Electronic Colloquium on Computational Complexity, Report No. 25 (2012)

Approximating the minmax value of 3-player games within a constant is as hard as detecting planted cliques

Kord Eickmeyer[1], Kristoffer Arnstfelt Hansen[2], and Elad Verbin[2]

[1]National Institute of Informatics, Tokyo
[2]Aarhus University*

March 23, 2012

We consider the problem of approximating the minmax value of a multiplayer game in strategic form. We argue that in 3-player games with 0-1 payoffs, approximating the minmax value within an additive constant smaller than $\xi/2$, where $\xi = \frac{3-\sqrt{5}}{2} \approx 0.382$, is not possible by a polynomial time algorithm. This is based on assuming hardness of a version of the so-called planted clique problem in Erdős-Rényi random graphs, namely that of *detecting* a planted clique. Our results are stated as reductions from a promise graph problem to the problem of approximating the minmax value, and we use the detection problem for planted cliques to argue for its hardness. We present two reductions: a randomized many-one reduction and a deterministic Turing reduction. The latter, which may be seen as a derandomization of the former, may be used to argue for hardness of approximating the minmax value based on a hardness assumption about *deterministic* algorithms. Our technique for derandomization is general enough to also apply to related work about ϵ-Nash equilibria.

*Hansen and Verbin acknowledge support from the Danish National Research Foundation and The National Science Foundation of China (under the grant 61061130540) for the Sino-Danish Center for the Theory of Interactive Computation, within which this work was performed. They also acknowledge support from the Center for Research in Foundations of Electronic Markets (CFEM), supported by the Danish Strategic Research Council.

ISSN 1433-8092

1. Introduction

We consider games in strategic form between 3 players. These are given by a finite strategy space for each player, S_1,S_2, and S_3 (also called the *pure strategies*), together with utility functions $u_1, u_2, u_3 : S_1 \times S_2 \times S_3 \to \mathbb{R}$. We can identify the strategy spaces with the sets $[n_1]$,$[n_2]$, and $[n_3]$, where $n_i = |S_i|$. We shall refer to this as a $n_1 \times n_2 \times n_3$ game. In this paper only the utilities for Player 1 are relevant.

Let Δ_1,Δ_2, and Δ_3 be the sets of probability distributions over S_1,S_2, and S_3 respectively; these are also called *mixed strategies*. The minmax value (also known as the threat value) for Player 1 is given by:

$$\min_{(\sigma_2,\sigma_3)\in\Delta_2\times\Delta_3} \max_{\sigma_1\in\Delta_1} \mathop{\mathrm{E}}_{a_i\sim\sigma_i}[u_1(a_1, a_2, a_3)]$$

A strategy profile (σ_2, σ_3) for Player 2 and Player 3 for which this values is obtained is called an optimal minmax profile. It is not hard to see that Player 1 may always obtain the maximum by a pure strategy, i.e., the minmax value is equal to:

$$\min_{(\sigma_2,\sigma_3)\in\Delta_2\times\Delta_3} \max_{a_1\in S_1} \mathop{\mathrm{E}}_{\substack{a_2\sim\sigma_2\\a_3\sim\sigma_3}}[u_1(a_1, a_2, a_3)] \tag{1}$$

The corresponding notion of minmax value in finite two-player games is a fundamental notion of game theory. Minmax values have been studied much less in multi-player player games, but are arguably also here of fundamental interest. In particular the minmax value of such games is crucial for the statements as well as proofs of the so-called *folk theorems* that characterize the Nash equilibria of *repeated games*. The problem of *computing* the minmax value of a multi-player game was first considered only recently by Borgs et al. [5], exactly in the context of studying computational aspects of the folk theorem. In particular they show that approximating the minmax value of a 3 player game within a specific inverse polynomial additive error is NP hard.

Here, to be able to talk meaningfully about approximation within an additive error, we assume that all payoffs have been *normalized* to be in the interval between 0 and 1. The question of approximating the minmax value was considered further by Hansen et al. [12]. Using a "padding" construction it was observed that the NP hardness result of Borgs et al. extends to any inverse polynomial additive error. This was complemented by a quasipolynomial approximation algorithm obtaining an approximation to within an arbitrary $\epsilon > 0$, which was obtained using a result of Lipton and Young [18], stating that in a $n \times n$ matrix game with payoffs between 0 and 1, each player can guarantee a payoff within any $\epsilon > 0$ of the value of the game using strategies that simply consist of a uniform choice from a multiset of $\lceil \ln n/(2\epsilon^2) \rceil$ pure strategies. We summarize these results by the following theorem.

