or approximate) the optimization versions of the problem, and the extension complexity model (by contrast with earlier hardness-based lower bounds) collapses the difference between the search and optimization versions. Is there a natural computational model which separates the complexity of these tasks?
- *Densest subgraph.* Techniques for proving upper and lower bounds on the complexity of the approximate Nash equilibrium problem have also been applied to the problem of finding the densest k -subgraph of a given graph. The lower bound of [16] rules out a PTAS for additive approximations to this problem on bipartite graphs, while the algorithm of [6] achieves a quasipolynomial algorithm. Can we match these results in the SoS or extension complexity settings? One barrier is that the hardness of [16] uses a reduction from the planted clique problem, for which SoS lower bounds are not as well understood as they are for CSPs.

1.4 Overview

In Section 2, we introduce some basic definitions and give a more technical overview of our contribution. In Section 3, we prove a lower bound on SoS relaxations for the set of two-player approximate Nash equilibria. Following this, in Section 4, we introduce an LP relaxation for approximate Nash that matches our lower bound as well as handling multiplayer games, and in Section 5, we compare its performance to the LMM algorithm. The appendices contain further background on SoS proofs (Appendix A), reductions for SDPs (Appendix B) and information theory (Appendix C).

2 Definitions and Preliminaries

2.1 Games

Consider a two-player game with strategy sets $[n_1], [n_2]$ and payoff vectors $f_1, f_2 \in \mathbb{R}^{n_1 n_2}$. A Nash equilibrium is a pair of probability distributions $p_1 \in \Delta_{n_1}, p_2 \in \Delta_{n_2}$ such that

$$\langle e_x \otimes p_2, f_1 \rangle \leq \langle p_1 \otimes p_2, f_1 \rangle \quad \forall x \in [n_1] \quad (2.1a)$$

$$\langle p_1 \otimes e_y, f_2 \rangle \leq \langle p_1 \otimes p_2, f_2 \rangle \quad \forall y \in [n_2] \quad (2.1b)$$

Here $[n] = \{1, \dots, n\}$, $\Delta_n = \{p \in \mathbb{R}^n : p(x) \geq 0, \sum_x p(x) = 1\}$, e_x is the vector with a one in position x and zeroes elsewhere and $p \otimes q$ is the vector with x, y entry equal to $p(x)q(y)$.

Nash proved that (2.1) always has a solution (known as a Nash equilibrium, or NE), but finding one is known to be PPAD-complete (or NP-complete in some cases when additional constraints or optimizations are added) [10]. For this reason, it is natural to consider instead approximate NE. Assume for the rest of the paper that $\max_i \max_x |f_i(x)| \leq 1$. We say that the distributions p_1, p_2 (or equivalently the joint distribution $p_1 \otimes p_2$) are an ϵ -approximate NE (or ϵ -ANE) if they satisfy

$$\langle e_x \otimes p_2, f_1 \rangle \leq \langle p_1 \otimes p_2, f_1 \rangle + \epsilon \quad \forall x \in [n_1] \quad (2.2a)$$

$$\langle p_1 \otimes e_y, f_2 \rangle \leq \langle p_1 \otimes p_2, f_2 \rangle + \epsilon \quad \forall y \in [n_2] \quad (2.2b)$$

From these expressions, we can see that the problem of optimizing over Nash equilibria is a polynomial optimization problem, where the variables are the probabilities p and q . The constraints are the simplex constraints and the Nash conditions ((2.2)).

We consider also *correlated equilibria*, first proposed by Aumann in 1974 [3]. Let q^{XY} denote a probability distribution in $\Delta_{n_1 n_2}$. Then we say that q is a *correlated equilibrium* if q satisfies the following analogue of (2.1):

$$\sum_{y \in [n_2]} q(x, y)(f_1(x', y) - f_1(x, y)) \leq 0 \quad \forall x, x' \in [n_1] \quad (2.3a)$$

$$\sum_{x \in [n_1]} q(x, y)(f_2(x, y') - f_2(x, y)) \leq 0 \quad \forall y, y' \in [n_2] \quad (2.3b)$$

Since (2.3) is an LP, we can find correlated equilibria efficiently; i.e. in time $\text{poly}(n_1, n_2)$.

While our hardness results will focus on the two-player case, we also describe algorithms for finding ϵ -ANE for games with more than two players. Consider an m -player game, where the players have strategy sets $S_1 = [n_1], \dots, S_m = [n_m]$ (with $S = S_1 \times \dots \times S_m$), and payoff tensors $f_1, \dots, f_m \in \mathbb{R}^{n_1} \otimes \dots \otimes \mathbb{R}^{n_m}$. This means that if players use mixed strategies $p_1 \in \Delta_{n_1}, \dots, p_m \in \Delta_{n_m}$, then player i receives payoff

$$\langle p_1 \otimes \dots \otimes p_m, f_i \rangle = \sum_{x=(x_1, \dots, x_m) \in S} p_1(x_1) \dots p_m(x_m) f_i(x).$$

Let $N := \prod_{i=1}^m n_i$. A distribution $p \in \Delta_N$ is a correlated equilibrium if

$$\sum_{x_{-i} \in S_{-i}} p(x_i, x_{-i}) (f_i(x'_i, x_{-i}) - f_i(x_i, x_{-i})) \leq 0 \quad \forall i \in [m], \forall x_i, x'_i \in S_i. \quad (2.4)$$

Here the notation S_{-i} indicates the strategy set $S_1 \times \cdots \times S_{i-1} \times S_{i+1} \times \cdots \times S_m$ of all players except player i , and similarly $x_{-i} := (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_m)$. A Nash equilibrium is a correlated equilibrium that is also a product distribution: i.e. such that $p = p^{X_1} \otimes \cdots \otimes p^{X_m}$.

2.2 Norms

Define the 1-norm and ∞ -norm of vectors to be

$$\|v\|_1 := \sum_x |v(x)| \quad \text{and} \quad \|v\|_\infty = \max_x |v(x)|.$$

For two probability distributions p, q , the 1-norm distance $\|p - q\|_1$ is also called the *variational distance*, because of the following special case of Hölder's inequality:

$$\langle v, w \rangle \leq \|v\|_1 \|w\|_\infty \quad \text{with equality holding iff } v = \lambda w \text{ for } \lambda \geq 0. \quad (2.5)$$

2.3 Optimization

Both our upper and lower bounds apply to relaxations of the convex optimization problem APPROXIMATE NASH_ϵ . The promise version of this problem is as follows:

- **Definition 2.1.** The problem $(a, b) - \text{APPROXIMATE}\text{NASH}_{\epsilon, m}(f)$ is to determine, given a game G over m players, where each player has at most n deterministic strategies, and a function h (called the “externality”) mapping strategies to real numbers, whether either
- there exists an exact Nash equilibrium strategy for G for which $h \geq a$, or
 - for every ϵ -approximate Nash equilibrium to G , h is at most b .

given the promise that one of these cases holds. Here $b < a$ are constant parameters. As an important special case, we write APPROXIMATE NASH_ϵ to refer to the case when $m = 2$.

3 SoS Lower Bound

Braverman et al. [9] show a reduction from 3SAT to approximate Nash equilibrium, which shows hardness for this problem conditional on ETH. We are able to use their reduction to show an unconditional hardness result for the SoS hierarchy for approximate Nash. We follow the approach of [15], who show how to SoS hardness for several continuous-variable optimization problems arising in quantum information. The proof consists of two steps: a *pseudo-solution-preserving* reduction from 3XOR over n variables to the intermediate optimization problem HONEST NASH over $\{\pm 1\}^n$, followed by a reduction from HONEST NASH to approximate APPROXIMATE NASH . The following schematic diagram illustrates this:

$$3\text{XOR} \implies \text{HONEST}\text{NASH} \implies \text{APPROXIMATE}\text{NASH}.$$

- The problem HONEST NASH is a polynomial optimization problem over the boolean hypercube $\{\pm 1\}^n$. The objective function $h_\phi(x)$ is the total expected externality of two players in a particular strategic game, whose strategies are specified by the input variables x ; the externality function is induced by a CSP instance ϕ . We consider the strategic game introduced by [9], for which this objective function is a polynomial of degree $\tilde{O}(\sqrt{n})$. This game, in turn, is based on a free game introduced by [1].

3.1 Framework of Deriving SoS Lower Bounds

Our proof makes extensive use of reductions between optimization problems. Here, we will briefly give definitions of some useful notions, while we defer a full description to Appendix B. Throughout this section, we will use the so-called ± 1 notation for boolean variables: that is, we encode FALSE as 1 and TRUE as -1 .

We derive all of our integrality gaps from the following foundational result of Grigoriev for the problem 3XOR (defined in Section B.3):

► **Proposition 3.1** (Theorem 3.1 of [5], due to Grigoriev). *For any $\epsilon > 0$, for every n there exists a 3XOR instance Φ_n with n variables and $m = O(n/\epsilon^2)$ clauses, such that $\text{OPT}(\Phi_n) \leq \frac{1}{2} + \epsilon$, but there exists a degree- $\Omega(n)$ value-1 pseudo-solution $\tilde{\mathbb{E}}$.*

Here “value 1” means that for every clause $x_i x_j x_k = a_{ijk}$, it holds that $\tilde{\mathbb{E}}[(x_i x_j x_k - a_{ijk})p(x)] = 0$ for all polynomials $p(x)$ with degree at most $d - 3$.

