Electronic Colloquium on Computational Complexity, Report No. 92 (2014)

Approximating the best Nash Equilibrium in $n^{o(\log n)}$-time breaks the Exponential Time Hypothesis

Mark Braverman* Young Kun Ko † Omri Weinstein‡

Abstract

The celebrated PPAD hardness result for finding an exact Nash equilibrium in a two-player game initiated a quest for finding *approximate* Nash equilibria efficiently, and is one of the major open questions in algorithmic game theory.

We study the computational complexity of finding an ε-approximate Nash equilibrium with good social welfare. Hazan and Krauthgamer and subsequent improvements showed that finding an ε-approximate Nash equilibrium with good social welfare in a two player game and many variants of this problem is at least as hard as finding a planted clique of size $O(\log n)$ in the random graph $\mathcal{G}(n, 1/2)$.

We show that any polynomial time algorithm that finds an ε-approximate Nash equilibrium with good social welfare refutes (the worst-case) Exponential Time Hypothesis by Impagliazzo and Paturi. Specifically, it would imply a $2^{\tilde{O}(n^{1/2})}$ algorithm for SAT.

Our lower bound matches the quasi-polynomial time algorithm by Lipton, Markakis and Mehta for solving the problem.

Our key tool is a reduction from the PCP machinery to finding Nash equilibrium via free games, the framework introduced in the recent work by Aaronson, Impagliazzo and Moshkovitz. Techniques developed in the process may be useful for replacing planted clique hardness with ETH-hardness in other applications.

*Department of Computer Science, Princeton University, mbraverm@cs.princeton.edu, Research supported in part by an NSF Award CCF-1215990, an NSF CAREER award (CCF-1149888), a Turing Centenary Fellowship, a Packard Fellowship in Science and Engineering, and the Simons Collaboration on Algorithms and Geometry.

†Department of Computer Science, Princeton University, yko@CS.Princeton.edu,

‡Department of Computer Science, Princeton University, oweinste@cs.princeton.edu. Research supported by a Simons Fellowship in Theoretical Computer Science and NSF Award CCF-1215990.

ISSN 1433-8092

1 Introduction

An important aspect of algorithmic game theory is understanding whether it is possible to efficiently compute or reach a stable state in various economic scenarios. The solution concept which has attracted the most attention, from the algorithms, complexity and machine-learning communities, is that of Nash equilibrium. The celebrated results of [CDT06, DGP06] assert that computing a Nash equilibrium in a finite game with two players is PPAD-complete. An attempt to circumvent this impossibility result has initiated a line of research focused on the more modest goal of finding an *approximate* Nash equilibrium:

Is there a polynomial time algorithm that finds an ε-Nash equilibrium for arbitrarily small but fixed $\varepsilon > 0$?

An ε-Nash equilibrium (or simply, ε-equilibrium) is an equilibrium in which no party has an incentive of more than ε to deviate from its current strategy. Besides being potentially more tractable (computationally), the family of ε-Nash equilibria is significantly less restrictive than that of exact Nash equilibria, and thus allows one to capture a richer set of desirable behaviors observed in practice. An example of this is a n-round repeated Prisoners' Dilemma — where the only Nash equilibrium is repeated defection, but some of the strategies observed in experiments, such as *tit-for-tat* and *grim-trigger* form ε-Nash equilibria for $\varepsilon \to 0$ as $n \to \infty$ [Axe87, Axe00].

While every finite two-player game is guaranteed to have at least one equilibrium, typical games possess many equilibria, some of which are more desirable than others. A standard measure of the "economic efficiency" of an equilibrium is its social welfare (i.e., the sum of payers' payoffs). Thus, a particularly appealing goal is that of finding an equilibrium or an ε-approximate equilibrium with high social welfare. This goal captures the most efficient outcome one can achieve assuming the players are ε-sensitive. More generally, being able to solve this problem appears to be an important step in understanding the space of ε-Nash equilibria of a particular game.

Finding an (exact) equilibrium with maximal social welfare is known to be NP-hard [GZ89, CS03], and is thus probably more difficult than finding just any equilibrium. A simple yet surprising result of Lipton, Markakis and Mehta [LMM03] asserts that it is possible to find an ε-equilibrium whose welfare is within an additive factor ε of the optimal one, in quasi-polynomial time $O(n^{\varepsilon^{-2}\log n})$, assuming the game is represented by an $n \times n$ matrix. The [LMM03] argument first uses *random sampling* to prove that any ε-equilibrium can be approximated by an $(\varepsilon + \varepsilon')$-equilibrium with only *logarithmic* support size. Once this fact has been established, one can use exhaustive search to find such an approximate equilibrium in time $n^{O_{\varepsilon'}(\log n)}$. Note that such an exhaustive search not only finds one approximate equilibrium, but constructs an ε'-net covering *all* ε-equilibria. In particular, this allows one to find an approximate equilibrium that is close to maximizing social welfare.

The existence of a quasi-polynomial algorithm sparked the hope that a polynomial algorithm for this problem exists. The subsequent works of [DMP06, KPS06, DMP07, BBM07] made additional encouraging progress in this direction, and the current state of the art is a polynomial time algorithm that computes a 0.3393-equilibrium [TS07].

It is natural to ask whether one can give matching (quasi-polynomial) lower bounds for the ε-Nash problem. The work [CDT06] rules out strongly-polynomial algorithms (under PPAD-hardness). However, a weaker PTAS (one that can have tems of the form $n^{1/\varepsilon}$) cannot be ruled out at the moment. For variants of the ε-Nash problem whose exact version is NP-hard, such as approximating the best-value equilibrium lower bounds conditional on the *planted clique* assumption previously existed. Hazan and Krauthgamer [HK09] showed that there is a constant $\varepsilon > 0$ such that finding an ε-equilibrium with welfare within ε of the highest-utility Nash equilibrium is *planted clique* hard. Specifically, such an algorithm cane be used (in a black-box fashion) to find a planted $5\log n$-clique in the random graph $\mathcal{G}(n, 1/2)$. Subsequently, Austrin et al. [ABC13] generalized the results of [HK09] to tighter values of ε and to a number of other NP-hard variants of the Nash equilibrium finding problem.

In the hidden clique problem [Jer92, Kuc95], the input is a graph G obtained by planting a clique of size $t > 5\log n$ in the random graph $\mathcal{G}(n, 1/2)$ (note that the typical max-clique number of $\mathcal{G}(n, 1/2)$ is only $\approx 2\log_2 n$). The objective is to find the clique, or, more modestly, to distinguish between the random graph $\mathcal{G}(n, 1/2)$ and the "planted" graph. When $t = \Theta(\sqrt{n})$, this can be done in polynomial time [AKS98] (subsequent algorithms with similar performance appear in [FK00, FR10, DGGP10]). However, for $t = o(\sqrt{n})$, there is no known polynomial time algorithm that finds even a $(1 + \epsilon)\log_2 n$ clique, for any constant $\epsilon > 0$. Note that this problem is easily solvable in quasi-polynomial time, since a "seed" of $3\log_2 n$ nodes in the clique, which can be found using brute force enumeration, allows one to identify the entire planted clique. The Planted Clique Conjecture asserts that, indeed, this problem requires $n^{\Omega(\log n)}$ time [AAK+07, AKS98, Jer92, Kuc95, FK00, DGGP10]. Recently, a fair number of diverse applications based on this conjecture, such as [AAK+07, BBB+13, BR13], have been published. The main weakness of this hypothesis is that it is an average-case assumption (where the instance is a random instance). It is often the case that average-case instances, even of worst-case-hard problems, are in fact tractable. Moreover, the average-case nature of the hypothesis makes it unlikely that one could prove its hardness based on classical worst-case assumptions such as lower bounds on SAT.

In this work we provide a very strong piece of evidence for ruling out a polynomial-time algorithm for finding an ε-equilibrium with an ε-additive approximation to the optimal welfare. We show that solving this this bi-criteria problem in $n^{o(\log n)}$ time, would yield a $2^{o(n)}$ algorithm for SAT, thereby refuting the well known Exponential Time Hypothesis:

Conjecture 1.1 (Exponential Time Hypothesis [IP01])**.** *Any deterministic algorithm for* **3SAT** *requires* $2^{\Omega(n)}$ *time.*

We refer the reader to [LMS+13] for a more thorough background on this conjecture, and its broader role within complexity theory.