Theorem 1 ([5, 12])**.** *For any constant $\epsilon > 0$ it is NP hard to approximate the minmax value of an $n \times n \times n$ game with 0-1 payoffs within additive error $1/n^\epsilon$. On the other hand, there is an algorithm that, given $\epsilon > 0$ and a $n \times n \times n$ game with payoffs between 0 and 1, approximates the minmax value from above with additive error at most ϵ in time $n^{O(\log(n)/\epsilon^2)}$.*

This naturally raises the question of whether it is possible to approximate the minmax value within any constant $\epsilon > 0$ in polynomial time, or even whether it is possible to approximate the minmax value within *some* nontrivial additive constant $0 < \epsilon < 1/2$ in polynomial time. Due to the quasipolynomial time algorithm above, it is unlikely that the theory of NP completeness can shed light on this question.

A similar situation is present for the problem of computing a Nash equilibrium in two player bimatrix games. Celebrated recent results [8, 6] show that this problem is complete for the complexity class PPAD. On the other hand several works provide algorithms for computing an ϵ-Nash equilibrium. An ϵ-Nash equilibrium in a $n \times n$ bimatrix game with payoffs between 0 and 1 can be computed in time $n^{O(\log(n)/\epsilon^2)}$ [17], by an algorithm similar to the one described above for the minmax value. As for polynomial time algorithms, several algorithms have been devised for decreasing additive error ϵ (see e.g. [20] for references).The current best such algorithm achieves $\epsilon = 0.3393$ [20]. How well a Nash equilibrium can be approximated in the sense of ϵ-Nash equilibria is a major open question. Having a polynomial time algorithm, polynomial also in $1/\epsilon$, or in other words having a fully polynomial time approximation scheme (FPTAS), would imply that every problem in the class PPAD would be solvable in polynomial time [6]. Currently there is no evidence for or against the existence of a polynomial time algorithm for any fixed $\epsilon > 0$, or in other words a polynomial time approximation scheme (PTAS) for computing ϵ-Nash equilibria.

The planted clique problem

Our result depends on assuming hardness of the so-called planted clique problem (more precisely, the *detection* variant). Let $G_{n,p}$ denote the distribution of Erdős-Rényi random graphs on n vertices where each potential edge is included in the graph independently at random with probability p. Most frequently the case of $p = 1/2$ is considered, but we will be interested in having $p > 0$ be a small constant. This choice is made in order to get a conclusion as strong as possible from our proof (cf. Remark 11).

It is well known that in almost every graph from $G_{n,p}$ the largest clique is of size $2\log_{1/p} n - O(\log\log n)$ [4]. The hidden clique problem is defined using the distribution $G_{n,p,k}$ [14, 16] of graphs on n vertices defined as follows: A graph G is picked according to $G_{n,p}$, then a set of k vertices are chosen uniformly at random and connected to form a clique. Thus apart from the planted k-clique the graph is completely random. The (*search* variant of the) planted clique problem is then defined as follows: Given a graph G chosen at random from $G_{n,p,k}$, find a k-clique in the graph G. Note that when the parameter k is significantly larger than $2\log_{1/p} n$, the planted clique is with high probability the unique maximum clique in the graph, and thus it also makes sense to talk about finding the planted clique, with high probability.

The planted clique problem is known as a difficult combinatorial problem. Indeed the current best polynomial time algorithms for solving the planted clique problem [1, 10] are only known to work when $k = \Omega(\sqrt{n})$. We may compare this with the observation due to Kučera [16] that for $k \geq C\sqrt{n \log n}$ when C is a suitably large constant, the vertices of the clique would almost surely be the vertices of largest degree, and hence easy to find. The planted clique problem has also been proposed as a basis for a cryptographic one-way function [15]. For this application, however, the size of the planted clique is $k = (1+\epsilon)\log_{1/p} n$, which is smaller than the expected size of the largest clique.

The planted clique *detection* problem is defined as follows: Given a graph G chosen at random from either (i) $G_{n,p}$, or (ii) $G_{n,p,k}$, decide which is the case.

1.1. Our Results

We show a relationship between the task of approximating the minmax value in a 3-player game and the planted clique detection problem. Our result builds heavily on the ideas of the work of Hazan and Krauthgamer in [13] (see also [19]). These are described in the next section.

In our results we prove hardness of approximating the minmax value, and aim to obtain a conclusion as strong as possible, while maintaining a reasonable hardness assumption.

We will actually state our results using the following promise[1] graph problem Gap-DBS, parametrized by numbers $0 < c_1 < c_2$ and $\eta > 0$. Let $G = (V_1, V_2, E)$ be a bipartite graph. For $S \subseteq V_1$, $T \subseteq V_2$ the *density* of the subgraph induced by S and T is given by $\mathrm{d}(S,T) = \frac{|E(S,T)|}{|S||T|}$. Note that if we let A denote the adjacency matrix of A and let u_S and u_T be the probability vectors that are uniform on the sets S and T, then we have $\mathrm{d}(S,T) = u_S^{\mathsf{T}} A u_T$.