The instance of Φ_n produced by Grigoriev has a constraint graph which is a good expander. However, there is no upper bound on the degree of the constraint graph, i.e. the number of clauses each variable can participate in. We remedy this issue by transforming the instance to an instance Φ'_n of the problem 3XOR+EQ, where we allow both 3XOR constraints and equality constraints between pairs of variables.

► **Proposition 3.2.** *For every n , there exists an instance Φ'_n of 3XOR+EQ on $O(n)$ variables where the constraint graph is $(\delta n, \alpha)$ -expanding and has degree at most d , for constants $\delta < 1, \alpha > 1, d$, such that the maximum fraction of clauses satisfiable is $\omega(\Phi'_n) \leq \frac{1}{2} + \epsilon$. Furthermore, there exists a degree- $\Omega(n)$, value-1 pseudo-solution to Φ'_n . That is, there exists a pseudo-expectation operator $\tilde{\mathbb{E}}[\cdot]$ with degree $D = \Omega(n)$, such that $\tilde{\mathbb{E}}[C(x)q(x)] = 1$ for every clause $C(x)$ in the instance Φ'_n and every polynomial $q(x)$ with degree $\deg(C(x)q(x)) \leq D$.*

Proof. We start with the instance Φ_n in Proposition 3.1, and then apply the degree reduction procedure of [29] to produce the new instance Φ'_n . This procedure consists of replacing each high-degree variable in the original instance with many copies, connected by equality constraints laid out according to an expander graph. It is shown in [29] that this procedure has constant soundness, so $\omega(\Phi_n) \leq \frac{1}{2} + \epsilon$ for some constant ϵ . Now, to produce the desired pseudosolution, we define $\tilde{\mathbb{E}}[p(x)]$ for any polynomial $p(x)$ to be what we get by first *identifying* all of the replicated variables with each other, and then evaluating the pseudoexpectation according to the operator $\tilde{\mathbb{E}}[\cdot]$ produced by Proposition 3.1. ◀

3.2 Consistent Sample Game

Inspired by the construction in [9], we now present a game called the Consistent Sample Game, whose Nash equilibria will be easy to characterize. As a warm-up, consider the following simple game

► **Definition 3.3** (Consistent Bit Game). The consistent bit game is a two-player strategic game, where Alice and Bob are each allowed to play a single bit 0 or 1. They win if they their bits agree and lose otherwise.

It is easy to see that this game has exactly two Nash equilibria: either Alice and Bob both play 0, or both play 1. Our consistent sample game is a scaled up version of this game, where Alice and Bob have access to n -bit strings. However, the strategies they play consist of assignments to a subset of the variables of size \sqrt{n} . To force the players to choose their subsets with close to uniform probability, we also add a zero-sum uniformity test. This test

and its analysis are one of the main technical contributions of [9], and fortunately we will be able to mostly reuse their analysis without change.

► **Definition 3.4** (Consistent Sample Game). For a given size n , the consistent sample game $G_{n,k,\ell,d}$ is specified by the following:

- The set of Alice’s possible pure strategies consists of: all tuples (S, s) consisting of a subset S of k variables and an assignment s to these variables, and all subsets Y of $\rho\sqrt{n}$ variables. We refer to the former as “tuple strategies” and the latter as “subset strategies.”
- The set of Bob’s possible pure strategies consist of: all tuples (T, t) consisting of a subset T of ℓ variables and an assignment t to the variables in T , and all subsets Z of $\rho\sqrt{n}$ clauses. We likewise refer to these two types of strategies as “tuple strategies” and “subset strategies.”
- The payoff matrix has a block structure. If both players play with tuple strategies (S, s) and (T, t) , the payoff for Alice is 1, and the payoff for Bob is

$$f'(S, T, s, t) = \begin{cases} \beta_T f(S, T, s, t) & \text{if } |S \cap \mathcal{N}(T)| > \frac{d}{10(\epsilon^*)^2}, \\ 0 & \text{otherwise} \end{cases},$$

where the notation $\mathcal{N}(T)$ denotes the neighbors of the variables in T in the constraint graph, i.e. the set of clauses involving variables in T . Here, $f(S, T, s, t)$ is 1 if there are no inconsistencies in Alice and Bob’s assignments to the variables and 0 otherwise, and $\beta_T \equiv \frac{1}{\Pr[|S \cap \mathcal{N}(T)| > \frac{d}{10(\epsilon^*)^2}]}$.

- If Alice plays with a tuple strategy (S, s) and Bob plays with a subset strategy Z , then if $S \cap Z \neq \emptyset$, Bob receives a payoff of K and Alice receives $-K$. Likewise, if Bob plays with a tuple strategy (T, t) and Alice plays with a subset strategy Y , then if $T \cap Y \neq \emptyset$, Alice receives a payoff of K and Bob receives $-K$.
- If both players play with a subset strategy, they both receive a payoff of 0.

In the above, $K > 1$ and $\epsilon^* < \frac{1}{2}$ are constant parameters, and $\rho = (\epsilon^*)/(c_2 \cdot K)$ where c_2 is an appropriately chosen constant.

As in Braverman et al., we choose k and ℓ to be $\Theta(\sqrt{n})$. The parameter d is equal to the degree of the constraint graph of the CSP ϕ , which we will use later to construct our externality function. Henceforth, we will denote the game simply by G_n . An important difference between our game and the one in Braverman et al. is that in our construction, when both players use tuple strategies, Alice always receives a payoff of 1.

We now define a Boolean optimization problem HONESTNASH by considering a restricted subset of strategies which we call “honest strategies.” We assume a fixed CSP instance ϕ that is known to both players.

► **Definition 3.5.** For every $x \in \{\pm 1\}^n$, we define the *honest strategy according to x* for G_n by the following prescription

- Alice follows a mixed strategy: she chooses her subset of variables S to be $k/3$ clauses chosen uniformly at random from ϕ , and her assignment s to be that given by x .
- Bob also follows a mixed strategy: he chooses his subset T of ℓ variables uniformly at random from ϕ , and his assignment t to be that given by x .

► **Definition 3.6.** For every 3XOR instance ϕ , the problem HONESTNASH $_\phi$ is a Boolean optimization problem $\max_{x \in \{\pm 1\}^n} h_\phi(x)$, where the objective function $h_\phi(x)$ is the expected value of the following externality function over the honest strategy induced by x (note that in an honest strategy, Alice and Bob always play tuple strategies):

22:10 Tight SoS-Degree Bounds for Approximate Nash Equilibria

- The externality is 1 if both players' assignments are consistent with each other and all of Alice's assignments satisfy their respective clauses in ϕ .
- The externality is 0 otherwise.

This objective function is a degree- $O(\sqrt{n})$ polynomial in the variables x .

To see why the function h_ϕ is a degree $O(\sqrt{n})$ polynomial, note that h can be written as an average of terms, where each term corresponds to the externality for a specific choice of S and T . This is a boolean function depending on only the $O(\sqrt{n})$ variables that appear in S and T . So the whole function $h_\phi(x)$ can be written as a sum of terms for each S, T , each of which has degree $O(\sqrt{n})$. Below in the proof of theorem 3.8, we will give an explicit expression for $h_\phi(x)$.

► **Theorem 3.7.** *If ϕ is satisfiable, then there exists an honest strategy for G_n that is an exact Nash equilibrium, and achieves expected externality 1. Moreover, there are fixed constants $\delta, \epsilon^* < 1/2$ independent of ϵ , such that if at most $(1 - \delta)$ -fraction of the clauses of ϕ are satisfiable, then all ϵ^* -approximate Nash equilibria for G_n have expected externality at most $O(\epsilon)$.*

Proof. This is the main result of [9]; we need to argue that is preserved under our modification of the payoff function. Because of the similarity between our proof and theirs, we only sketch our proof. For the completeness case, the desired Nash equilibrium is simply the honest strategy playing according to the satisfying assignment of ϕ . The argument presented in Lemma 3.2 of [9] goes through without change.