The starting point of our reduction from SAT to Nash is a recent result, by Aaronson, Impagliazzo and Moshkovitz [AIM14] that connects classical PCP machinery to free two-prover games. In their work, they first take a view of the PCP theorem as stating that approximating the value of so called "two-prover" games (associated with the underlying SAT instance) is NP-hard. Informally, a two-prover game is a cooperative constraint satisfaction game where two provers Alice and Bob are given a pair of challenges (x, y) and output responses a and b, respectively. Alice and Bob win if the verifier accepts their responses — that is if $V(x, y, a, b) = 1$, where $V(\cdot)$ is a predicate known in advance. The value of a game is the highest probability of acceptance Alice and Bob can achieve under a given distribution of challenges. Two-prover games have played an important role in the theory of PCP and hardness of approximation.

Aaronson et al. explore a special type of two-prover games, called *free games*, in which the challenges of Alice and Bob are *independent* of each other (distributed uniformly). While the PCP theorem states that approximating the value of a general two prover game is NP-hard, free games turn out to be much easier to approximate. Indeed, [AIM14] and [BH13] independently give algorithms for approximating the value of a free game of size[1] N to within an additive ε factor in time $O(N^{\varepsilon^{-2}\log N})$, the same as the running rime of LMM's algorithm for finding a (high-welfare) ε-approximate Nash equilibrium! This match is not a coincidence — the technique used by

[1]See Section 2.2 for the precise definitions.

[AIM14] for proving their upper bound highly resembles that of [LMM03]: reducing the support-size of the player's strategy space by random sampling, and then exhaustively searching for the best logarithmic-support strategies.

The main result of [AIM14] is a reduction which converts general games into free games, with a sub-exponential blowup in the size of the game. A general game with challenge-sets of size n can be converted into an (essentially) equivalent free game of size $N := 2^{\tilde{O}(\sqrt{n})}$. Combined with the above result, this shows that it is possible to approximate the value of a *general* two-prover game to within arbitrary additive precision in time $N^{O(\log(N))} = 2^{\tilde{O}(n)}$. Notice that the ETH conjecture asserts this is tight — any asymptotic improvement on the quasi-polynomial running time for approximating the value of free games wold imply a $2^{o(n)}$ algorithm for approximating the original constraint-satisfaction game, and therefore for SAT (by PCP). We elaborate more on this entire line of results in Section 2.2.

We should note that it is not hard to construct a (randomized) reduction from the Planted Clique problem to the free game vaule approximation problem. Thus, an $N^{o(\log N)}$ algorithm for the latter would (in addition to breaking the ETH), also break the planted clique hypothesis, providing further evidence for the superiority of relying on free-game value hardness over planted-clique hardness.

1.1 Our results

Our main technical result is the following.

Theorem 1.2 (Main reduction)**.** *There exist global constants $1 > \varepsilon', \varepsilon^* > 0$ such that the following holds. There is a reduction running in time $2^{\tilde{O}(\sqrt{n})}$ which takes as an input an n-variable* **3SAT** *formula ϕ, and outputs a two player strategic game $\mathcal{G}$, such that:*

- *If ϕ is satisfiable, then $\mathcal{G}$ has an equilibrium of value 1.*
- *If ϕ is not satisfiable, then any ε^*-equilibrium of $\mathcal{G}$ has value at most ε'.*

The above theorem directly implies

Corollary 1.3. *Assuming the (deterministic)* **ETH** *conjecture (Conjecture 1.1), there is a constant $\varepsilon^* > 0$ such any algorithm for finding an ε^*-approximate Nash equilibrium whose welfare is lower by at most ε^* than the optimal welfare of a Nash equilibrium of the game, requires $n^{\tilde{\Omega}(\log n)}$ time.*

To best of our knowledge, this is the first computational relationship established between two-prover games (as in 'Unique Games Conjecture') and strategic games (as in 'Game Theory').

2 Preliminaries

Throughout the paper we use capital letters to denote sets and calligraphic letters to denote set families. We sometimes use calligraphic letters for probability distributions as well (For example, $x \sim \mathcal{D}$). In particular, we denote by $\mathcal{U}(S)$ the *uniform* distribution over S (or simply $\mathcal{U}$, when the support is clear from context). We use $\|\cdot\|_1$ as a conventional ℓ_1 norm, that is the sum of the absolute values of the entries. The rest of notation is set forth in the following two introductory subsections.

2.1 Approximate Nash equilibria in games

We restrict our attention to symmetric games, hence our definition assumes square matrices for the payoff. A (square) two player bi-matrix game is defined by two payoff matrices $R, C \in \mathbb{R}^{n\times n}$, such that if the row and column players choose pure strategies $i \in \mathcal{S} := [n],\ j \in \mathcal{T} := [n]$, respectively, the payoff to the row and column players are $R(i,j)$ and $C(i,j)$, respectively. The game is called *symmetric* if $C = R^\top$. A *mixed strategy* for a player is a distribution over pure strategies (i.e. rows/columns), and for brevity we may refer to it simply as a strategy. An *ε-approximate Nash equilibrium* (or simply, *ε-equilibrium*) is a pair of mixed strategies (x, y) such that:

$$\forall \tilde{x} \in \Delta(\mathcal{S})\ ,\ \tilde{x}^\top R y \le x^\top R y + \varepsilon, \qquad \text{and} \qquad \forall \tilde{y} \in \Delta(\mathcal{T})\ ,\ x^\top C \tilde{y} \le x^\top C y + \varepsilon,$$

where here and throughout the paper, $\Delta(S)$ denotes the family of all probability distributions over the set S. If $\varepsilon = 0$, the strategy pair (x, y) is called a *Nash equilibrium* (NE).

For a pair of (mixed) strategies (x, y), the (expected) *payoff* of the row player is $x^\top R y$ and similarly $x^\top C y$ for the column player. The *value* of an (approximate) equilibrium (μ, ν) for the game $\mathcal{G}$, denoted $\pi_{x,y}(\mathcal{G})$, is the average (expected) payoff of the players when the strategy pair (x, y) is played. Notice that is equivalent to the *social-welfare* (the sum of players payoffs) of the equilibrium up to a factor of 2. The *value of the game* $\mathcal{G}$, denoted $\pi(\mathcal{G})$, is the maximum, over all equilibrium strategy pairs (x, y), of the average payoff of the two players. $\mathbf{Nash}_\varepsilon$ is the problem of approximating the value of a two player game to within an additive constant ε:

Definition 2.1 ($\mathbf{Nash}_\varepsilon$)**.** *Given as input a game $\mathcal{G}$ described by two bounded*[2] *payoff matrices R, C, estimate $\pi(\mathcal{G})$ to within an additive factor $\pm\varepsilon$ (the input size of the problem is $n := |R||C|$ and, ε is an arbitrarily small constant, unless otherwise stated).*

Note that we give a lower bound for an easier problem (thus making the lower bound stronger): in a bounded-payoff game, distinguish the case when there is a welfare-1 Nash equilibrium from the case when *all* ε-Nash equilibria have welfare $< 1/2$.

In [LMM03], the authors exhibited an algorithm for $\mathbf{Nash}_\varepsilon$ that runs in $n^{O((\log n)/\varepsilon^2)}$ time and approximates the value of $\mathcal{G}$ within an additive error of ε. The idea behind their algorithm is to show the existence of a small-support equilibrium, by randomly subsampling $O((\log n)/\varepsilon^2)$ pure strategies for each player. Then they argue, via standard concentration bounds and the probabilistic method, that the uniform distribution over these randomly chosen strategies is an ε-equilibrium with high probability, and furthermore this set contains (w.h.p) an equilibrium within ε of $\pi(\mathcal{G})$. Finding such an ε-equilibrium can in turn be done by an exhaustive search over all subsets of the reduced support, yielding the claimed running time of $n^{O(\log n/\varepsilon^2)}$.

The hardness result of [HK09] implies that if one believes the "Planted Clique Conjecture", i.e, that finding a planted (but hidden) clique of size of $O(\log n)$ in the random graph $\mathcal{G}(n, 1/2)$ takes $n^{\Omega(\log n)}$ time, then the above algorithm is tight. Our result can thus be cast as reproving this tight lower bound, under the **ETH** conjecture instead of the Planted Clique conjecture.

2.2 Two-prover games and free games

In an effort to make this writeup as self-contained as possible, we henceforth introduce the necessary background and previous results leading to our main reduction. We begin with the following definition which is central to this paper:

[2] Since we are concerned with an additive notion of approximation, we assume that the entries of the matrices are in the range $[0, M]$, for M which is a constant independent of all the other parameters (cf. [HK09]).