> Gap Dense Bipartite Subgraph (Gap-DBS)
>
> *Input:* Bipartite graph $G = (V_1, V_2, E)$, $|V_1| = |V_2| = n$
>
> *Promise:* Either
>
> (i) There exist $S \subseteq V_1$, $T \subseteq V_2$, $|S| = |T| = c_2 \ln n$, such that $\mathrm{d}(S,T) \geq 1 - \eta$, or
>
> (ii) For all $S \subseteq V_1$, $T \subseteq V_2$, $|S| = |T| = c_1 \ln n$, it holds that $\mathrm{d}(S,T) \leq \eta$.
>
> *Problem:* Decide which of these is the case

We also introduce the following gap problem for the minmax value of 3-player games with 0-1 payoffs, parametrized by numbers $0 \leq \alpha < \beta \leq 1$

[1] Clearly if there exist sets S and T with $|S| = |T| = c_2 \ln n$ and $\mathrm{d}(S,T) \geq 1 - \eta$, there exist subsets $S' \subseteq S$ and $T' \subseteq T$ with $|S'| = |T'| = c_1 \ln n$ and $\mathrm{d}(S',T') \geq 1 - \eta$ as well.

Gap 3 Player minmax (Gap-minmax)	
Input:	$n \times n \times n$ game G with 0-1 payoffs
Promise:	Either (i) The minmax value for Player 1 in G is at most α, or (ii) The minmax value for Player 1 in G is at least β.
Problem:	Decide which of these is the case

We are now ready to state our main results. Throughout the paper $\xi = \frac{3-\sqrt{5}}{2} \approx 0.382$ is the smaller of the two roots of $x^2 - 3x + 1 = 0$, which is also known as $1 - \varphi$, where φ is the conjugate golden ratio.

Theorem 2. *There exist reductions from the Gap-DBS problem to the Gap-minmax problem as follows.*

1. *For every $0 < \eta < 0.1$ and $0 < c_1 < c_2$ satisfying $\frac{c_2}{c_1} > \frac{2\ln(1/\eta)}{(1-\eta)\eta^2}$ there is a randomized many-one reduction from the Gap-DBS problem to the Gap-minmax problem with parameters $(\eta, \xi - \eta/5)$.*

2. *For every $0 < \eta < 0.1$ and $0 < c_1 < c_2$ satisfying $\frac{c_2}{c_1} > 1/\eta$ there is a deterministic Turing reduction from the Gap-DBS problem to to the Gap-minmax problem with parameters $(\eta, \xi - \eta/5)$.*

We prove the two parts of this theorem as two separate theorems, stated as Theorem 9 and Theorem 15. We note that, interestingly, the constant ξ has previously turned up as the additive error $\xi + \delta$, for arbitrary $\delta > 0$, obtained by an approximation algorithm for computing ϵ-Nash equilibria [9].

One can view the second reduction in Theorem 2 as a derandomization of the first reduction in Theorem 2. However, this derandomization comes at the cost of turning the many-one reduction into a Turing reduction. On the other hand the required ratio between c_1 and c_2 is actually much smaller.

We will use the planted clique problem to argue that the Gap-DBS problem is hard for certain settings of parameters (c_1, c_2, η). For this we use similar arguments as in [13, 19]. Given a graph H that is an input to the planted clique detection problem, we let A be the adjacency matrix of H and let G be the bipartite graph that also has A as adjacency matrix. We wish to have the following property: If H was chosen from $G_{n,p,k}$, then with high probability G belongs to case (i) of the Gap-DBS problem, and if H was instead chosen from $G_{n,p}$ then with high probability G belongs to case (ii) of the Gap-DBS problem. This can indeed be obtained with an appropriate assumption about the clique detection problem. We have the following statement.

Proposition 3. *For any $\eta > 0$ there exist $p > 0$ and $c_1 > 0$ such that for $k = c_2 \ln n$, with $c_2 > c_1$, Gap-DBS with parameters (c_1, c_2, η) is as hard as the hidden clique detection problem for $G_{n,p,k}$.*

To prove this proposition we need the following basic lemma, whose proof is presented in Appendix A.

Lemma 4. *Let $\eta > 0$ be arbitrary. Then there exists a choice of $p > 0$ and $c_1 > 0$ such that with high probability a graph $G = (V, E)$ chosen from $G_{n,p}$ satisfies the following: Let A be the adjacency matrix of G, let $S, T \subseteq V$ be of size $|S| = |T| = c_1 \ln n$. Then $u_S^\intercal A u_T \leq \eta$, where u_S and u_T are probability vectors uniform on S and T.*

With this the proof is immediate:

Proof of Proposition 3. We choose $p > 0$ and $c_1 > 0$ according to Lemma 4. This ensures that with high probability graphs from $G_{n,p}$ (when considered as the bipartite graphs with the same adjacency matrix) satisfy case (ii) of Gap-DBS. Also, the choice of $k = c_2 \ln n$ automatically ensures that graphs from $G_{n,p,k}$ (as bipartite graphs) satisfy case (i) of Gap-DBS, since if S is the set of k nodes where the clique is placed, we have $\mathrm{d}(S, S) = 1 - 1/k$. $\square$

The choice of $c_2 > c_1$ of interest for us will be dictated by the choice of reduction we wish to use from Theorem 2, and in turn dictates the precise hardness assumption for the planted clique detection problem needed. However we find it natural to assume that the planted clique detection problem is hard for $G_{n,p,k}$ for any $p > 0$ and any $k = c_2 \ln n$, $c_2 > -2/\ln p$, i.e., with k significantly greater than the largest clique in $G_{n,p}$ (with high probability). Thus our results can be stated as follows.