For the soundness, we again follow the proof strategy of [9]. In particular, we note that the proof of Lemma 3.4, which states that all Nash equilibria must choose the subsets S, T with roughly uniform probability over the clauses in ϕ , holds unchanged with our payoff function. Thus, we can reproduce the argument of Lemma 3.5 with our modified payoff function, to upper bound Bob's externality by $O(\epsilon)$. Moreover, the same argument applied to our payoff function upper-bounds Alice's externality by $O(\epsilon)$. Thus, we obtain an average payoff of $O(\epsilon)$ as desired. ◀

► **Theorem 3.8.** *For ϕ be the 3XOR+EQ instance from Proposition 3.2, there exists a degree- $\Omega(n)$ pseudosolution to HONESTNASH with externality h_ϕ achieving value 1.*

Proof. Let $\tilde{\mathbb{E}}[\cdot]$ be the degree- $\Omega(n)$ pseudoexpectation operator associated with a value-1 pseudosolution for our 3XOR+EQ instance ϕ . (Such a pseudosolution exists by proposition 3.2.) We claim that this yields a value-1 pseudosolution for HONESTNASH. To prove this, let us examine the objective function $f_\phi(x)$ for this problem. First, define the polynomial

$$\text{AND}(x_1, x_2, \dots, x_k) = 1 + \frac{1}{2^{k-1}}(1 - x_1)(1 - x_2) \dots (1 - x_k)$$

This polynomial evaluates to -1 when all of the input variables are -1 , and 1 otherwise. Next, we define a polynomial function for $g_C(x)$ for each clause $C = ax_ix_jx_k$:

$$g_C(x) = 1 - \frac{1}{2}(x_ix_jx_k - a)^2$$

This evaluates to -1 when the clause is satisfied and 1 otherwise. By our definition of honest strategies, the consistency tests always pass with probability 1. So the objective function

$f_\phi(x)$ is a function of just the clauses $\{C_1, \dots, C_k\}$ appearing in the random sets of variables S and T :

$$f_\phi(x) = \frac{1}{2} - \frac{1}{2} \mathbb{E}_S \mathbb{E}_T \beta_T \text{AND}(g_{C_1}(x), g_{C_2}(x), \dots, g_{C_k}(x)) = \frac{1}{2} - \frac{1}{2^k} \mathbb{E}_S \prod_{i=1}^k (1 - C_i(x)).$$

Here the expectations are taken over the uniform distribution over sets S, T of the appropriate size. Note that $k = \Theta(\sqrt{n})$. Now, we know that $\mathbb{E}'[C_i(x)q(x)] = 0$ for all $q(x)$ such that $\deg(q(x)) \leq d$ where $d = \Omega(n)$. Therefore,

$$\mathbb{E}'[\text{AND}(g_{C_1}(x), g_{C_2}(x), \dots, g_{C_k}(x))] = 1.$$

So $\mathbb{E}'[f_\phi(x)] = \mathbb{E}_S \mathbb{E}_T \beta_T = 1$. ◀

3.3 Embedding HonestNash in ApproximateNash

We now pass from HONESTNASH to APPROXIMATENASH by broadening the space of strategies searched over to include *all* mixed strategies, not just honest ones. In order to preserve our SoS lower bound, it will help to show that every feasible point of HONESTNASH corresponds to a feasible point of APPROXIMATENASH achieving exactly the same value.

► **Theorem 3.9.** *Every honest strategy for G_n is a Nash equilibrium.*

Proof. First, we will show that there is no incentive for either player to switch to another tuple strategy. We will then invoke the soundness analysis of Braverman et al. [9] to show that there is no incentive to switch to subset strategies either.

For tuple strategies, there are two cases to consider. First, let's suppose we fix Alice's strategy and allow Bob's strategy to deviate. Bob's payoff depends only on whether Alice's clauses are satisfied, and whether Alice and Bob are inconsistent on any variables. If Bob were to deviate from the honest strategy, the number of satisfied clauses would be unaffected, and the chance of inconsistencies can only go up. So Bob has no incentive to deviate. Now, if we fix Bob's strategy and allow Alice to deviate, note that Alice's payoff is always 1 regardless of which strategy she chooses, so she has no incentive to deviate either.

Now, we need to show that neither party has an incentive to switch to a subset strategy. We will use the following fact, which is shown in the course of the proof of Lemma 3.2 of [9].

► **Fact 3.10.** *Suppose Alice plays honestly and Bob plays with a (deterministic) subset strategy Z . Then Bob's expected payoff is upper bounded by*

$$v = K \mathbb{E}_{S \sim \mathcal{U}} \mathbf{1}(S \cap Z \neq \emptyset) \leq \frac{2}{0.9 \cdot c_2}.$$

So for $c_2 > 2/(0.9\epsilon^*)$, Bob has no incentive to deviate. A symmetric argument applies to Alice. ◀

We note that the previous theorem would be false in the original version of the game given by [9], without the modification to Alice's payoff. This is because for any honest assignment to the variables with 3XOR value at least $\frac{1}{2} + \epsilon$, Alice can find some subset S' of \sqrt{n} clauses that are perfectly satisfied by that assignment. So Alice will always have an incentive to switch to the deterministic strategy that always answers with S' : this strategy would achieve a payoff of 1 for Alice. We are able to remove this incentive by making Alice's payoff 1 independent of her choice of strategy.

It turns out that the preceding argument to show that each honest strategy is a Nash equilibrium can be itself converted into a sum of squares proof. This enables us to “lift” the SoS pseudosolution we constructed for HONESTNASH to one for APPROXIMATE NASH, and achieve our main lower bound.

► **Theorem 3.11.** *There exists a constant $\epsilon^* < \frac{1}{2}$ such that for any constant $\epsilon < \frac{1}{2}$, there exists a game G_n of size $N = O(n^{\sqrt{n}})$ and externality function h such that the expected value of h is at most ϵ over all ϵ^* -Nash equilibria. At the same time, there is a degree- $\Omega(\log N)$ pseudosolution that satisfies all the Nash equilibrium constraints for G_n and achieves externality value 1.*

Proof. Choose G_n to be the consistent sample game, and take the externality function h to be the one induced by the 3XOR+EQ instance from Proposition 3.2. Then it follows from Theorem 3.7 that the expected value of h is at most ϵ over all ϵ^* -Nash equilibria. Now, to construct the pseudosolution, first, let us define a polynomial formulation of the problem APPROXIMATE NASH. Our variables will be $p_{(S,s)}$, representing the probability that Alice plays the tuple strategy (S, s) ; p_Y , representing the probability that Alice plays the subset strategy Y ; $q_{(T,t)}$ representing the probability that Bob plays the tuple strategy (T, t) , and q_Z representing the probability that Bob plays the subset strategy Z . Let $\tilde{\mathbb{E}}[\cdot]$ be the pseudosolution for HONESTNASH derived in Theorem 3.8, which is defined up to degree $D = \Omega(\sqrt{n})$. We need to lift this to a pseudoexpectation $\tilde{\mathbb{E}}'[\cdot]$ on variables p, q . We do so as follows: when evaluating a pseudoexpectation $\tilde{\mathbb{E}}'[p \dots pq \dots q]$ of a monomial term, first perform the following substitutions, and then evaluate the resulting polynomial in x according to $\tilde{\mathbb{E}}[\cdot]$:

- Replace p_Y or q_Z with 0 (this is because honest strategies have no support over subset strategies).
- Replace $p_{(S,s)}$ by $c_a \text{AND}(x_{S_1} = s_1, \dots, x_{S_k} = s_k)$, where c_a is the probability of Alice choosing the subset S . This term evaluates to c_a if the assignment in s matches the assignment in x , and 0 otherwise.
- Likewise, replace $q_{(T,t)}$ by $c_b \text{AND}(x_{T_1} = t_1, \dots, x_{T_k} = t_k)$, where c_b is the probability of Bob choosing the subset T .

This pseudoexpectation “automatically” satisfies the positive semidefinite constraint, since it arises from a valid pseudoexpectation for x . Moreover, under the substitution process, the degree of a polynomial can only increase by a multiplicative fact of at most $O(\sqrt{n})$, so $\tilde{\mathbb{E}}'[\cdot]$ is defined up to degree $\Omega(D/\sqrt{n}) = \Omega(\sqrt{n})$. It remains to check that it satisfies the Nash constraints, and gives a high objective value for the externality function. To check the former, we need to show that the pseudoexpectation of the advantage gained by switching to any other strategy, multiplied by any squared polynomial, is non-positive. First, let us consider Bob. Suppose Bob switches to a tuple strategy (T', t') . The advantage gained can be written as

$$\text{advantage} = \sum_{(S,s),(T,t)} p_{(S,s)} q_{(T,t)} (f_{(S,s),(T',t')} - f_{(S,s),(T,t)}). \quad (3.1)$$

The term $f_{(S,s),(T,t)}$ checks whether the assignments s and t are consistent, so it is a degree- $O(\sqrt{n})$ polynomial in the variables x (essentially the AND of a number of equality checks). We want to check that

$$\tilde{\mathbb{E}}'[\text{advantage} \cdot P^2(p, q)] \leq 0$$

for an arbitrary polynomial $P(p, q)$. By the construction of the honest strategies, all (S, s, T, t) in the support of the strategy are consistent and so $\tilde{\mathbb{E}}'[(f_{(S,s),(T,t)} - 1)Q(p, q)] = 0$ for any

such (S, s, T, t) and any polynomial $Q(p, q)$. Another way to see this is that any *inconsistent* tuples (S, s, T, t) should disagree on some bit x_i , meaning that one AND term contains a $1 + x_i$ and the other AND term contains a $1 - x_i$. In this case

$$\tilde{\mathbb{E}}'[p_{(S,s)}q_{(T,t)}Q(p, q)] = \tilde{\mathbb{E}}[(1 - x_i)(1 + x_i)Q'(x)] = 0$$

where $Q'(x)$ is some other polynomial of x_1, \dots, x_n and the last equality comes from the $x_i^2 = 1$ constraint satisfied by the original pseudoexpectation $\tilde{\mathbb{E}}$. Thus the second term in (3.1) always evaluates to $\tilde{\mathbb{E}}'[P^2]$ under the pseudoexpectation. On the other hand we claim the first term is $\leq \tilde{\mathbb{E}}'[P^2]$. Here the constraint $\text{AND}^2 = 1$ is implied by the $x_i^2 = 1$ constraints, thus we have the SOS proof

$$1 - \text{AND} = \frac{1 - 2 \text{AND} + \text{AND}^2}{2} = \frac{(1 - \text{AND})^2}{2} \geq 0.$$

This implies that $\tilde{\mathbb{E}}[\text{AND}^2 P^2] \leq \tilde{\mathbb{E}}[P^2]$, as desired.