Definition 2.2 (Two-Prover Game)**.** *A two-prover game $G = (X, Y, A, B, V)$ is a game between two provers (Alice and Bob) and a referee (Charlie), consiststing of:*

1. *Finite challenge sets X, Y (one for Alice, one for Bob) and answer sets A, B,*
2. *A probability distribution $\mathcal{D}$ over challenge pairs $(x, y) \in X \times Y$, and*
3. *A verification predicate $V : X \times Y \times A \times B \longrightarrow [0, 1]$.*

The value of the game, denoted $\omega(G)$, is

$$\max_{a:X\to A\ ,\ b:Y\to B} \quad \mathbb{E}_{(x,y)\sim\mathcal{D}}[V(x, y, a(x), b(y))],$$

where the maximum ranges over all pairs of response strategies of Alice and Bob.

A two-prover game can be interpreted as follows: The cooperating provers, Alice and Bob, can agree on a strategy ($a(X)$ and $b(Y)$) in advance but cannot communicate once the game starts. First Charlie chooses a pair of questions $(x, y) \sim \mathcal{D}$, and sends x to Alice and y to Bob. The provers then send back responses $a = a(x)$ and $b = b(y)$ respectively. Finally, Charlie declares the provers to have won with probability equal to $V(x, y, a, b)$. $\omega(G)$ is therefore the probability that the provers win if they use an optimal strategy.

A two-prover game is called *free* if the challenges (x, y) are chosen independently, that is, if $\mathcal{D}$ is a *product* distribution over $X \times Y$ (for our purposes, we assume without loss of generality that this is the *uniform* distribution on challenges). $\textbf{FreeGame}_\varepsilon$ is the problem of approximating the value of a free game up to an additive constant ε:

Definition 2.3 ($\textbf{FreeGame}_\varepsilon$)**.** *Given as input a description of a free game $G = (X, Y, A, B, V)$, estimate $\omega(G)$ to within an additive factor $\pm\varepsilon$ (unless otherwise stated, the input size is $n := |X||Y||A||B|$ and, ε is an arbitrarily small constant).*

Computing the exact or even the approximate value of a general two prover game is NP-hard — in fact, the celebrated PCP theorem can be restated as a hardness of approximation result for computing the value of two-prover games, namely that approximating the value of such game to within an additive *constant* error ε is NP-hard. Free games, however, turn out to be much easier to approximate; [AIM14] and [BH13] independently gave algorithms for $\textbf{FreeGame}_\varepsilon$ that runs in time $n^{O(\log n/\varepsilon^2)}$. The work of [BH13] is based on an LP-relaxation algorithm while in [AIM14], the authors show that a similar-in-spirit approach to that of [LMM03], of subsampling $O(\log n/\varepsilon^2)$ random challenges from each X and Y and restricting the free game to this smaller support, will only change $\omega(G)$ by $\pm\varepsilon$.

The authors in [AIM14] prove the above algorithm is in fact tight, assuming the **ETH** conjecture. Since our reduction exploits *the structure of the hard instance* produced by the reduction of [AIM14] (i.e, it is not a "black-box" reduction from Free Games), we give a brief overview of their result. This reduction is the content of the next subsection.

2.3 An ETH-hardness result for $\textbf{FreeGame}_\varepsilon$ ([AIM14])

The **ETH** conjecture involves a **3SAT** instance. The following sequence of reductions described in [AIM14] will allow us to consider a free game instance instead, as the starting point of our main reduction. The first step in this sequence is the following result of Moshkovitz and Raz, who showed how to map a **3SAT** instance to a **2CSP** instance, with a very mild blowup in the instance size:

$$\textbf{3SAT} \longrightarrow \textbf{2CSP} \longrightarrow \text{Clause/Variable Game} \longrightarrow \textbf{FreeGame}_\varepsilon$$

$$\varphi \longrightarrow \phi \longrightarrow H_\phi \longrightarrow H^{\sqrt{n}\times\sqrt{n}}$$

$$n \dashrightarrow n^{1+o(1)} \dashrightarrow n^{1+o(1)} \dashrightarrow 2^{\tilde{O}(\sqrt{n})}$$

Figure 1: A schematic outline of [AIM14]'s reduction from **3SAT** to $\textbf{FreeGame}_\varepsilon$.

Theorem 2.4 (PCP Theorem, [MR08])**.** *Given a* **3SAT** *instance φ of size n as well as $\delta > 0$, it is possible in* $\text{poly}(n)$ *time to produce a* **2CSP** *instance ϕ, with $n^{1+o(1)}\text{poly}(1/\delta)$ variables and constraints, and over an alphabet $A = B$ of size $|A| \le 2^{\text{poly}(1/\delta)}$, such that*

- *(Completeness) If* $\text{OPT}(\varphi) = 1$ *then* $\text{OPT}(\phi) = 1$*.*
- *(Soundness) If* $\text{OPT}(\varphi) < 1$ *then* $\text{OPT}(\phi) < \delta$*.*
- *(Balance) The constraint graph of ϕ is bipartite, and every variable appears in exactly d constraints, for some $d = \text{poly}(1/\delta)$.*

The above **2CSP** instance can in turn be viewed as a natural two-prover game:

Definition 2.5 (Clause/Variable Game H_ϕ (for **2CSP**), [AIM14])**.** *Given a* **2CSP** *instance ϕ, given by Theorem 2.4, with X and Y as the set of left and right variables respectively, we define the following two-prover game as a respective* Clause/Variable Game H_ϕ*. Charlie chooses a random constraint C_k which has $i \in X$ and $j \in Y$ as its variable. Then Charlie sends i to Alice and j to Bob. Charlie accepts if and only if Alice and Bob gives a satisfying assignment to C_k, that is $C_k(a(i), b(j)) = 1$.*

Notice, however, that H_ϕ is not a free game, as the event of Alice receiving $i \in X$ as her challenge is highly dependent on the event of Bob receiving $j \in Y$ as his challenge (challenge i for Alice leaves only $deg(i)$ possible challenges for Bob). In [AIM14], the authors introduce a "hardness amplification" technique ("Birthday Repetition") which converts H_ϕ into a *free game* H, while (essentially) preserving its value.

The idea is that instead of choosing one challenge for Alice and one for Bob, Charlie chooses a *tuple* of $\Theta(\sqrt{n})$ of the n challenges uniformly at random, say S for Alice and T for Bob. Then Alice gives assignments to all the challenges in S, and sends them back to Charlie. Similarly, Bob with T. Charlie accepts if and only if all the edges between S and T are satisfied by the assignments given by Alice and Bob. If there is no edge between S and T, Charlie just accepts regardless of the assignment (and the players "win for free"). Intuitively, the birthday paradox ensures that by choosing tuples of size $\Theta(\sqrt{n})$, this event is very unlikely, and thus the value of the original game is preserved (this is also the source of the term "birthday repetition").

Making the game free, however, comes at a substantial cost: the size of the game blows up. New challenges are $O(\sqrt{n})$-sized sets of singleton challenges, therefore the answers to these challenges must be also be $O(\sqrt{n})$-long tuples. As a result we get the following blowup size:

$$|H| = (|X||A|)^{|S|}(|Y||B|)^{|T|} = n^{O(\sqrt{n})} = 2^{O(\sqrt{n}\log n)} = 2^{\tilde{O}(\sqrt{n})}.$$

Combining all the aforementioned steps (see Figure 1), one obtains a mapping from **3SAT** to a hard instance of $\textbf{FreeGame}_\varepsilon$, with the following particular structure:

Theorem 2.6 (essentially [AIM14]). *For some constant $\varepsilon > 0$, there exists a reduction running in time $2^{\tilde{O}(\sqrt{n})}$ that maps* **3SAT** *instances of size n to a* **FreeGame**$_\varepsilon$ *instance $H = (\mathcal{S}, \mathcal{T}, \mathcal{A}, \mathcal{B}, V)$ of the following form:*

- *(Challenge and Answer sets) The challenge sets $\mathcal{S}$ and $\mathcal{T}$ are all $\sqrt{n}/\varepsilon$-sized subsets (of the original challenge sets X and Y of H_ϕ defined above). Accordingly, the answer sets are $\mathcal{A} = A^{\sqrt{n}/\varepsilon}$ and $\mathcal{B} = B^{\sqrt{n}/\varepsilon}$.*
- *(Game Size) The size of H is thus $|\mathcal{S}||\mathcal{T}||\mathcal{A}||\mathcal{B}| = 2^{\tilde{O}(\sqrt{n}/\varepsilon)}$.*
- *(Symmetry) $X = Y$ and $A = B$. Therefore $\mathcal{S} = \mathcal{T}$ and $\mathcal{A} = \mathcal{B}$.*
- *(Degree concentration) For any challenge $S \in \mathcal{S}$, let $\mathcal{N}(S)$ denote the set of challenges $T \in \mathcal{T}$ such that there exists some constraint $C(x, y)$ in original ϕ such that $x \in S$ and $y \in T$. Then $|\mathcal{N}(S)| \le d|S|$ and :*
$$\Pr_{S \sim \mathcal{U}(\mathcal{S})}\left[|\mathcal{N}(S)| > \frac{9}{10}d|S|\right] \ge 1 - o(1). \tag{1}$$
  Furthermore, if S satisfies (1) *then:*
$$\Pr_{T}\left[|\mathcal{N}(S) \cap T| < \frac{d}{10\varepsilon^2}\right] \le c_\varepsilon = \mathrm{poly}(1/\varepsilon)e^{-\mathrm{poly}(1/\varepsilon)} < 0.1 \tag{2}$$

The first two properties are by definition of the reduction of [AIM14]. The last two properties (symmetry and concentration) are additional properties that will be used in our main reduction, and we provide a complete proof of them in Section A.1 of the appendix.