Theorem 5. *For every $\epsilon > 0$, there is no randomized polynomial time algorithm that with high probability approximates the minmax value of a $n \times n \times n$ game with payoffs between 0 and 1 within an additive error $\xi/2 - \epsilon$, unless there exist $p > 0$ and $c_2 > -2/\ln p$ and a randomized polynomial time algorithm that solves the planted clique detection problem for $G_{n,p,k}$ with high probability, for $k = c_2 \ln n$.*

Theorem 6. *For every $\epsilon > 0$, there is no polynomial time algorithm approximating the minmax value of a $n \times n \times n$ game with payoffs between 0 and 1 within an additive error $\xi/2 - \epsilon$, unless there exist $0 < c_1 < c_2$ satisfying $c_2 > c_1/\eta$ and a (deterministic) polynomial time algorithm that solves the Gap-DBS problem with parameters (η, c_1, c_2), for $\eta = 5/3\epsilon$.*

1.1.1. Further Results

Our technique for derandomization used for obtaining the second part of Theorem 2 (and the restatement of Theorem 6) is general enough to apply to the related work about ϵ-Nash equilibria that is to be described in the next section. As an example of this, we present in Section 3.1 a derandomization of a result by Minder and Vilenchik [19]. For this we introduce another gap problem denoted Gap-ANE, reflecting the problem of approximating the maximum social welfare of a ϵ-Nash equilibrium. We then give a (deterministic) polynomial time Turing reduction from Gap-DBS to Gap-ANE.

Given that our result on the minmax value as well these results about ϵ-Nash equilibria are in fact based on the same hardness assumption, it is natural to ask if one can compare the hardness of these problems. We are not able to give an answer to this, but we are able to give some *evidence* that computing the minmax value in three player games is at least as hard as finding ϵ-Nash equilibria of high social welfare. We present this in Section 3.2.

1.2. Related Work

The problem of computing a Nash equilibrium in a bimatrix is PPAD complete. However, there are many different properties such that asking for a Nash equilibrium that satisfies the property is an NP hard problem [7, 11]. In particular it is NP hard to compute a Nash equilibrium maximizing the social welfare, i.e. maximizing the sum of the two players payoffs.

Hazan and Krauthgamer [13], motivated by the question of whether there is a PTAS for computing ϵ-Nash equilibria, considered an "ϵ-Nash" variant of the problem of maximizing social welfare, namely that of computing an ϵ-Nash equilibrium whose social welfare is no less than the maximal social welfare achievable by a Nash equilibrium, minus ϵ. In order to describe all the results in the following, say that an ϵ-Nash equilibrium is *δ-good* if its social welfare is no less than the maximal social welfare achievable by a Nash equilibrium, minus δ.

Remark 7. For the notion introduced by Hazan and Krauthgamer, Minder and Vilkenchik [19] use the terminology "ϵ-best ϵ-Nash equilibrium". However we feel this is somewhat of a misnomer, since the social welfare is compared to the largest achievable by a Nash equilibrium rather than an ϵ-Nash equilibrium. Indeed, a simple example[2] shows that for any $\epsilon > 0$ one may have a game where the (unique) Nash equilibrium has social welfare ϵ, but there exist an ϵ-Nash equilibrium of social welfare 1. For this reason we will instead call it "ϵ-good". In fact, we find it useful to generalize the notion to an ϵ-Nash equilibrium being called *δ-good* as defined above.

Hazan and Krauthgamer gave a randomized polynomial time reduction from the planted clique problem to the problem of computing an ϵ-good ϵ-Nash equilibrium. More precisely, they show there are constants $\epsilon, c > 0$ such that if there is a polynomial time algorithm that computes in a two-player bimatrix game an ϵ-good ϵ-Nash equilibrium, then there is a randomized polynomial time algorithm that solves the planted clique problem in $G_{n,1/2}$ for $k = c \log_2 n$ with high probability.

This result was sharpened by Minder and Vilenchik [19], making the constant c smaller. In particular they obtain $c = 3 + \delta$, for arbitrary $\delta > 0$ (here $\delta > 0$ dictates an upper bound on ϵ), and for the similar problem of detecting a planted clique they obtain $c = 2 + \delta$. Essentially the goal of Minder and Vilenchik was the opposite of ours. Namely, viewing their result as arguing for hardness, their goal was to obtain an assumption as weak as possible, while maintaining a nontrivial conclusion.

[2]Consider just the bimatrix game given by the two 1×2 matrices $\begin{bmatrix} 1 & 0 \end{bmatrix}$ for the row player and $\begin{bmatrix} 0 & \epsilon \end{bmatrix}$ for the column player.