Next, suppose Bob switches to a subset strategy Z . Then his advantage is

$$\text{advantage} = \sum_{(S,s),(T,t)} p_{(S,s)}q_{(T,t)}(f_{(S,s),Z} - f_{(S,s),(T,t)}). \quad (3.2)$$

Note that the first term in the difference is independent of s . Indeed the second term is as well, since $f_{(S,s),(T,t)} = 1$ whenever $p_{(S,s)}q_{(T,t)}$ are not constrained to equal 0 by our construction. Thus we can rewrite the advantage as

$$\text{advantage} = \sum_{S,T} p_S q_T (K \mathbf{1}(S \cap Z \neq \emptyset) - 1), \quad (3.3)$$

with $p_S := \sum_s p_{(S,s)}$ and likewise for q_T . We now claim that these terms factor out of a pseudoexpectation; i.e. that

$$\tilde{\mathbb{E}}'[p_S P(p, q)] = \tilde{\mathbb{E}}'[p_S] \tilde{\mathbb{E}}'[P(p, q)] \quad (3.4)$$

for any $P(p, q)$ (and likewise for q_T). To see this observe that our substitution replaces p_S with

$$\sum_{s_1, \dots, s_k} c_a \left(1 + \frac{1}{2^{k-1}} (1 - s_1 x_1)(1 - s_2 x_2) \dots (1 - s_k x_k)\right) = 2^k c_a.$$

Thus for any polynomial $P(p, q)$ we have

$$\tilde{\mathbb{E}}'[\text{advantage} \cdot P^2(p, q)] = \tilde{\mathbb{E}}'[\text{advantage}] \tilde{\mathbb{E}}'[P^2(p, q)]. \quad (3.5)$$

The first term is nonpositive due to Fact 3.10 and the second term is nonnegative, as we have argued above, because it equals $\tilde{\mathbb{E}}[\tilde{P}^2(x)]$ for some polynomial \tilde{P} and because $\tilde{\mathbb{E}}$ is a pseudoexpectation.

We conclude that $\tilde{\mathbb{E}}'$ satisfies the Nash equilibrium constraints. It achieves externality value 1 because it inherits the property from $\tilde{\mathbb{E}}$ of satisfying all the 3XOR constraints. ◀

4 Approximate Nash equilibria via linear programming

In this section we will describe an LP hierarchy for ANE.

We first introduce a new form of equilibrium, called an ϵ -correlated equilibrium. If p is a correlated equilibrium for some m -player game f then we say it is an ϵ -correlated equilibrium if

$$\|p^{X_1 \dots X_m} - p^{X_1} \otimes \dots \otimes p^{X_m}\|_1 \leq \epsilon. \quad (4.1)$$

Denote the set of Nash equilibria for game f by \mathcal{N}_f , the ϵ -correlated Nash equilibria by $\tilde{\mathcal{N}}_{f,\epsilon}$ and the ϵ -ANE by $\mathcal{N}_{f,\epsilon}$. Of course $\mathcal{N}_f = \mathcal{N}_{f,0} = \tilde{\mathcal{N}}_{f,0}$. These will allow us to state a result that will imply Theorem 1.2.

► **Theorem 4.1.**

1. Fix a two-player game $f = (f_1, f_2)$ with strategy sets of size n_1, n_2 and all payoffs in $[-1, 1]$. Let $D = \exp(O(\ln(n_1) \ln(n_2)/\epsilon^2))$. There exists a polytope $P_f \subset \mathbb{R}^D$ such that its projection onto $\mathbb{R}^{n_1 n_2}$, called Q_f , satisfies

$$\mathcal{N}_f \subseteq Q_f \subseteq \text{conv}(\tilde{\mathcal{N}}_{f,\epsilon}).$$

P_f is defined by poly(D) explicit constraints and thus we can test membership in it in time poly(D).

2. Now fix an m -player game with strategy sets of size n_1, \dots, n_m and payoffs in $[-1, 1]$. Choose positive integers k_1, \dots, k_m . Then the same result holds with $D = n_1^{k_1} n_2^{k_2} \dots n_m^{k_m}$ and

$$\epsilon = \sqrt{2 \sum_{1 \leq i < j \leq m} \frac{\ln(n_i)}{k_j} \max_{\substack{i \in [m] \\ x \in S}} |f_i(x)|}. \quad (4.2)$$

If we specialize to $n_1 = \dots = n_m =: n$ then we can take $D = \exp(O(m^3 \ln^2(n)/\epsilon^2))$

► **Corollary 4.2.** Use the same parameters as in Theorem 4.1 and let h be an efficiently computable concave function such that $|h(p) - h(q)| \leq \eta \|p - q\|_1$ for p, q any pair of probability distributions over joint strategies. Then given some threshold T and in time $D^{O(n^2)}$ we can distinguish between the cases

- $\max_{p \in \mathcal{N}_f} h(p) \geq T$; or
- $\max_{p \in \mathcal{N}_{f,\epsilon}} h(p) \leq T - \epsilon$. (We could equivalently replace this with $\max_{p \in \tilde{\mathcal{N}}_{f,\epsilon}} h(p) \leq T - \epsilon$.)

Corollary 4.2 follows from Theorem 4.1 and the fact that optimizing over a convex set has a poly-time reduction to the problem of testing membership in that set [13, Theorem 4.3.2 and Remark 4.2.5].

4.1 Proof for two players

While the two-player proof is a special case of the multiplayer proof, and uses very similar ideas, the notation is much simpler and it is a good warmup to the general case.

First we observe that any ϵ -correlated equilibrium q can be rounded to an ϵ -approximate NE by replacing q with the pair of marginal distributions q^X, q^Y , i.e.

$$q^X := \sum_{x \in [n_1], y \in [n_2]} \langle e_x \otimes e_y, q \rangle e_x \quad \text{and} \quad q^Y := \sum_{x \in [n_1], y \in [n_2]} \langle e_x \otimes e_y, q \rangle e_y. \quad (4.3)$$

Call this “marginal rounding.”

In general marginal rounding can produce pairs of strategies that are far from equilibria. One situation in which it works well is when q is already nearly of product form; i.e. when $\|q - q^X \otimes q^Y\|_1$ is small.

► **Lemma 4.3.** *If q is a correlated equilibrium then (q^X, q^Y) is an ϵ -approximate NE for*

$$\epsilon = \|q - q^X \otimes q^Y\|_1 \cdot \max\{\|f_1\|_\infty, \|f_2\|_\infty\}. \quad (4.4)$$

Proof. From (2.5) we have

$$\langle q - q^X \otimes q^Y, f_1 \rangle \leq \|q - q^X \otimes q^Y\|_1 \|f_1\|_\infty \leq \epsilon.$$

Combining with (2.3a) yields (2.2a). Repeating the argument for f_2 yields (2.2b). ◀

To obtain uncorrelated q , we will consider a variant of correlated equilibria in which there are k copies of player 2 for some positive integer k . If $q \in \Delta_{n_1 n_2^k}$ then we can interpret q as a probability distribution on random variables X, Y_1, \dots, Y_k . We use the abbreviations:

$$Y_{<j} := Y_1, \dots, Y_{j-1}$$

$$Y_{>j} := Y_{j+1}, \dots, Y_k$$

$$Y_{-j} := Y_{<j}, Y_{>j}$$

For $y_{<j} \in [n_2]^{j-1}$, let $q_{Y_{<j}=y_{<j}}^{XY_j}$ be the distribution on XY_j obtained by conditioning on $Y_i = y_i$ for $i < j$. Explicitly

$$q_{Y_{<j}=y_{<j}}^{XY_j}(x, y_j) = \frac{q^{XY_{\leq j}}(x, y_{<j}, y_j)}{\sum_{x', y'_j} q^{XY_{\leq j}}(x', y_{<j}, y'_j)}. \quad (4.5)$$

Now define the “ k -extendable relaxation” of NE to be the following LP:

$$q \in \Delta_{n_1 n_2^k} \quad (4.6a)$$

$$q_{Y_{<j}=y_{<j}}^{XY_j} \text{ satisfies (2.3)} \quad \forall j \in [k], \forall y_{<j} \in [n_2]^{j-1} \text{ such that } q^{Y_{<j}}(y_{<j}) > 0 \quad (4.6b)$$

This is a linear program since (4.6b) is equivalent to the following uglier-but-manifestly-linear conditions:

$$\sum_{\substack{x \in [n_1] \\ y_j \in [n_2] \\ y_{>j} \in [n_2]^{k-j}}} q(x, y)(f_1(x', y_j) - f_1(x, y_j)) \leq 0 \quad \forall x' \in [n_1], \forall j \in [k], \forall y_{<j} \in [n_2]^{j-1} \quad (4.7a)$$

$$\sum_{\substack{x \in [n_1] \\ y_j \in [n_2] \\ y_{>j} \in [n_2]^{k-j}}} q(x, y)(f_2(x, y'_j) - f_2(x, y_j)) \leq 0 \quad \forall y_j \in [n_2], \forall j \in [k], \forall y_{<j} \in [n_2]^{j-1} \quad (4.7b)$$

In other words, we ask for a distribution on $XY_1 \dots Y_k$ such that each XY_j are in a correlated equilibrium even when conditioned on the actions of $Y_{<j}$. To see that this is indeed a relaxation, observe that if p_1, p_2 is a NE, then $q = p_1 \otimes p_2^{\otimes k}$ is a valid solution to (4.6). Along with Nash’s theorem, this also implies that the LP is always feasible.