Convention. To avoid confusion, from now on whenever considering the free game instance H from Theorem 2.6 above, we shall refer to challenges $(S, T) \in \mathcal{S} \times \mathcal{T}$ in H as *challenge-tuples*. The challenges X and Y of the original game (H_ϕ) shall remain under the name *challenges*. This distinction will be needed in the proof of Theorem 1.2. Now, assuming Conjecture 1.1, which gives a natural lower bound for solving **3SAT**, Theorem 2.6 directly implies the following.

Corollary 2.7 (Hardness of Free Games). *Assuming the deterministic (randomized)* **ETH** *conjecture, any deterministic (randomized) algorithm for* **FreeGame**$_\varepsilon$ *(of the type described in Theorem 2.6) requires $n^{\tilde{\Omega}(\varepsilon^2 \log n)}$ time, for all $\varepsilon \ge 1/n$ bounded below some constant.*

3 From FreeGame$_\varepsilon$ to Nash$_\varepsilon$: The main reduction

In this section we prove Theorem 1.2. By the previous discussion, it suffices to produce a gap-preserving reduction from the **FreeGame**$_\varepsilon$ instance from Theorem 2.6 to **Nash**$_{\varepsilon^*}$ with $\varepsilon^* = 1 - \Omega(1)$. Indeed, we show that

Lemma 3.1 (Main reduction). *There is a reduction which runs in $2^{\tilde{O}(\sqrt{n})}$-time, which maps the Free Game H from Theorem 2.6 to a two-player strategic game $\mathcal{G}$, such that:*

- *(Completeness) If $\omega(H) = 1$, then $\mathcal{G}$ has an equilibrium of value 1.*
- *(Soundness) If $\omega(H) \le \varepsilon$, then any ε^*-equilibrium of $\mathcal{G}$ has value at most ε^* , for $\varepsilon^* = \Omega(1)$, $\varepsilon^* = 1 - \Omega(1)$.*

We begin with an intuitive overview of our reduction.

3.1 Overview of the reduction

Consider the free game instance H from Theorem 2.6. A natural idea is to convert this constraint satisfaction game into a cooperative strategic game where the payoff is the value of the constraint satisfaction game. The naïve way of doing so is to create a matrix whose columns are indexed by all $|\mathcal{A}|^{|\mathcal{S}|}$ possible (pure) strategy functions Alice may use (similarly, columns are indexed by all $|\mathcal{B}|^{|\mathcal{T}|}$ strategies Bob might use). For any pair of strategies (a,b) chosen by the players, we can then assign both players an (identical) payoff of $\mathbb{E}_{(S,T)\sim\mathcal{U}}[V(S,T,a(S),b(T))]$. Clearly, any equilibrium strategy played by the players in the resulting game induces a response strategy for the original free-game with the same payoff, and in particular any ε-approximate Nash equilibrium provides an ε-approximation to $\omega(H)$. The problem, of course, is that the size of this game is prohibitively large – $2^{\Omega(N)}$, while our reduction is only allowed to run in time (and space) comparable to the size of H which is $2^{\tilde{O}(\sqrt{n})}$. One may try to apply the more clever subsampling technique of Barak et. al. [BHHS09] to reduce the instance size. Unfortunately, even this technique can at best reduce the size to $2^{\Omega(\log^2(N))} = 2^{\Omega(n)}$, while our reduction must run in $2^{o(n)}$ time in order to refute **ETH**.

The first step of our reduction is therefore to use a much more efficient encoding of the players' response strategies in the payoff matrix – the strategy space for Alice consists of all pairs (S,a) of her (challenge-tuple, response) in H, and, similarly, the strategy space for Bob are pairs (T,b). The payoff is $(1,1)$ if $V(S,T,a(S),b(T)) = 1$ and $(0,0)$ otherwise. Notice that the above encoding yields a game of size only $|\mathcal{S}||\mathcal{A}| \cdot |\mathcal{T}||\mathcal{B}| = 2^{\tilde{O}(\sqrt{n})}$, as desired.

If $\omega(H) = 1$, Alice and Bob can always achieve an equilibrium with welfare equal to 1: Alice will choose S uniformly, and will choose $a = a(S)$ according to her strategy in the constraint satisfaction game. Bob will choose (T,b) in a similar fashion, and neither player has an incentive to defect. This ensures that the value of the cooperative game is *at least* the value of the free game[3]. For the reduction to work, however, it also needs to be sound: if the value of H is low, the strategic game must *not* have a high-welfare equilibrium. The naïve construction above clearly fails at this task: Alice and Bob can ensure the maximum payoff of $(1,1)$ by choosing any tuple (S_0,T_0,a_0,b_0) satisfying $V(\cdot)$ and playing the pure strategy pair $((S_0,a_0),(T_0,b_0))$.

To address this problem, we would like to modify the game above in a way that forces players to play a (near) *uniform* strategy on the respective challenge sets, in which case finding a high-welfare Nash equilibrium entails approximating the optimal response strategy in H. This is also the main technical challenge of the proof. As in [HK09], we would like to add a negative-sum part to the game which allows e.g. Bob to "punish" Alice if she plays far from uniformly on her challenges. Unfortunately, following the [HK09] directly would cause an exponential blow-up in the size of the game (we can't afford even a quasi-polynomial blow-up). The planted clique reduction of [HK09] only needed to rule out distributions of support $O(\log n)$, while we need to enforce statistical closeness to the uniform distribution (although we manage to weaken this requirement). Luckily, H is not arbitrary, but has a very particular structure: Every challenge-tuple $S \in \mathcal{S}$ ($T \in \mathcal{T}$) induced by Alice's (Bob's) choice is a *subset* of $\sqrt{n}/\varepsilon$ original challenges from the ground set X (Y).

In order to implement the aforementioned idea, we will append additional payoff matrices to the game which allows Bob (reps. Alice) to specify a (distribution over) subset Z of size $\rho\sqrt{n}$ from the ground set X, and we will reward him with a positive payoff in the event that $Z \cap S \neq \emptyset$ (where S is the tuple chosen by Alice). Alice, on the other hand, will get penalized by a negative payoff if this event occurs (we make sure this part of the game is zero-sum, hence any high-welfare equilibrium should have a significant mass on the "original" part of the game). We set the size of the "penalzing"

[3]We remark that since we are only considering uncoordinated, simultaneous-move strategic games, the *joint* distribution on $\mathcal{S} \times \mathcal{T}$ obtained by players' mixed strategies, is always a *product* distribution ($p(S,T) = \mu(S)\nu(T)$), the same as in free games. Thus, the freeness of the game is inherent to the reduction above.

subsets Z small enough ($\rho\sqrt{n}$) so that if a player plays fairly (i.e, the induced distribution on his challenge-tuples is *uniform*), the opponent cannot gain by defecting to this auxiliary part of the game, in particular the optimal response strategy in H remains a Nash equilibrium.

Notice that the number of ($\rho\sqrt{n}$)-sized subsets from $X = [n]$ is $2^{\tilde{O}(\sqrt{n})}$, so encoding the above payoff matrices is within our budget. However, it is not a priori clear that adding the above "intersection constraints" on subsets of size only $O(\sqrt{n})$ is a good enough proxy to imply near-uniform behavior of the players. We show that the above construction is enough to enforce closeness to the uniform distribution *only on a certain set of marginal probability distributions*, thus our reduction is actually from a (general) two-prover game to a strategic game of size $n^{O(\sqrt{n})}$. Making this reduction work, and analyzing the equilibria of the resulting game (mainly showing that indeed if the two player game has a low success value then the strategic game has no good-welfare equilibria) constitutes the bulk of our construction and proof.