Austrin et al. [3] considered the other goal of obtaining strong hardness conclusions for computing δ-good ϵ-Nash equilibria (as well as ϵ-Nash versions of computing second equilibria and small support equilibria, and approximating pure Bayes Nash equilibria), assuming hardness for the planted clique problem. For this reason this work is the most relevant to use for comparing with our results. With the goal of obtaining strong hardness conclusions for computing δ-good ϵ-Nash equilibria in mind, one now needs to consider both of the parameters, ϵ and δ, and their relationship. Austrin et al. consider the extreme cases for both of these parameters individually and obtain the following results.

Theorem 8 (Austrin et al.). *1. For any $\eta > 0$ there exists $\delta = \Omega(\eta^2)$ such that computing a δ-good ϵ-Nash equilibrium is as hard as the planted clique problem, for $\epsilon = 1/2 - \eta$.*

2. For any $\eta > 0$ there exists $\epsilon = \Omega(\eta^2)$ such that computing a δ-good ϵ-Nash equilibrium is as hard as the planted clique problem, for $\delta = 2 - \eta$.

Furthermore Austrin et al. give a simple polynomial time algorithm that computes a $\frac{1}{2}$-Nash equilibrium with social welfare at least as large as any Nash equilibrium, showing that the first part of Theorem 8 is tight. Clearly the second part is tight as well. On the other hand it appears that the tightness of these results were possible due to the focus on a single parameter at a time, and the exact trade-off possible between these two parameters still seems unclear.[3]

1.3. Techniques and comparison with related work

While Hazan and Krauthgamer [13] consider the specific setting of computing ϵ-good ϵ-Nash equilibria, we can describe their approach in a general way for an unspecified computational problem. Let A be an $n \times n$ matrix with entries belonging to the interval $[0, 1]$. Let $0 < a < b < c < 1$ and $\gamma > 0$ be constants to be discussed later. From this matrix a particular instance for the problem in hand is constructed. By specifics of the problem considered and by properties of the instance constructed, a solution of the instance gives rise to probability distributions x, y satisfying the following two properties:

1. $x^\mathsf{T} A y \geq 1 - a$.
2. For any subset $S \subseteq [n]$ with $|S| \leq c_1 \ln n$ we have $\Pr_x[S] \leq \gamma$ and $\Pr_y[S] \leq \gamma$.

Define $T = \{i \mid x^\mathsf{T} A e_i \geq 1 - b\}$. Using Markov's inequality one gets $\Pr_y[T] \geq 1 - a/b$. Assuming $a/b < 1 - \gamma$ one concludes $|T| \geq c_1 \ln n$. Further, define $S = \{i \mid {e_i}^\mathsf{T} A u_T \geq 1 - c\}$, where u_T is the uniform distribution on the set T. Again using Markov's inequality one gets $\Pr_x[S] \geq 1 - b/c$, and assuming $b/c < 1 - \gamma$ one concludes $|S| \geq c_1 \ln n$. The conclusion of this argument is that assuming the inequalities $a/b < 1 - \gamma$

[3] While the statements of Theorem 8 are given using asymptotic notation, the proofs provide concrete (albeit not particularly optimized) constants. For instance the proof of the first part gives $\delta = 1/288$ for $\epsilon = 1/4$, and the proof of the second part gives $\epsilon = 1/288$ for $\delta = 3/2$.

and $b/c < 1-\gamma$ we may find sets S and T of size at least $c_1 \ln n$ such that $u_S^\mathsf{T} A u_T \geq 1-c$. Now if A is the adjacency matrix of a random graph G picked according to $G_{n,p}$ and if $1-c > p$ and c_1 is a sufficently large constant, such sets S and T would not exist with high probability. Alternatively if G was chosen at random from $G_{n,p,k}$ where $k = c_2 \ln n$ for a sufficently large constant $c_2 > c_1$, such sets S and T can be used to recover the hidden k-clique with high probability in polynomial time.

Consider now the setting of 2-player games [13, 19, 3]. The idea above is implemented by taking the adjacency matrix A as the payoff matrix for both players, and then augmenting these with additional pure strategies for both players. The new payoffs are constructed in such a way that in any ϵ-Nash equilibrium of sufficiently large social welfare, both players must place most of their probability mass on the pure strategies corresponding to A. The additional pure strategies are then meant to enable a player the possibility of achieving a payoff at least ϵ larger than the current by switching to an appropriate pure strategy if his opponent places probability more than probability mass γ on a set of $c_1 \ln n$ vertices. The probability distributions x and y above are obtained by first restricting the support of the strategies of the two players to those pure strategies corresponding to A and then normalizing.