The rounding algorithm is as follows.

1. Enumerate over all $j \in [k]$ and $y_{<j} \in [n]^{j-1}$.
2. For each $j, y_{<j}$, let $p_1 = q_{Y_{<j}=y_{<j}}^X$ and $p_2 = q_{Y_{<j}=y_{<j}}^{Y_j}$.
3. Output the (p_1, p_2) that is an ϵ -ANE for the lowest value of ϵ .

In the final step, we are using the fact that given p_1, p_2 , it is easy to find the smallest value of ϵ for which (p_1, p_2) is an ϵ -ANE.

► **Theorem 4.4.** *The above procedure returns a $\sqrt{\frac{2 \ln(n_1)}{k}}$ -approximate NE.*

22:16 Tight SoS-Degree Bounds for Approximate Nash Equilibria

Thus, setting $k = 2 \ln(n_1)/\epsilon^2$ yields an algorithm that finds an ϵ -ANE and runs in time $\text{poly}(mn^k) = \exp(O(\ln(n_1) \ln(n_2)/\epsilon^2))$. To prove Theorem 4.4 we will need a few basic facts from information theory, reviewed in Appendix C.

Proof of Theorem 4.4. Let $q^{XY_1 \dots Y_k}$ be a solution of (4.6). Observe that

$$\ln(n_1) \geq H(X)_q \geq I(X; Y_1 \dots Y_k)_q = \sum_{j=1}^k I(X; Y_j | Y_1 \dots Y_{j-1}). \quad (4.8)$$

Introduce the abbreviations $\alpha = (j, y_{<j})$ for $y_{<j} \in [n_2]^{j-1}$ and $q_\alpha = q^{XY_j | Y_{<j}=y_{<j}}$. Then we can rewrite (4.8) as

$$\mathbb{E}_{j \in [k]} \mathbb{E}_{y_{<j} \sim q^{Y_{<j}}} I(X; Y)_{q_\alpha} \leq \frac{\ln(n_1)}{k}. \quad (4.9)$$

Thus there is a choice of $\alpha = (j, y_{<j})$ for which $I(X; Y)_{q_\alpha} \leq \frac{\ln(n_1)}{k}$. By Pinsker's inequality (specifically (C.3)), we have

$$\|q_\alpha^{XY} - q_\alpha^X \otimes q_\alpha^Y\|_1 \leq \sqrt{\frac{2 \ln(n_1)}{k}}. \quad (4.10)$$

Finally (4.6b) forces q_α to be a valid correlated equilibrium so we can use Lemma 4.3 to obtain that (q_α^X, q_α^Y) is a $\sqrt{\frac{2 \ln(n_1)}{k}}$ -approximate NE. \blacktriangleleft

In the above algorithm, we could have chosen $j, y_{<j}$ randomly instead of enumerating over all possibilities. However, this would not improve the asymptotic runtime. We also could have replaced (4.6b) with the stronger constraint that $q_{Y_{-j}=y_{-j}}^{XY_j}$ is a correlated equilibrium, with no asymptotic increase in run-time, but also without any provable performance improvement.

4.2 Proof for multiplayer games

We can quantify the distance of a distribution p to a product distribution by using the *multipartite mutual information* (which was first proposed in 1954, but has since been reinvented multiple times [23, 35, 22])

$$I(X_1 : \dots : X_m)_p := \sum_{i=1}^m H(X_i)_p - H(X_1 \dots X_m)_p \quad (4.11a)$$

$$= D(p^{X_1 \dots X_m} \| p^{X_1} \otimes \dots \otimes p^{X_m}) \quad (4.11b)$$

$$= \sum_{i=2}^m I(X_{<i} : X_i)_p \quad (4.11c)$$

Due to (4.11b) and Pinsker's inequality (see (C.1) in Appendix C), we can use the multipartite mutual information to bound the distance of a distribution to product:

$$\|p^{X_1 \dots X_m} - p^{X_1} \otimes \dots \otimes p^{X_m}\|_1 \leq \sqrt{2I(X_1 : \dots : X_m)}. \quad (4.12)$$

To define the relaxation we will need to introduce some more notation. Let k_1, \dots, k_m be positive integers that we will choose later. Define $\vec{k} = (k_1, \dots, k_m)$ and $\vec{n} = (n_1, \dots, n_m)$. Introduce random variables Y_i^j with $i \in [m]$ and $j \in [k_i]$. Let $\vec{n}^{\vec{k}} := [n_1]^{k_1} \times \dots \times [n_m]^{k_m}$

and let $q \in \Delta_{\vec{n}_k}$ be a distribution on $Y := (Y_1^1, \dots, Y_m^{n_m})$. Define $\Phi_{\vec{k}} = [k_1] \times \dots \times [k_m]$. If $\phi = (\phi_1, \dots, \phi_m) \in \Phi_{\vec{k}}$, then we define

$$\begin{aligned} Y^\phi &:= Y_1^{\phi_1} \dots Y_m^{\phi_m} \\ Y_{<j}^{\phi_{<j}} &:= Y_1^{\phi_1} \dots Y_{j-1}^{\phi_{j-1}} \\ Y^{<\phi} &:= Y_1^{<\phi_1} \dots Y_m^{<\phi_m} \\ Y_{-j}^{<\phi_{-j}} &:= Y_1^{<\phi_1} \dots Y_{j-1}^{<\phi_{j-1}} Y_{j+1}^{<\phi_{j+1}} \dots Y_m^{<\phi_m} \\ Y_j &:= Y_j^1 \dots Y_j^{k_j} \end{aligned}$$

It will also be convenient to define

$$\vec{n}^{<\phi} := [n_1]^{\phi_1-1} \times \dots \times [n_m]^{\phi_m-1}$$

and likewise for $\vec{n}^\phi, \vec{n}^{\geq\phi}$, etc. We can now define the LP relaxation:

$$q \in \Delta_{\vec{N}} \tag{4.13a}$$

$$q_{Y_{<\phi}^{\phi}} \text{ is a correlated equilibrium} \quad \forall \phi \in \Phi_{\vec{k}}, \forall y_{<\phi} \in \vec{n}^{<\phi} \tag{4.13b}$$

Again the constraint (4.13b) is only defined when we condition on an event with positive probability.

The rounding algorithm is similar to the bipartite case:

1. Enumerate over all $\phi \in \Phi_{\vec{k}}$ and $y_{<\phi} \in [n_1]^{\phi_1-1} \times \dots \times [n_m]^{\phi_m-1}$.
2. For each $\phi, y_{<\phi}$, let

$$p = \bigotimes_{i=1}^m q_{Y_{<\phi}^{\phi_i}}^{Y_i^{\phi_i}}.$$

3. Output the p that is an ϵ -ANE for the lowest value of ϵ .

► **Theorem 4.5.** *The above procedure returns an ϵ -approximate NE, where*

$$\epsilon = \sqrt{2 \sum_{1 \leq i < j \leq m} \frac{\ln(n_i)}{k_j} \max_{\substack{x \in S \\ i \in [m]}} |f_i(x)|}. \tag{4.14}$$

Proof. We begin by bounding

$$\mathbb{E}_{\phi \in \Phi_{\vec{k}}} I(Y_1^{\phi_1} : \dots : Y_m^{\phi_m} | Y^{<\phi})_p \tag{4.15a}$$

$$= \mathbb{E}_{\phi \in \Phi_{\vec{k}}} \sum_{j=2}^m I(Y_{<j}^{\phi_{<j}} : Y_j^{\phi_j} | Y^{<\phi})_p \quad \text{using (4.11c)} \tag{4.15b}$$

$$= \sum_{j=2}^m \mathbb{E}_{\phi_{-j}} \mathbb{E}_{\phi_j} I(Y_{<j}^{\phi_{<j}} : Y_j^{\phi_j} | Y_j^{<\phi_j} Y_{-j}^{<\phi_{-j}})_p \quad \text{writing } \phi = (\phi_j, \phi_{-j}) \tag{4.15c}$$

$$= \sum_{j=2}^m \mathbb{E}_{\phi_{-j}} \frac{1}{k_j} I(Y_{<j}^{\phi_{<j}} : Y_j | Y_{-j}^{<\phi_{-j}})_p \quad \text{chain rule (C.6)} \tag{4.15d}$$

$$\leq \sum_{j=2}^m \frac{\ln(n_1 \dots n_{j-1})}{k_j} = \sum_{1 \leq i < j \leq m} \frac{\ln(n_i)}{k_j} \tag{4.15e}$$