3.2 Construction

Given as input the free game instance H from Theorem 2.6, we create the following payoff matrices R and C for Alice and Bob respectively : Each row of R, C corresponds to a pair $(S, a) \in \mathcal{S} \times \mathcal{A}$, and each column of R, C corresponds to a pair $(T, b) \in \mathcal{T} \times \mathcal{B}$. To describe the payoffs of each entry of R and C, recall (from Theorem 2.6) that each challenge-tuple $S \in \mathcal{S}$ ($T \in \mathcal{T}$) of H is a *subset* of $\sqrt{n}/\varepsilon$ original challenges from X (Y). For each entry of R, we set

$$R_{(S,a),(T,b)} = \begin{cases} 0 & \text{if } |\mathcal{N}(S)| < 0.9d|S| \text{ or } |\mathcal{N}(T)| < 0.9d|T| \\ \alpha_S \cdot V(S, T, a(S), b(T)) & \text{if } |\mathcal{N}(S) \cap T| > \frac{d}{10\varepsilon^2} \\ 0 & \text{otherwise} \end{cases},$$

where $\alpha_S := \frac{1}{\Pr_{\mathcal{D}_T}[|\mathcal{N}(S)\cap T| > \frac{d}{10\varepsilon^2}]} \in [1, 1.2)$ is a normalization factor we will need to ensure the existence of an equilibrium with welfare 1 in case that $\omega(H) = 1$, and $\mathcal{D}_T$ is defined as the uniform distribution over $T \in \mathcal{T}$ such that $|\mathcal{N}(T)| \geq 0.9d|T|$. Similarly, we define each entry of C as:

$$C_{(S,a),(T,b)} = \begin{cases} 0 & \text{if } |\mathcal{N}(S)| < 0.9d|S| \text{ or } |\mathcal{N}(T)| < 0.9d|T| \\ \beta_T \cdot V(S, T, a(S), b(T)) & \text{if } |S \cap \mathcal{N}(T)| > \frac{d}{10\varepsilon^2} \\ 0 & \text{otherwise} \end{cases},$$

where $\beta_T := \frac{1}{\Pr_{\mathcal{D}_S}\left[|S\cap\mathcal{N}(T)| > \frac{d}{10\varepsilon^2}\right]}$ with $\mathcal{D}_S$ defined similarly on $\mathcal{S}$.

To enforce uniformity over the challenges, we further define the following payoff matrix. Let $K > 1$ be a constant to be defined shortly. Define the matrix D whose rows are indexed by $\mathcal{S} \times \mathcal{A}$ (the same as R), and whose columns are indexed by all subsets $Z \subseteq X$ of size $\rho\sqrt{n}$, for $\rho := \frac{\varepsilon}{c_2 \cdot K}$ with c_2 a universal constant to be defined shortly. Define the entries of D as follows:

$$D_{(S,a),Z} = \begin{cases} K & \text{if } S \cap Z \neq \emptyset \\ 0 & \text{otherwise} \end{cases}.$$

Notice that since the game H is symmetric ($\mathcal{S} = \mathcal{T}$), the rows of the matrix $D^\top$ can be similarly interpreted as indexing all subsets of $W \subseteq Y$ of size $\rho\sqrt{n}$. Thus:

$$D^\top_{W,(T,b)} = \begin{cases} K & \text{if } W \cap T \neq \emptyset \\ 0 & \text{otherwise} \end{cases}.$$

The final two-player game $\mathcal{G}$ constructed by our reduction is described by the following payoff matrices (for Alice and Bob respectively):

$$P = \begin{pmatrix} R & -D \\ D^\top & 0 \end{pmatrix} \quad ; \quad Q = \begin{pmatrix} C & D \\ -D^\top & 0 \end{pmatrix}.$$

Notice that the above reduction runs in time and space $2^{\tilde{O}(\sqrt{n})}$. We now turn to analyze the equilibria of the above game.

Notation. We use the following shorthands: The $|\mathcal{S}||\mathcal{A}| \times |\mathcal{T}||\mathcal{B}|$ upper left entries of P, Q are called the (R, C)-*part* of $\mathcal{G}$, while the remaining entries of P, Q shall be called the D-*part*. For a (row player) Alice strategy μ, we denote by $\mu|_{R,C}$ the strategy μ projected[4] to the entries in the (R, C)-part of $\mathcal{G}$, and by $\mu|_D$ the strategy μ projected to the entries in the D-part. Define $\nu|_{R,C}, \nu|_D$ analogously for Bob's (the column player) strategy ν.

3.3 Completeness

In this section, we show the completeness part of our reduction :

Lemma 3.2 (Completeness)**.** *If $\omega(H) = 1$, then $\mathcal{G}$ has a Nash equilibrium of welfare at least 1.*

Proof. Consider an optimal strategy pair (a, b) achieving value 1 in the free-game H, i.e., under the uniform distribution over $\mathcal{S} \times \mathcal{T}$, and recall the distributions $\mathcal{D}_S, \mathcal{D}_T$ on challenge-tuples, defined in Section 3.2. Consider the following strategy pair for players in $\mathcal{G}$:

- Alice's (mixed) strategy μ is obtained by choosing $S \sim \mathcal{D}_S$, and setting $a = a(S)$. This specifies a mixed strategy over rows (S, a) of R, C.
- Bob's (mixed) strategy ν is obtained by choosing $T \sim \mathcal{D}_T$, and setting $b = b(T)$. This specifies a mixed strategy over columns (T, b) of R, C.

By the definition of R, C and third and fourth item in Theorem 2.6, we have $\pi_{\mu,\nu}(\mathcal{G}) = 1$ (this is due to the way we defined α_S, β_T). In particular, both Alice's and Bob's payoff is 1. It remains to show that (μ, ν) is an equilibrium, that is neither Alice nor Bob has an incentive to deviate to some other strategy. To show this, it suffices to argue that Bob does not have a better response to the strategy μ played by Alice (the argument will be symmetric for the reverse case). Indeed, suppose Bob has a better response strategy. By the averaging principle, this strategy can assumed to be pure, i.e., a single column e_i of $\mathcal{G}$. There are two cases:

- If e_i is a column of C, then the column e_i corresponds some fixed challenge-tuple T^* of size $\sqrt{n}/\varepsilon$ from the set Y with an assignment b on T^*. Then the payoff of Bob should be $\sum_{S \in \mathcal{S}} C_{(S,a),(T^*,b)}$. By definition, if $|\mathcal{N}(T^*)| < 0.9d|T^*|$, the payoff is 0. Hence without loss of generality, T^* is such that $|\mathcal{N}(T^*)| > 0.9d|T^*|$. But for such T^* :

$$\mathbb{E}_{S \sim \mathcal{D}_S} C_{(S,a),(T^*,b)} \leq \frac{1}{\Pr_{\mathcal{D}_S}[|S \cap \mathcal{N}(T^*)| > \frac{d}{10\varepsilon^2}]} \mathbb{E}_{S \sim \mathcal{D}_S}[\mathbb{I}_{|S \cap \mathcal{N}(T^*)| > \frac{d}{10\varepsilon^2}}] = 1$$

  where $\mathbb{I}$ is the indicator variable of the event "$|S \cap \mathcal{N}(T^*)| > \frac{d}{10\varepsilon^2}$". Therefore Bob does not have an incentive to deviate to e_i.

[4] Unless otherwise stated, this is not a restriction but rather just the measure obtained by truncating all the entries not in R, C. It would need to be normalized to become a restriction.