In our setting we let the matrix A with 0 entries and 1 entries exchanged give the payoffs to player 1, corresponding to a single of his pure strategies. We may also think of the matrix A as giving penalties for player 1 rather than rewards. The pure strategies of player 2 and player 3 correspond exactly to rows and columns of A, and thus the probability distributions x and y directly corresponds to strategies (of player 2 and player 3) in the game we construct. This fact makes our reduction and analysis technically simpler compared to the 2-player setting. To implement the idea above we provide a number of new strategies for player 1, such that if either player 2 or player 3 places probability mass γ on a set of $c_1 \ln n$ vertices, switching to an appropriate pure strategy may ensure him a payoff of γ. Note here that in our setting we have only one parameter available, namely the minmax value. Because of this fact we must perform the analysis of our reduction in a tight way in order to obtain best possible results.

2. The Reductions

We collect the utilities for Player 1 in matrices, one for each pure strategy. Thus we define $n_2 \times n_3$ matrices $A^{(1)}, \ldots, A^{(n_1)}$ by $a_{j,k}^{(i)} = u_1(i,j,k)$. In this notation, if Player 1 plays the pure strategy i and Player 2 and Player 3 play by mixed strategies x and y, the expected payoff to Player 1 is given by $x^\mathsf{T} A^{(i)} y$.

2.1. The randomized reduction

In this section we present a randomized reduction from approximate planted clique to minmax-value in three player games. To be precise, we prove the following result:

Theorem 9. *Let $0 < \eta < 0.1$ and $0 < c_1 < c_2$ and such that $\frac{c_2}{c_1} > \frac{2\ln(1/\eta)}{(1-\eta)\eta^2}$. Then there is a randomized polynomial time many-one reduction which, given as input the* *adjacency matrix* $A \in \{0,1\}^{n\times n}$ *of a bipartite graph* G*, outputs a three-player game* G_A *such that with high probability*

- *if there are subsets* $S,T \subseteq [n]$ *of size at least* $c_2 \ln n$ *such that* $\mathrm{d}(S,T) \geq 1-\eta$*, then* $\mathrm{minmax}_1 G_A \leq \eta$*.*
- *if* $\mathrm{d}(S,T) < \eta$ *for every* $S,T \subseteq [n]$ *of size at least* $c_1 \ln n$*, then* $\mathrm{minmax}_1 G_A > \xi - \frac{\eta}{5}$*.*

We will need the following lemma.

Lemma 10. *Let* $0 < \delta < 1$*, and* $k_1 = c_1 \ln n$*,* $k_2 = c_2 \ln n$*, where* $0 < c_1 < c_2$ *satisfy* $c_2 > \frac{2\ln(1/\delta)}{(1-\delta)\delta^2} \cdot c_1$*. Let* $D \subseteq [n]$ *be a fixed subset of size* $|D| = k_2$*. Then there is a constant* c *such that if we we choose at random* $m = n^c$ *subsets* $S_1, \ldots, S_m \subseteq [n]$*, by letting* $j \in S_i$ *with probability* $1-\delta$*, independently for every* i *and* j*, with probability at least* $1 - n^{-\Omega(1)}$ *the sets satisfy the following properties.*

(a) For all i*,* $|S_i \cap D| \geq (1-\delta)^2 k_2$*.*

(b) For every set $S \subseteq [n]$ *of size* $|S| = k_1$*, there exists* i *such that* $S_i \cap S = \emptyset$*.*

Proof. By assumption we can pick c such that

$$c_1 \cdot \ln(1/\delta) < c < \frac{(1-\delta)\delta^2}{2} \cdot c_2 \ .$$

We first prove property (a) holds with the claimed probability. We have $E[|S_i \cap D|] = (1-\delta)k_2$. By the Chernoff bound for the lower tail we have

$$\Pr[|S_i \cap D| < (1-\delta)^2 k_2] < \exp(-(1-\delta)\delta^2 k_2/2) = n^{-\frac{(1-\delta)\delta^2}{2}\cdot c_2}$$

Hence

$$\Pr[\exists i : |S_i \cap D| < (1-\delta)^2 k_2] < m \cdot n^{-\frac{(1-\delta)\delta^2}{2}\cdot c_2} = n^{-\Omega(1)} \ .$$

We next prove property (b) also holds with the claimed probability. Consider $S \subseteq [n]$ of size $|S| = k_1$. Then $\Pr[S_i \cap S \neq \emptyset] = 1 - \delta^{k_1}$, and

$$\Pr[\forall i : S_i \cap S \neq \emptyset] = (1-\delta^{k_1})^m < \exp(-\delta^{k_1} m) = \exp(-n^{c - c_1 \ln(1/\delta)})$$

Hence

$$\begin{aligned}\Pr[\exists S \subseteq [n], |S| = k_1 : \forall i : S_i \cap S \neq \emptyset] &< \binom{n}{k_1} \exp(-\delta^{k_1} m) \\ &\leq \exp(c_1 \ln^2(n) - n^{c - c_1 \ln(1/\delta)}) < \exp(-n^{\Omega(1)}) \ .\end{aligned}$$