Thus there exists a ϕ and a $y^{<\phi}$ for which

$$I(Y_1^{\phi_1} : \dots : Y_m^{\phi_m})_{p_{Y^{<\phi=y^{<\phi}}}} \leq \sum_{1 \leq i < j \leq m} \frac{\ln(n_i)}{k_j}. \quad (4.16)$$

Set $q^{X_1 \dots X_m} := p_{Y^{<\phi=y^{<\phi}}}^{Y_1^{\phi_1} \dots Y_m^{\phi_m}}$. Then (4.12) implies that

$$\|p^{X_1 \dots X_m} - p^{X_1} \otimes \dots \otimes p^{X_m}\|_1 \leq \sqrt{2 \sum_{1 \leq i < j \leq m} \frac{\ln(n_i)}{k_j}} = \epsilon. \quad (4.17)$$

Enumerating over all ϕ will find a ϕ satisfying (4.16) and thereby also (4.17). (While this suffices for our purposes, note that concavity of the square root means that (4.17) holds in expectation even if ϕ and $y^{<\phi}$ is randomly chosen.) By an easy generalization of Lemma 4.3, we conclude that the marginal distributions of p form an ϵ -approximate Nash equilibrium. ◀

► **Corollary 4.6.** *For an m -player game where each player has a strategy set of size n and $\max_{i \in [n]} \max_{x \in S} |f_i(x)| \leq 1$, an ϵ -ANE can be found in time $\exp(O(m^3 \ln^2(n)/\epsilon^2))$.*

Proof. Set $k_j = \kappa \sqrt{m-j}$ in Theorem 4.5 for κ to be chosen later. The error is

$$\leq \sqrt{2 \sum_{j=1}^m \frac{(m-j) \ln(n)}{\kappa \sqrt{m-j}}} \leq \sqrt{\frac{2m^{3/2} \ln(n)}{\kappa}}.$$

For this to be $\leq \epsilon$, we set $\kappa = \frac{2m^{3/2} \ln(n)}{\epsilon^2}$. The size of the LP is

$$n^{k_1 + \dots + k_m} = \exp\left(\ln(n) \kappa \sum_{j=1}^m \sqrt{m-j}\right) \leq \exp\left(\frac{2m^3 \ln^2(n)}{\epsilon^2}\right),$$

and the run-time of the algorithm is polynomial in this dimension. ◀

5 Comparison to the LMM algorithm for approximate NE

Lipton, Markakis and Mehta [21] gave a method to find an ϵ -ANE in quasipolynomial time; specifically $\exp(O(\log(n_1) \log(n_2)/\epsilon^2))$. Their strategy was to prove that for any NE (p_1, p_2) , there exists an ϵ -ANE (\hat{p}_1, \hat{p}_2) of the form

$$\hat{p}_1 = \frac{1}{|S_X|} \sum_{x \in S_X} e_x \quad \text{and} \quad \hat{p}_2 = \frac{1}{|S_Y|} \sum_{y \in S_Y} e_y \quad (5.1)$$

for some multisets S_X, S_Y satisfying $|S_X| = \lceil 12 \log(n_2)/\epsilon^2 \rceil, |S_Y| = \lceil 12 \log(n_1)/\epsilon^2 \rceil$. (Their paper states a bound that is slightly worse when n_1 and n_2 are far apart, but it is not hard to improve their analysis here.) Indeed, S_X, S_Y can be obtained by randomly sampling from p_1, p_2 respectively. Such S_X, S_Y can be found deterministically by checking (2.2) for all possible choices of S_X, S_Y . This requires time

$$O(n_1^{|S_X|} n_2^{|S_Y|}) \leq O(\exp(24 \log(n_1) \log(n_2)/\epsilon^2)),$$

which matches the performance of our algorithm up to the constant term in the $O()$.

In the multiplayer case, let us consider for simplicity the case of m players each with n strategies. Here LMM find that an ϵ -ANE exists with all probabilities integer multiples of $1/k$

with $k = 3m^2 \ln(m^2n)/\epsilon^2$. The resulting runtime is $O(n^{mk}) = \exp(O(m^3 \ln(n) \ln(mn)/\epsilon^2))$. Our run-time is essentially the same, but with the $\ln(mn)$ term replaced by a $\ln(n)$ term. (However, we note that when $m \gg n$ the algorithm of [4] achieves a better runtime of $\exp(O(m(\ln(n) + \ln(m))))$.)

Our algorithm can be used to achieve a slightly different and stronger notion of approximation than [21]. Specifically, it could be used to output an ϵ -correlated equilibrium. By Lemma 4.3, this can be used to obtain an ϵ -ANE, but the reverse direction is not known.

On the other hand, if a Nash equilibrium exists with small support, then LMM will find it exactly. It does not appear that our method would take advantage of the existence of small-support Nash equilibria. Our method would outperform its worst-case bounds under a somewhat different condition: if Alice's strategy had low entropy in all correlated equilibria. Fortunately, this can be checked quickly, since correlated equilibria form a polytope and entropy is a concave function that we can maximize efficiently using standard techniques. Unfortunately, this condition does not seem to be a very natural one. As mentioned in the introduction, both methods are compatible with maximizing linear objective functions (LMM works because a Chernoff bound can also be used to show the value of the objective function can be approximated by sparse solutions), but only our method works for maximizing general concave functions, such as entropy. Entropy maximization has been discussed before in the context of games [14], but we are not aware of algorithmic implications of this.

Our algorithm also has the disadvantage (compared with LMM) of requiring quasipolynomial space, whereas LMM requires only polynomial space. On the other hand, it is possible that our LP could be approximately solved using the multiplicative weights method to reduce this space requirement.

Acknowledgments. We are grateful for helpful discussions with Costis Daskalakis, Matt Hastings and Ben Recht. AWH was funded by NSF grants CCF-1111382 and CCF-1452616. AWH and AN were funded by ARO contract W911NF-12-1-0486.

References

- 1 S. Aaronson, R. Impagliazzo, and D. Moshkovitz. AM with multiple merlins. In *Computational Complexity (CCC), 2014 IEEE 29th Conference on*, pages 44–55, June 2014. doi:10.1109/CCC.2014.13.
- 2 Noga Alon, Troy Lee, Adi Shraibman, and Santosh Vempala. The approximate rank of a matrix and its algorithmic applications: Approximate rank. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing, STOC'13*, pages 675–684, 2013. doi:10.1145/2488608.2488694.
- 3 Robert Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974. URL: <http://EconPapers.repec.org/RePEc:eee:mateco:v:1:y:1974:i:1:p:67-96>.
- 4 Yakov Babichenko, Siddharth Barman, and Ron Peretz. Simple approximate equilibria in large games. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation, EC'14*, pages 753–770, New York, NY, USA, 2014. ACM. doi:10.1145/2600057.2602873.
- 5 Boaz Barak. Sum of squares upper bounds, lower bounds, and open questions. <http://www.boazbarak.org/sos/files/all-notes.pdf>, 2014.
- 6 Siddharth Barman. Approximating Nash equilibria and dense bipartite subgraphs via an approximate version of Caratheodory's theorem. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing, STOC'15*, pages 361–369, New York, NY, USA, 2015. ACM. doi:10.1145/2746539.2746566.