- If e_i is a column of D, then e_i corresponds a fixed set of size $\rho\sqrt{n} = \frac{\varepsilon}{c_2 \cdot K}\sqrt{n}$ of X, call it Z. By definition of the payoffs of D and since Alice's strategy μ is *uniform* over $\mathcal{S}$, Bob's payoff from defecting to e_i is $K \cdot \Pr_{S\sim\mathcal{D}_S}[S \cap Z \neq \emptyset]$. But we have

$$\begin{aligned}
&\Pr_{S\sim\mathcal{D}_S}[S \cap Z \neq \emptyset] = 1 - \Pr_{S\sim\mathcal{D}_S}[S \cap Z = \emptyset] \leq 1 - \Pr_{S\sim\mathcal{U}(\mathcal{S})}[S \cap Z = \emptyset] + o(1) \\
&= 1 - \left(\frac{\binom{n-|Z|}{|S|}}{\binom{n}{|S|}}\right) + o(1) = 1 - \frac{(n-|S|)(n-|S|-1)\dots(n-|Z|-|S|+1)}{n(n-1)\dots(n-|Z|+1)} + o(1) \\
&\leq 1 - \left(\frac{n-|Z|-|S|}{n-|Z|}\right)^{|Z|} + o(1) = 1 - \left(1 - \frac{|S|}{n-|Z|}\right)^{|Z|} + o(1) \\
&\leq 1 - \left(1 - \frac{|S|}{0.9 \cdot n}\right)^{|Z|} + o(1) \leq 1 - e^{-\frac{|S||Z|}{0.9 \cdot n}} + o(1) \leq \frac{2}{0.9 \cdot c_2 \cdot K}.
\end{aligned}$$

where we have used the fact that $1 + x \leq e^x \leq 1 + x + x^2/2$ and $\Pr_{S\sim\mathcal{U}(\mathcal{S})}[S \cap Z = \emptyset] - o(1) \leq \Pr_{S\sim\mathcal{D}_S}[S \cap Z = \emptyset]$. Here, α is some universal constant such that $\alpha < 1$. Therefore, Bob's payoff from the strategy e_i is upper bounded by $\frac{2K}{0.9 \cdot c_2 \cdot K} = \frac{2}{0.9 \cdot \alpha c_2}$. Thus if we set c_2 so that $2/0.9 \leq c_2$, say $c_2 = 5/2$, Bob does not have an incentive to deviate to e_i.

Repeating a symmetric argument for Alice, we conclude that (μ, ν) is a (perfect) Nash equilibrium with social welfare $= 1$, as desired. □

3.4 Soundness

In this section, we give a proof of soundness part of our reduction :

Theorem 3.3 (Soundness)**.** *There is some constant $\varepsilon^* = \Omega(1)$ such that if $\omega(G) \leq \varepsilon$, then any ε^*-equilibrium of $\mathcal{G}$ has value at most $O(\varepsilon)$.*

We first argue that the projection of any approximate Nash Equilibrium to the (R, C)-part of $\mathcal{G}$ can be decomposed into "uniform+ small noise" components. To this end, for any Alice strategy μ, let

$$\tilde{\mu}(i) := \frac{1}{|S|} \sum_{S: S \ni i} \mu(S)$$

denote the marginal frequency of challenge $i \in X$ under μ. Define $\tilde{\nu}(i)$ analogously with respect to any Bob strategy ν.

Lemma 3.4 (Decomposition lemma)**.** *Suppose (μ, ν) is an ε^*-Nash Equilibrium of $\mathcal{G}$, where $\varepsilon^* < 1/2$. Then one can write $\mu|_{R,C} = \mu_1 + \mu_2$ such that*

- *$\tilde{\mu}_1$ is near uniform; that is, for all $i \in X$, $\tilde{\mu}_1(i) < 5/n$.*
- *$\|\mu_2\|_1 < \frac{2}{K}$.*

An analogues decomposition holds for $\nu|_{R,C}$.[5]

Proof. By Proposition A.5, we know that the entries in R and C are at most $1 + O(c_\varepsilon) < 1.2$. Therefore, by convexity, the payoff for Alice and Bob is bounded above by 1.2 for any (μ, ν).

For $\mu|_{R,C}$, consider the following recursive procedure for decomposition:

[5] Note that $\mu_1, \mu_2, \mu|_{R,C}$ are not distributions, but non-negative vectors over rows of payoff matrices, P and Q

- Set $k = 1$, $\mu_1^1 = \mu, \mu_2^1 = 0, B \longleftarrow \emptyset$.
- If (a) $\|\tilde{\mu}_1^k\|_\infty < 5/n$ or (b) $\|\mu_2\|_1 \geq 2/K$ then halt. Otherwise pick $i_k = \arg\max \tilde{\mu}_1^k(i)$. Add i_k to B.
- Define $\mu_1^{k+1} = \mu_1^k - \mu_1^k|_{i_k \in S}$ and $\mu_2^{k+1} = \mu_2^k + \mu_1^k|_{i_k \in S}$
- Set $k \longleftarrow k + 1$ and repeat.

First we argue that above algorithm must terminate, that is if (a) does not hold at any round, then (b) must hold after some number of rounds. So long as (a) does not hold, we know that there exists some i_k for each k such that $\tilde{\mu}_1^k(i_k) \geq 5/n$. Then,

$$\|\mu_1^k|_{i_k \in S}\|_1 = \sum_{i_k \in S} \mu_1^k(S) = |S| \cdot \tilde{\mu}_1^k(i_k) > |S| \cdot \frac{5}{n} = \frac{5}{\varepsilon\sqrt{n}},$$

that is, the ℓ_1 norm of μ_2 increases in each round by at least $\frac{5}{\varepsilon\sqrt{n}}$. Thus after $T = \frac{2\varepsilon}{5K}\sqrt{n}$ rounds,

$$\|\mu_2^T\|_1 > \frac{5}{\varepsilon\sqrt{n}} \cdot T = \frac{2}{K}$$

which violates (b). Therefore, our algorithm must halt before $k = T$.

Now, we show that if our algorithm terminates, it should either output a valid decomposition, or (μ, ν) is not a valid ε^*-equilibrium. If our algorithm halts via (a) before reaching (b), μ_1 and μ_2 become our desired decomposition. Now suppose instead our algorithm halts via (b). Then consider the set of indices B in our algorithm. If Bob deviates to the corresponding column for B in D, which we know that exists since $|B| \leq T = \frac{2\varepsilon}{5K}\sqrt{n}$. (Choosing any set in D that contains B suffices) Then his new payoff is

$$K \cdot \Pr_{S \sim \mu}[B \cap S \neq \emptyset] \geq K \cdot \|\mu_2\|_1 \geq 2,$$

Thus, Bob's payoff increases by more than 1/2 by deviating to such B. Therefore, (μ, ν) is not an ε^*-equilibrium.

Repeating a symmetric argument for Alice, we conclude that any equilibrium with distribution (μ, ν) must satisfy above decomposition properties. □

Now if ε^*-Nash Equilibrium indeed satisfies above decomposition, we then show that its payoff cannot be large:

Lemma 3.5 (Payoff Bound)**.** *Suppose μ and ν satisfies the condition in Lemma 3.4, and forms an ε^*-Nash Equilibrium. Then the payoff of the game is at most $O(\varepsilon)$.*

Proof. Let $M = \frac{1}{2} \cdot (R + C)$. Then we can decompose the value of the game in a following way:

$$\begin{aligned}
\mu \cdot M \cdot \nu^T &\leq \mu_1 M \nu_1^T + \mu_2 M \nu_1^T + \mu_1 M \nu_2^T + \mu_2 M \mu_2^T \\
&\leq \mu_1 M \nu_1^T + (1 + c_\varepsilon)\left(\|\mu_2\|_1\|\nu_1\|_1 + \|\mu_1\|_1\|\nu_2\|_1 + \|\mu_2\|_1\|\nu_2\|_1\right) \\
&\leq \mu_1 M \nu_1^T + \frac{6}{K}
\end{aligned}$$

Note that any mass on $\mu|_D$ and $\nu|_D$ will contribute a zero payoff by our construction. Here c_ε is a small constant dependent on ε introduced in Proposition A.5, which can be assumed to be < 0.1.