□

Proof of Thm. 9. We use Lemma 10 with c_1 and c_2 as in the problem description and $\delta = 1 - \sqrt{1-\eta} = \eta/2 + O(\eta^2)$. Let m be as in the lemma. The reduction first picks $2m$ subsets $S_1^{(\mathrm{r})}, \ldots, S_m^{(\mathrm{r})}, S_1^{(\mathrm{c})}, \ldots, S_m^{(\mathrm{c})}$ at random as in the lemma. It then outputs a 3-player game G_A as follows:

- Players 2 and 3 have n strategies each.
- Player 1 has $2m+1$ strategies given by matrices B, $R^{(1)},\dots,R^{(m)}$, and $S^{(1)},\dots,S^{(m)}$. The matrix B is defined as $B = 1 - A$, and $R^{(k)}$ and $C^{(k)}$ for $k = 1,\dots,m$, are given by

$$(R^{(k)})_{ij} = \begin{cases} 1 & \text{if } i \notin S_k^{\mathrm{r}} \\ 0 & \text{if } i \in S_k^{\mathrm{r}} \end{cases} \quad \text{and} \quad (C^{(k)})_{ij} = \begin{cases} 1 & \text{if } j \notin S_k^{\mathrm{c}} \\ 0 & \text{if } j \in S_k^{\mathrm{c}} \end{cases}$$

We claim that this game satisfies our assumptions.

For the first part, let $S, T \subseteq [n]$ be sets of size at least $c_2 \log n$ such that $\mathrm{d}(S,T) \geq 1-\eta$. By choosing appropriate subsets, we may assume that, in fact, $|S| = |T| = c_2 \log n$. Furthermore, by Lemma 10, with high probability $|S_i^{\mathrm{r}} \cap S| \geq (1-\delta)^2 c_2 \ln n$ and $|S_i^{\mathrm{c}} \cap T| \geq (1-\delta)^2 c_2 \ln n$. Thus if players 2 and 3 play strategies u_S and u_T, respectively, by playing any of the strategies corresponding to $R^{(k)}$ and $C^{(k)}$ player 1 will receive payoff at most $1-(1-\delta)^2 = \delta(2-\delta) = (1-\sqrt{1-\eta})(1+\sqrt{1-\eta}) = \eta$, while by playing the strategy corresponding to B will give player 1 payoff $1 - \mathrm{d}(S,T) < \eta$.

For the second part, we assume to the contrary that G has density $\mathrm{d}(S,T) < \eta$ for all sets S, T of size at least $c_1 \ln n$, but $\min\max G_A \leq a$. Let (σ_2, σ_3) be an optimal strategy profile, i.e., $\max\left\{\sigma_2^\intercal B\sigma_3, \sigma_2^\intercal R^{(k)}\sigma_3, \sigma_2^\intercal C^{(k)}\sigma_3\right\} \leq a$. We first show that on any support of size at most k_1 each of σ_2 and σ_3 places probability at most a: Suppose $S \subseteq [n]$ and $|S| \leq k_1$ with $\Pr_{\sigma_2}[S] = p$. Then by switching to an appropriate set action corresponding to $R^{(k)}$, player 1 might increase his payoff to at least p. Thus $p \leq a$. The proof for σ_3 is the same, replacing $R^{(k)}$ with $C^{(k)}$. We set, with foresight, $a = \xi - \frac{\eta}{5}$, $b = 1 - \xi - \frac{\eta}{2}$, and $c = 1 - \eta$. Direct calculations show that for $0 < \eta < 0.1$, these values satisfy

$$a < b < c < 1 \qquad (1-a)b > a \qquad \text{and} \qquad (1-a)c > b\ . \tag{2}$$

We show that there exist sets S and T of size at least $c_1 \ln n$ such that $u_S^\intercal A u_T \geq 1 - c$: Define $T = \{i \mid \sigma_2^\intercal B e_i \leq b\}$, and let $p = \Pr_{\sigma_3}[T]$. Then $a \geq \sigma_2^\intercal B \sigma_3 \geq (1-p)b$, and therefore $(1-p)b < a$, which means $1 - p < a/b$. But we have $1 - a > a/b$, which then implies $p > a$, and therefore $|T| \geq c_1 \ln n$ as argued above. Furthermore, by definition of T we have $\sigma_2^\intercal B u_T \leq b$. Next, define $S = \{i \mid {e_i}^\intercal B u_T \leq c\}$, and let $p = \Pr_{\sigma_2}[S]$. Similarly to before we then have $b \geq \sigma_2^\intercal B u_T \geq (1-p)c$ which means $(1-p)c < b$, and thus $1 - p < b/c$. But we have $1 - a > b/c$, which then implies $p > a$, and again we obtain that $|S| \geq c_1 \ln n$. Furthermore, by definition of S and $B = 1 - A$ we have $u_S^\intercal A u_T \geq 1 - c = \eta$. ☐

Remark 11. We see in the above proof our reason for considering the planted clique problem $G_{n,p,k}$ in the setting of having $p > 0$ a small constant. In order to make a as large as possible and still satisfy the inequalities (2), $c = 1 - \eta$ should be made as large as possible. Since the conclusion $u_S^\intercal A u_T \geq \eta$ should not hold with high probability if A is the adjacency matrix of a graph chosen at random from $G_{n,p}$ this means we need $\eta > p$. The argument of the proof could still be performed in the setting of $p = \frac{1}{2}$, but that would then require us to have $a < 2 - \sqrt{3} \approx 0.268$, instead of having $a < \xi \approx 0.382$ as above.