- 7 Fernando G. S. L. Brandão and Aram W. Harrow. Quantum de Finetti theorems under local measurements with applications. In *Proceedings of the 45th annual ACM Symposium on theory of computing*, STOC'13, pages 861–870, 2013.
- 8 Fernando G. S. L. Brandão and Aram W. Harrow. Estimating operator norms using covering nets, 2015.
- 9 Mark Braverman, Young Kun Ko, and Omri Weinstein. Approximating the best Nash equilibrium in $n^{o(\log n)}$ -time breaks the Exponential Time Hypothesis, 2014. ECCC TR14-092.
- 10 Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player Nash equilibria. *J. ACM*, 56(3):14:1–14:57, May 2009. doi:10.1145/1516512.1516516.
- 11 T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Series in Telecommunication. John Wiley and Sons, New York, 1991.
- 12 Dima Grigoriev. Linear lower bound on degrees of Positivstellensatz calculus proofs for the parity. *Theoretical Computer Science*, 259:613–622, 2001.
- 13 M. Grötschel, L. Lovász, and A. Schrijver. *Geometric algorithms and combinatorial optimization*. Springer-Verlag, 1993.
- 14 Peter D. Grünwald and A. Philip Dawid. Game theory, maximum entropy, minimum discrepancy and robust Bayesian decision theory. *Ann. Statist.*, 32(4):1367–1433, 08 2004. doi:10.1214/009053604000000553.
- 15 Aram W Harrow, Anand Natarajan, and Xiaodi Wu. Limitations of monogamy, Tsirelson-type bounds, and other semidefinite programs in quantum information, 2015.
- 16 Elad Hazan and Robert Krauthgamer. How hard is it to approximate the best nash equilibrium? *SIAM Journal on Computing*, 40(1):79–91, 2011. doi:10.1137/090766991.
- 17 Sam Hopkins, Pravesh Kothari, Aaron Potechin, Prasad Raghavendra, and Tselil Schramm. On the integrality gap of degree-4 sum of squares for planted clique. In *SODA '16*, 2016.
- 18 Jean B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11(3):796–817, 2001.
- 19 M. Laurent. Sums of squares, moment matrices and optimization over polynomials. *Emerging applications of algebraic geometry*, 149:157–270, 2009.
- 20 James R. Lee, Prasad Raghavendra, and David Steurer. Lower bounds on the size of semidefinite programming relaxations, 2014.
- 21 Richard J. Lipton, Evangelos Markakis, and Aranyak Mehta. Playing large games using simple strategies. In *EC'03*, pages 36–41, 2003. doi:10.1145/779928.779933.
- 22 Francesco M. Malvestuto. Theory of random observables in relational data bases. *Information Systems*, 8(4):281–289, 1983.
- 23 William J. McGill. Multivariate information transmission. *Psychometrika*, 19(2):97–116, 1954.
- 24 John F. Nash. Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950. doi:10.1073/pnas.36.1.48.
- 25 John F. Nash. Non-cooperative games. *Annals of mathematics*, pages 286–295, 1951.
- 26 Y. Nesterov. Squared functional systems and optimization problems. *High performance optimization*, 13:405–440, 2000.
- 27 Ryan O'Donnell, John Wright, Chenggang Wu, and Yuan Zhou. Hardness of robust graph isomorphism, lasserre gaps, and asymmetry of random graphs. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA'14, pages 1659–1677. SIAM, 2014. doi:10.1137/1.9781611973402.120.
- 28 Christos H. Papadimitriou and Tim Roughgarden. Computing correlated equilibria in multi-player games. *J. ACM*, 55(3):14:1–14:29, August 2008. doi:10.1145/1379759.1379762.

- 29 Christos H. Papadimitriou and Mihalis Yannakakis. Optimization, approximation, and complexity classes. *Journal of Computer and System Sciences*, 43:425–440, 1991.
- 30 Pablo A. Parrilo. Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization. Technical report, MIT, 2000. Ph.D thesis.
- 31 Prasad Raghavendra and Ning Tan. Approximating CSPs with global cardinality constraints using SDP hierarchies. In *SODA'12*, pages 373–387, 2012.
- 32 Grant Schoenebeck. Linear level Lasserre lower bounds for certain k -CSPs. In *Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science, FOCS'08*, pages 593–602, Washington, DC, USA, 2008. IEEE Computer Society. doi:10.1109/FOCS.2008.74.
- 33 N.Z. Shor. An approach to obtaining global extremums in polynomial mathematical programming problems. *Cybernetics and Systems Analysis*, 23(5):695–700, 1987.
- 34 Madhur Tulsiani. Csp gaps and reductions in the lasserre hierarchy. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 303–312. ACM, 2009.
- 35 Satoshi Watanabe. Information theoretical analysis of multivariate correlation. *IBM Journal of research and development*, 4(1):66–82, 1960.

A Polynomial Optimization and Sum-of-Squares Proofs

In this section, we lay out the basics of the sum-of-squares (SoS) optimization algorithms. They were introduced in [33, 26, 30, 18] and reviewed in [19, 5].

A.1 Polynomials

Let $\mathbb{R}[x] := \mathbb{R}[x_1, \dots, x_n]$ be the set of real-valued polynomials over n variables, and let $\mathbb{R}[x]_d$ be the subspace of polynomials of degree $\leq d$. The set of polynomials $\mathbb{R}[x]_d$ can be viewed as $\bigoplus_{d' \leq d} \text{Sym}^{d'} \mathbb{R}^n$, where $\text{Sym}^{d'} V$ denotes the symmetric subspace of $V^{\otimes d'}$.

A.2 Polynomial optimization

Given polynomials $f, g_1, \dots, g_m \in \mathbb{R}[x]$, the basic polynomial optimization problem is to find

$$f_{\max} := \sup_{x \in \mathbb{R}^n} f(x) \text{ subject to } g_1(x) = \dots = g_m(x) = 0. \quad (\text{A.1})$$

Equivalently we could impose inequality constraints of the form $g'_i(x) \geq 0$ but we will not explore this option here.

A.3 Sum-of-Squares (SoS) proofs

Although (A.1) is in general NP-hard to compute exactly, the SoS hierarchy is a general method for approximating f_{\max} from above. This complements simply guessing values of x or (ρ, X) which provides lower bounds on f_{\max} when they satisfy the constraints. A SoS proof is a bound that makes use of the fact that $p(x)^2 \geq 0$ for any $p \in \mathbb{R}[x]$. In particular, a SoS proof that $f(x) \leq c$ for all valid f is a collection of polynomials $p_1, \dots, p_k, q_1, \dots, q_m \in \mathbb{R}[x]$ such that

$$c - f = \sum_{i=1}^k p_i^2 + \sum_{i=1}^m q_i g_i. \quad (\text{A.2})$$

Observe that the RHS is ≥ 0 when evaluated on any x satisfying $g_i(x) = 0, \forall i$; for this reason, we refer to (A.2) as a Sum-of-Squares (SoS) proof, in particular, a proof that $c - f(x) \geq 0$ whenever $g_i(x) = 0$ for all i . This is a degree- d SoS proof if each term p_i^2 and $q_i g_i$ is in $\mathbb{R}[x]_d$. Finding an SoS proof of degree $\leq d$ can be done in time $n^{O(d)} m^{O(1)}$ using semidefinite programming [19].

If we find the minimum c for which (A.2) holds, then we obtain a hierarchy of upper bounds on f_{\max} , referred to as the SoS hierarchy or the Lasserre hierarchy. Denote this upper bound by f_{SoS}^d . Given mild assumptions on the constraints g_1, \dots, g_m one can prove that $\lim_{d \rightarrow \infty} f_{\text{SoS}}^d = f_{\max}$ [19]. The tradeoff between degree d and error ($f_{\text{SoS}}^d - f_{\max}$) is the key question about the SoS hierarchy. We can also express this tradeoff by defining $\text{deg}_{\text{SoS}}(c - f)$ to be the minimum d for which we can find a solution to (A.2). Note that deg_{SoS} has an implicit dependence on the g_1, \dots, g_m .

A.4 Pseudo-expectations

We will work primarily with a dual version of SoS proofs that have an appealing probabilistic interpretation. A degree- d pseudo-expectation $\tilde{\mathbb{E}}$ is an element of $\mathbb{R}[x]_d^*$ (i.e. a linear map from $\mathbb{R}[x]_d$ to \mathbb{R}) satisfying

- **Normalization.** $\tilde{\mathbb{E}}[1] = 1$.
- **Positivity.** $\tilde{\mathbb{E}}[p^2] \geq 0$ for any $p \in \mathbb{R}[x]_{d/2}$.

We further say that $\tilde{\mathbb{E}}$ satisfies the constraints g_1, \dots, g_m if $\tilde{\mathbb{E}}[g_i q] = 0$ for all $i \in [m]$ and all $q \in \mathbb{R}[x]_{d - \text{deg}(g_i)}$. Then SDP duality³ implies that

$$f_{\text{SoS}}^d = \max\{\tilde{\mathbb{E}}[f] : \tilde{\mathbb{E}} \text{ is a degree-}d \text{ pseudo-expectation satisfying } g_1, \dots, g_m\}. \quad (\text{A.3})$$

The term ‘‘pseudo-expectation’’ comes from the fact that for any distribution μ over R^n we can define a pseudo-expectation $\tilde{\mathbb{E}}[f] := \mathbb{E}_{x \sim \mu}[f(x)]$. Thus the set of pseudo-expectations can be thought of as the low-order moments that could come from a ‘‘true’’ distribution μ or could come from a ‘‘fake’’ distribution. Indeed an alternate approach (which we will not use) proceeds from defining ‘‘pseudo-distributions’’ that violate the nonnegativity condition of probability distributions but in a way that cannot be detected by looking at the expectation of polynomials of degree $\leq d$ [20].

A.5 The boolean cube

Throughout this work, we will be interested in the special case of pseudo-expectations over the boolean cube $\{\pm 1\}^n$. This set is defined by the constraints $x_i^2 - 1 = 0, i = 1, \dots, n$, and thus we say that $\tilde{\mathbb{E}}$ is a degree- d pseudo-expectation over $\{\pm 1\}^n$ if for any variable x_i and polynomial q of degree at most $d - 2$,

$$\tilde{\mathbb{E}}[(x_i^2 - 1)q] = 0. \quad (\text{A.4})$$

This means we can define $\tilde{\mathbb{E}}$ entirely in terms of its action on multilinear polynomials.