Now we wish to bound payoff given by R and C under (μ_1, ν_1) separately. Recall that our original assumption about the **2CSP** instance gives us :

$$\frac{1}{|E|} \sum_{(i,j) \in E} V(i,j) < \varepsilon$$

where $V(i,j)$ refers to the "squashed" value of the constraint between $i \in X$ and $j \in Y$, in a sense that it is whether (i,j) is being satisfied within the challenge-tuples, and E is the set of constraints. Then the following chain of inequalities hold :

$$\begin{aligned}
\mu_1 \cdot R \cdot \nu_1^T &\leq (1 + c_\varepsilon) \cdot \frac{10\varepsilon^2}{d} \sum_{(i,j) \in E} \Pr_{S \sim \mu_1, T \sim \nu_1} [i \in S, j \in T, V(i,j,a,b) = 1] \\
&= (1 + c_\varepsilon) \cdot \frac{10\varepsilon^2}{d} \sum_{(i,j) \in E} \Pr_{S \sim \mu_1, T \sim \nu_1} [i \in S, j \in T] V(i,j) \\
&= (1 + c_\varepsilon) \cdot \frac{10\varepsilon^2}{d} \sum_{(i,j) \in E} V(i,j) \Pr_{S \sim \mu_1} [i \in S] \Pr_{T \sim \nu_1} [j \in T] \\
&\leq (1 + c_\varepsilon) \cdot \frac{10\varepsilon^2}{d} \sum_{(i,j) \in E} V(i,j) \frac{5|S|}{n} \frac{5|T|}{n} \leq (1 + c_\varepsilon) \cdot \frac{10\varepsilon^2}{d} \frac{\varepsilon \cdot 5^2 \cdot |E| \cdot |S| \cdot |T|}{n^2} \\
&= (1 + c_\varepsilon) \cdot 250 \cdot \varepsilon \leq O(\varepsilon)
\end{aligned}$$

The first inequality holds since we need at least $\frac{d}{10\varepsilon^2}$ pairs of (i,j) to gain the payoff. Second equality holds since our game is a free game, that is μ_1 and ν_1 are independent.

Repeating a symmetric argument for C, $\mu_1 M \nu_1^T \leq O(\varepsilon)$. Setting $K = O(1/\varepsilon)$, the payoff is at most $O(\varepsilon)$. □

Proof of Theorem 3.3. Suppose (μ, ν) forms an ε^*-equilibrium for $\varepsilon^* = 1 - \Omega(1)$, and $\varepsilon^* < 1/2$, and our original PCP had value $< \varepsilon$. By Lemma 3.4, we know that $\mu|_{R,C}$ and $\nu|_{R,C}$ can be decomposed, if they are indeed an ε^*-equilibrium. But then by Lemma 3.5, $\pi_{\mu,\nu}(\mathcal{G}) = O(\varepsilon)$, which completes the proof. □

References

[AAK+07] Noga Alon, Alexandr Andoni, Tali Kaufman, Kevin Matulef, Ronitt Rubinfeld, and Ning Xie. Testing k-wise and almost k-wise independence. In *STOC*, pages 496–505. ACM, 2007.

[ABC13] Per Austrin, Mark Braverman, and Eden Chlamtác. Inapproximability of NP-complete variants of Nash equilibrium. *Theory of Computing*, 9(3):117–142, 2013.

[AIM14] Scott Aaronson, Russell Impagliazzo, and Dana Moshkovitz. Am with multiple merlins. *CoRR*, abs/1401.6848, 2014.

[AKS98] Noga Alon, Michael Krivelevich, and Benny Sudakov. Finding a large hidden clique in a random graph. In *SODA*, pages 594–598, 1998.

[Axe87] Robert Axelrod. The evolution of strategies in the iterated prisoners dilemma. *The dynamics of norms*, pages 1–16, 1987.

[Axe00] Robert Axelrod. On six advances in cooperation theory. *Analyse & Kritik*, 22(1):130–151, 2000.

[BBB+13] Maria-Florina Balcan, Christian Borgs, Mark Braverman, Jennifer Chayes, and Shang-Hua Teng. Finding endogenously formed communities. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 767–783. SIAM, 2013.

[BBM07] Hartwig Bosse, Jaroslaw Byrka, and Evangelos Markakis. New algorithms for approximate nash equilibria in bimatrix games. In Xiaotie Deng and FanChung Graham, editors, *Internet and Network Economics*, volume 4858 of *Lecture Notes in Computer Science*, pages 17–29. Springer Berlin Heidelberg, 2007.

[BH13] Fernando G. S. L. Brandão and Aram Wettroth Harrow. Quantum de finetti theorems under local measurements with applications. In *STOC*, pages 861–870, 2013.

[BHHS09] Boaz Barak, Moritz Hardt, Thomas Holenstein, and David Steurer. Subsampling semidefinite programs and max-cut on the sphere. *CoRR*, abs/0911.5526, 2009.

[BR13] Quentin Berthet and Philippe Rigollet. Complexity theoretic lower bounds for sparse principal component detection. In *Conference on Learning Theory*, pages 1046–1066, 2013.

[CDT06] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Computing nash equilibria: Approximation and smoothed complexity. *CoRR*, abs/cs/0602043, 2006.

[CS03] Vincent Conitzer and Tuomas Sandholm. Complexity results about Nash equilibria. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI-03)*, pages 765–771, 2003.

[DGGP10] Yael Dekel, Ori Gurel-Gurevich, and Yuval Peres. Finding hidden cliques in linear time with high probability. *CoRR*, abs/1010.2997, 2010.

[DGP06] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a nash equilibrium. pages 71–78. ACM Press, 2006.

[DMP06] Constantinos Daskalakis, Aranyak Mehta, and Christos Papadimitriou. A note on approximate nash equilibria. In Paul Spirakis, Marios Mavronicolas, and Spyros Kontogiannis, editors, *Internet and Network Economics*, volume 4286 of *Lecture Notes in Computer Science*, pages 297–306. Springer Berlin Heidelberg, 2006.

[DMP07] Constantinos Daskalakis, Aranyak Mehta, and Christos Papadimitriou. Progress in approximate nash equilibria. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, EC '07, pages 355–358, New York, NY, USA, 2007. ACM.

[FK00] Uriel Feige and Robert Krauthgamer. Finding and certifying a large hidden clique in a semirandom graph. *Random Struct. Algorithms*, 16(2):195–208, 2000.

[FR10] Uriel Feige and Dorit Ron. Finding hidden cliques in linear time. In *AOFA*, 2010.

[GZ89] I. Gilboa and E. Zemel. Nash and correlated equilibria: Some complexity considerations. page 1:8093. Games and Economic Behavior, 1989.

[HK09] Elad Hazan and Robert Krauthgamer. How hard is it to approximate the best nash equilibrium? In *SODA*, pages 720–727, 2009.

[IP01] Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-sat. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001.

[Jer92] Mark Jerrum. Large cliques elude the metropolis process. *Random Struct. Algorithms*, 3(4):347–360, 1992.

[KPS06] Spyros C. Kontogiannis, Panagiota N. Panagopoulou, and Paul G. Spirakis. Polynomial algorithms for approximating nash equilibria of bimatrix games. In Paul Spirakis, Marios Mavronicolas, and Spyros Kontogiannis, editors, *Internet and Network Economics*, volume 4286 of *Lecture Notes in Computer Science*, pages 286–296. Springer Berlin Heidelberg, 2006.

[Kuc95] Ludek Kucera. Expected complexity of graph partitioning problems. *Discrete Applied Mathematics*, 57(2-3):193–212, 1995.

[LMM03] Richard J. Lipton, Evangelos Markakis, and Aranyak Mehta. Playing large games using simple strategies. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, EC '03, pages 36–41, New York, NY, USA, 2003. ACM.

[LMS+13] Daniel Lokshtanov, Dániel Marx, Saket Saurabh, et al. Lower bounds based on the exponential time hypothesis. *Bulletin of EATCS*, 3(105), 2013.

[MR08] Dana Moshkovitz and Ran Raz. Sub-constant error low degree test of almost-linear size. *SIAM J. Comput.*, 38(1):140–180, 2008.

[TS07] Haralampos Tsaknakis and Paul G. Spirakis. An optimization approach for approximate nash equilibria. In Xiaotie Deng and FanChung Graham, editors, *Internet and Network Economics*, volume 4858 of *Lecture Notes in Computer Science*, pages 42–56. Springer Berlin Heidelberg, 2007.

A Appendix

A.1 Symmetry and Concentration properties of the free-game of [AIM14]

In this section, we prove third and fourth properties in Theorem 2.6, which will be used in Section 3.3

Proposition A.1 (Symmetry)**.** *Let $\phi = (X, Y, E)$ be a* **2CSP** *instance formed by [MR08], where X, Y and E refers to the set of left, right vertices and edges between them respectively when viewed as a bipartite graph. Define $\tilde{\phi} = (\tilde{X}, \tilde{Y}, \tilde{E})$ as "symmetrized" version of ϕ (viewed as a bipartite graph), which is formed by setting $\tilde{X} = \tilde{Y} = X \cup Y$ and*

$$\tilde{E} = \{(u, v)|(u, v) \in E \text{ or } (v, u) \in E\} \cup \{(v, v)|v \in X \text{ or } v \in Y\}$$

and setting a uniform distribution over such $\tilde{E}$. $e \in \{(u, v)|(u, v) \in E$ or $(v, u) \in E\}$ is considered satisfied if the assignment to u and v is considered satisfied in the original ϕ. $e \in \{(v, v)|v \in X$ or $v \in Y\}$ is considered satisfied if the assignment to both v are equal.