³ Certain regularity conditions (e.g. the Archimedean condition) are needed for strong duality to hold; for more details, see section 6.2 of [19]. These conditions hold in particular when the feasible set is a subset of the Boolean hypercube, which is the only setting we use in this paper.

B Framework of Deriving Lower Bounds

In this section, for the sake of completeness, we demonstrate our framework of deriving sum-of-squares (SoS) or semidefinite programming (SDP) lower bounds for optimization problems formulated in [15]. To this end, we formalize the familiar notions of optimization problem, SDP relaxations and integrality gaps. Then we show general methods for reducing optimization problems to each other as well as mapping integrality gaps for one problem/relaxation pair to another.

B.1 Optimization problems and integrality gaps

To prove a hardness result for an optimization problem, we would like to find instances where the SoS hierarchy and other SDP relaxations fail. These examples are known as “integrality gaps,” where the terminology comes from the idea of approximating integer programs with convex relaxations. For our purposes, an integrality gap will be an example of an optimization (maximization) problem in which the true answer is lower than the output of the SDP relaxation. To achieve this, we need to demonstrate a feasible point of the SDP with a value that is larger than the true answer. These feasible points are called *pseudo-solutions*, and we will define them for any polynomial optimization problem as follows.

► **Definition B.1** (Pseudo-Solution). Let A be a polynomial optimization problem. Let $\Phi_m^A \in \Delta_m^A$ be an instance of optimization A for some m . A *degree- d value- c pseudo-solution* for Φ_m^A is a degree- d pseudo-expectation $\tilde{\mathbb{E}}$ satisfying the constraints of \mathcal{P}_n^A such that

$$\tilde{\mathbb{E}}[\Phi_m^A(x)] \geq c$$

A single degree- d value- c pseudo-solution for an instance Φ_m^A implies the sum-of-squares approach (up to degree d) believes the optimum value of Φ_m^A is at least c . If the true optimum value of Φ_m^A is smaller than c , then such a pseudo-solution serves as an integrality gap for the SoS approach, i.e. an example where the SoS hierarchy gives the wrong answer. To refute the power of the SoS hierarchy, we need to establish such pseudo-solutions as well as small true optimum values for any large m .

► **Definition B.2** (integrality gap). Let A be any polynomial optimization problem. Let $d = d(n), c = c(n), s = s(n)$ be functions of n such that $0 \leq s < c \leq 1$. A degree- d value- (c, s) integrality gap for A is a collection of $\Phi_n^A \in \Delta_n^A$ for each $n \geq n_0$, s.t.

- The true optimum value $\text{OPT}(\Phi_n^A) \leq s$.
- For each $n \geq n_0$, there exists a degree- d value- c pseudo-expectation $\tilde{\mathbb{E}}_n$ for Φ_n^A such that $\tilde{\mathbb{E}}_n[\Phi_n^A(x)] \geq c$.

B.2 Reduction between optimization problems

To obtain SoS lower bounds for optimization problems, it suffices to establish integrality gaps. However, it is not clear how to obtain such integrality gaps in general, which might be a challenging task by its own. Here, we formulate an approach to establish such integrality gaps through reductions. Specifically, we start with some optimization problem with known integrality gaps and reduce it to an optimization problem that we want to establish integrality gaps.

► **Definition B.3** (Reductions). A *reduction* $R_{A \Rightarrow B}$ from optimization problem A to optimization problem B is a map from Δ^A to Δ^B ; i.e. $R(\Phi_n^A) \in \Delta_n^B$. It is called

- (s^B, s^A) -approximate if for any n and any Φ_n^A and its corresponding $\Phi_n^B = R(\Phi_n^A)$, we have

$$\text{OPT}(\Phi_n^A) = \max_{x \in \mathcal{P}_n^A} \Phi_n^A(x) \leq s^A \Rightarrow \text{OPT}(\Phi_n^B) = \max_{x \in \mathcal{P}_n^B} \Phi_n^B(x) \leq s^B.$$

Here s^A, s^B are understood to be functions of n .
The parameters (s^B, s^A) can be considered soundness parameters of the reduction.

B.3 3XOR with integrality gap

In this section, we will introduce the source of all hardness we have for this paper, which is the 3XOR problem first discovered by Grigoriev [12] and subsequently rediscovered by Schoenebeck [32]. It is analogous to the proof that 3-SAT is NP-hard, from which other hardness results can be derived by reducing those problems to 3-SAT. In our framework, 3XOR can be formulated as follows.

► **Definition B.4 (3XOR).** 3XOR is a boolean polynomial optimization problem with the following restriction:

- **Instances:** for any n , an instance is parameterized by a formula Φ_n that consists of $m = m(n)$ 3XOR clauses, the set of which denoted by \mathcal{C} , on n boolean variables (i.e., each clause is $x_i x_j x_k = a_{ijk}$ for some combination of (i, j, k) and $x_i, x_j, x_k \in \{\pm 1\}$). Thus, the objective function is

$$\Phi_n(x) = \frac{1}{m} \sum_{(i,j,k) \in \mathcal{C}} \frac{1 + a_{ijk} x_i x_j x_k}{2}, x \in \{\pm 1\}^n.$$

Thanks to the $x_i^2 = 1$ constraints, these terms are equivalent to ones of the form $(1 - (x_i x_j x_k - a_{ijk})^2)/2$.

Grigoriev's result [12] (reformulated by Barak [5]) implies the following integrality gaps. (Note that we have a slightly different formulation from [5] that is slightly stronger but guaranteed by [12].)

► **Proposition B.5** (Theorem 3.1 of [5], due to Grigoriev). *For any $\epsilon > 0$, for every n there exists a 3XOR instance Φ_n with n variables and $m = O(n/\epsilon^2)$ clauses, such that $\text{OPT}(\Phi_n) \leq \frac{1}{2} + \epsilon$, but there exists a degree- $\Omega(n)$ value-1 pseudo-solution $\tilde{\mathbb{E}}$.*

Here "value 1" means that for every clause $x_i x_j x_k = a_{ijk}$, it holds that $\tilde{\mathbb{E}}[(x_i x_j x_k - a_{ijk})p(x)] = 0$ for all polynomials $p(x)$ with degree at most $d - 3$.

In our framework, this implies a degree- $\Omega(n)$ value- $(1, \frac{1}{2} + \epsilon)$ integrality gap for the 3XOR problem.

C Information Theory

Proof of the claims in this section can be found in any information-theory textbook, such as [11]. For a distribution $p \in \Delta_n$ define its *entropy* to be

$$H(p) = - \sum_{x \in [n]} p(x) \ln(p(x)).$$

It can be shown that $0 \leq H(p) \leq \ln(n)$.

Another basic quantity is the relative entropy, which is defined for a pair of distributions $p, q \in \Delta_n$ to be

$$D(p\|q) := \sum_{x \in [n]} p(x) \ln \left(\frac{p(x)}{q(x)} \right).$$

The relative entropy has some distance-like properties. For our purposes, we will need only *Pinsker's inequality* [11, Lemma 11.6.1]:

$$D(p\|q) \geq \frac{1}{2} \|p - q\|_1^2. \quad (\text{C.1})$$

For a distribution over several random variables e.g. $p^{XYZ} \in \Delta_{n^3}$ where X, Y, Z are supported on $[n]$, we write the entropy of the marginals using a notation that emphasizes the variables rather than the distribution:

$$H(X)_p := H(p^X), H(XY)_p := H(p^{XY}), H(XYZ)_p := H(p), \text{etc..}$$

Using this notation we can define the mutual information between random variables X, Y to be

$$I(X; Y)_p := H(X)_p + H(Y)_p - H(XY)_p.$$

It is straightforward to show that $I(X; Y) \leq \min(H(X), H(Y))$. The mutual information is a measure of correlation, as can be seen by the following alternate characterization:

$$I(X; Y)_p = D(p^{XY} \| p^X \otimes p^Y). \quad (\text{C.2})$$

This can be verified by a quick calculation. Combined with (C.1), we obtain

$$\|p^{XY} - p^X \otimes p^Y\|_1 \leq \sqrt{2I(X; Y)_p}. \quad (\text{C.3})$$

Finally we will make use of the *conditional mutual information*, defined to be

$$I(X; Y|Z)_p := I(X; YZ)_p - I(X; Z)_p. \quad (\text{C.4})$$

The term conditional mutual information refers to the alternate characterization of the quantity $I(X; Y|Z)$ as the mutual information of conditional distribution XY averaged over all values of Z ; i.e.

$$I(X; Y|Z)_p = \sum_z p(Z = z) I(X; Y) |_{p_{Z=z}}, \quad (\text{C.5})$$

where $p_{Z=z}(x, y) := p(x, y, z) / \sum_{x', y'} p(x', y', z)$.

We will also find it useful to repeatedly apply (C.4) to obtain the *chain rule of mutual information*:

$$I(X; Y_1 \dots Y_k) = \sum_{j=1}^k I(X; Y_j | Y_1 \dots Y_{j-1}). \quad (\text{C.6})$$