Then if $\omega(\phi) = 1$ then $\omega(\tilde{\phi}) = 1$. If $\omega(\phi) < \delta$, then $\omega(\tilde{\phi}) < 2\delta$

Proof. If $\omega(\phi) = 1$, then $\omega(\tilde{\phi}) = 1$, since we can simply copy the assignment for ϕ to $\tilde{\phi}$. Now suppose $\omega(\phi) \leq \delta$. Suppose we pick a random $e \in \tilde{E}$. If $e \in \{(u, v)|(u, v) \in E$ or $(v, u) \in E\}$, the probability of e being satisfied is at most δ, and the probability of picking such e is $2d/(2d + 2)$. Then:

$$\omega(\tilde{\phi}) \leq \frac{2d}{2d+2}\delta + \frac{2}{2d+2} \cdot 1 = \frac{d\delta + 1}{d+1} \leq 2\delta$$

The last inequality holds since $1 \leq \delta d + 2\delta$, as $d > 1/\delta$ in [MR08]. □

Therefore, we can assume from now on that for any ϕ formed by [MR08] and the corresponding Clause/Variable game $G_\phi = (X, Y, A, B, V)$ as:

- $X = Y$ and $A = B$
- For any subset of size k say $S \subseteq X$, the number of variables in Y that shares a challenge (edge) with some variable in S denoted as $\mathcal{N}(S)$ is at least $|S|$ and at most $(d + 1)|S|$, thus $\Theta(|S|)$.

We henceforth denote by d the degree of the symmetric game G_ϕ.

Proposition A.2 (Expansion concentration)**.** *Under the uniform distribution over $\mathcal{S}$, the number of neighbors for $S \subseteq X$ where $|S| = \sqrt{n}/\varepsilon$ is concentrated.*

$$\Pr_{S \sim \mathcal{U}(\mathcal{S})} [0.9d|S| < |\mathcal{N}(S)| \leq d|S|] \geq 1 - o(1) \tag{3}$$

Proof. For each $i \in Y$ define the following random variable :

$$X_i = \begin{cases} 1 & \text{if } i \in \mathcal{N}(S) \\ 0 & \text{otherwise} \end{cases}$$

Then $|\mathcal{N}(S)| = \sum_{i\in Y} X_i$. Also observe that for $i \in Y$ to be in $\mathcal{N}(S)$, one of neighbors of i must be in S. Therefore

$$\begin{aligned}
\mathbb{E}[|\mathcal{N}(S)|] &= \sum_{i\in Y} \mathbb{E}[X_i] = \sum_{i\in Y} \Pr[X_i = 1] = n \cdot \left(1 - \frac{\binom{n-d}{|S|}}{\binom{n}{|S|}}\right) \\
&\geq n \cdot \left(1 - \left(1 - \frac{d}{n}\right)^{|S|}\right) \geq n \cdot \left(1 - e^{-\frac{d}{\varepsilon\sqrt{n}}}\right) \\
&= \frac{d\sqrt{n}}{\varepsilon} - \frac{d^2}{2\varepsilon^2} + o(1)
\end{aligned}$$

Now Proposition A.1 ensures that $|\mathcal{N}(S)| \leq d|S| = \frac{d\sqrt{n}}{\varepsilon}$. Thus we can view $d|S| - |\mathcal{N}(S)|$ as a nonnegative random variable and bound the concentration using Markov inequality :

$$\begin{aligned}
&\Pr[0.9d|S| \leq |\mathcal{N}(S)| \leq d|S|] = \Pr[d|S| - |\mathcal{N}(S)| \leq \frac{d|S|}{10}] \\
&= 1 - \Pr\left[d|S| - |\mathcal{N}(S)| > \frac{d|S|}{10}\right] \geq 1 - \frac{d^2}{2\varepsilon^2} \cdot \frac{10}{d|S|} \\
&= 1 - o(1)
\end{aligned}$$

This completes the proof. □

To show the final concentration property, we need to use two well-known facts.

Fact A.3 (Approximation of $\binom{n}{k}$)**.** *If n is large and $k \leq \sqrt{n}/\varepsilon$, then*

$$e^{-1/\varepsilon^2} \frac{n^k}{k!} \leq \binom{n}{k} \leq \frac{n^k}{k!} \tag{4}$$

Fact A.4 (Stirling's approximation)**.**

$$n! = \Theta\left(\sqrt{2\pi n}\left(\frac{n}{e}\right)^n\right) \tag{5}$$

Proposition A.5. *Fix S that satisfies expansion concentration, that is $|\mathcal{N}(S)| = \kappa_S d|S| > 0.9d|S|$. Then under the uniform distribution over $T \subseteq Y$ with $|T| = \sqrt{n}/\varepsilon$ satisfies the following concentration inequalities.*

$$\Pr_T\left[|\mathcal{N}(S) \cap T| < \frac{d}{c_1\varepsilon^2}\right] \leq \text{poly}(1/\varepsilon)e^{-\text{poly}(1/\varepsilon)} \tag{6}$$

where c_1 is some universal constant such that $\frac{\log c_1 + 1}{c_1} < 0.9$.

Proof. Set $t = \frac{d}{c_1\varepsilon^2}$ then

$$
\begin{aligned}
&\Pr\left[|\mathcal{N}(S)\cap T| < t\right] = \sum_{k=1}^{t} \Pr\left[|\mathcal{N}(S)\cap T| = k\right] \le t\cdot \Pr\left[|\mathcal{N}(S)\cap T| = t\right] \\
&= t\cdot \frac{\binom{n-|\mathcal{N}(S)|}{|T|-t}\cdot\binom{|\mathcal{N}(S)|}{t}}{\binom{n}{|T|}} \le t\cdot \frac{\frac{(n-|\mathcal{N}(S)|)^{|T|-t}}{(|T|-t)!}\cdot\frac{|\mathcal{N}(S)|^t}{t!}}{e^{-1/\varepsilon^2}\frac{n^{|T|}}{|T|!}} \\
&= t\cdot e^{1/\varepsilon^2}\cdot\binom{|T|}{t}\cdot\left(1-\frac{|\mathcal{N}(S)|}{n}\right)^{|T|-t}\cdot\left(\frac{|\mathcal{N}(S)|}{n}\right)^t \\
&\le t\cdot e^{1/\varepsilon^2}\cdot\frac{1}{t!}\cdot\left(1-\frac{|\mathcal{N}(S)|}{n}\right)^{|T|-t}\cdot\left(\frac{|\mathcal{N}(S)|\cdot|T|}{n}\right)^t \\
&= \frac{t\cdot e^{1/\varepsilon^2}}{t!}\cdot\left(1-\frac{|\mathcal{N}(S)|}{n}\right)^{|T|-t}\cdot\left(\frac{\kappa_S d}{\varepsilon^2}\right)^t \\
&\approx \frac{\sqrt{t}\cdot e^t\cdot e^{1/\varepsilon^2}}{\sqrt{2\pi}t^t}\cdot e^{-\frac{\kappa_S d|S|(|T|-t)}{n}}\cdot e^{t\cdot\log\frac{\kappa_S d}{\varepsilon^2}} \\
&= \frac{1}{\sqrt{2\pi}}\sqrt{t}\cdot e^{t\cdot(\log(\kappa_S c_1 t)+1-\log t)}\cdot e^{o(1)}\cdot e^{-\frac{\kappa_S d}{\varepsilon^2}} \\
&\le \frac{1}{\sqrt{2\pi}}\sqrt{t}\cdot e^{t\log(ec_1)}\cdot e^{o(1)}e^{-0.9c_1 t} = \mathrm{poly}(1/\varepsilon)e^{-\mathrm{poly}(1/\varepsilon)} = c_\varepsilon
\end{aligned}
$$

where second inequality uses (4) and the approximation uses (5). □

Therefore, we can set $c_1 = 10$.

Corollary A.6. *By choosing ε as a small enough constant, we can assume that for any $S \subseteq X$ of size $\sqrt{n}/\varepsilon$ such that $|\mathcal{N}(S)| > 0.9d|S|$:*

$$\Pr\left[|\mathcal{N}(S)\cap T| \ge \frac{d}{10\varepsilon^2}\right] \ge 1 - c_\varepsilon > 0.9$$

ECCC ISSN 1433-8092
http://eccc.hpi-web.de