2.1.1. Tightness of the analysis

We will observe here that the above analysis is tight, namely that in the case when $\mathrm{d}(S,T) < \eta$ for every $S,T \subseteq [n]$ of size at least $c_1 \ln n$, it is not possible to prove a lower bound on the minmax value better than ξ in the game constructed.

Proposition 12. *There is a $n \times n$ matrix A such that the game G_A given by the reduction of Theorem 9 satisfies the following:*

- *For all $S,T \subseteq [n]$ of size at least $c_1 \ln n$ we have $\mathrm{d}(S,T) < 2/c_1 \ln n$.*
- *With probability $1 - 2^{\Omega(-\eta^2/n)}$ we have $\mathrm{minmax}_1 G_A \leq \xi + \eta/2 + O(\eta^2)$.*

Proof. Define A to be the $n \times n$ matrix where $(A)_{ij} = 1$ if either $i = 1$ or $j = 1$, and $(A)_{ij} = 0$ otherwise. The first claim about the densities is then obvious.

For the second claim, let $\epsilon = \delta/3$ and define $\sigma_2 = \sigma_3 = (\xi - \epsilon, \frac{1-\xi+\epsilon}{n-1}, \ldots, \frac{1-\xi+\epsilon}{n-1})^\mathsf{T}$. In other words σ_2 and σ_3 place probability $\xi - \epsilon$ on the first strategy, and distributes the remaining probability mass uniformly on the remaining strategies.

By definition, the pair (σ_2, σ_3) establishes an upper bound on the minmax value. First

$$\sigma_2^\mathsf{T} B \sigma_3 = (1-\xi+\epsilon)^2 = (1-\xi)^2 + 2\epsilon(1-\xi) + \epsilon^2 \leq \xi + 3\epsilon = \xi + \delta \ .$$

Next,

$$E[\sigma_2^\mathsf{T} R^{(k)} \sigma_3] = (1-\delta)(\xi-\epsilon) + \delta(1-\xi+\epsilon) \leq \xi - \epsilon + \delta = \xi + 2\delta/3 \ ,$$

and by the Chernoff-Hoeffding bound we then have

$$\Pr[\sigma_2^\mathsf{T} R^{(k)} \sigma_3 > \xi + \delta] \leq \exp(-2(\delta/3)^2/n) \ .$$

We have the same bound for $C^{(k)}$ instead of $R^{(k)}$, and hence taking a union bound over all these $2m$ matrices we have $\sigma_2^\mathsf{T} R^{(k)} \sigma_3 \leq \xi + \delta$ and $\sigma_2^\mathsf{T} C^{(k)} \sigma_3 \leq \xi + \delta$, except with probability $2^{-\Omega(\delta^2/n)}$. ☐

2.2. Derandomization

In this section we derandomize our result in Theorem 9, at the price of turning our many-one reduction into a Turing reduction.

Recall that randomness was needed by our reduction for the construction of the sets $S_i^{(r)}$ and $S_i^{(c)}$. We now show how these sets can be constructed explicitly, giving a derandomized analogue of Lemma 10:

Lemma 13. *Let $0 < k_1 < k_2 < n \in \mathbb{N}$. Then there are families $A^{(1)}, \ldots, A^{(r)}$ of subsets of $[n]$ such that*

- *there are $r = 2^{O(k_2)} \log n$ families, and each family is of size $s = \binom{k_2}{k_1}$,*

- *for every set $M \subseteq [n]$ of size k_2, there is an index $j \in [r]$ such that $\left|A_i^{(j)} \cap M\right| = k_2 - k_1$, for all $i \in [s]$ and*
- *for every set $M \subseteq [n]$ of size k_1 and every $j \in [r]$, there is an index $i \in [s]$ such that $A_i^{(j)} \cap M = \emptyset$.*

These sets can be constructed in time polynomial in n and r.

Proof. In [2], Alon et al. gave a construction of a family $H = \{f_1, \ldots, f_r\}$ of *perfect hash functions* from $[n]$ to $[k_2]$. This means

- each f_j is a function from $[n]$ to $[k_2]$ and
- for each $M \subseteq [n]$ of size k_2, at least one of the f_j is injective on M.

Moreover, $r = 2^{O(k_2)} \log n$ and the functions can be constructed in time polynomial in n and r.

Let $s = \binom{k_2}{k_1} \leq 2^{k_2}$ and let $M_1, \ldots, M_s$ be an enumeration of the subsets of $[k_2]$ of size k_1. Define $A_i^{(j)} := \{x \in [n] \mid f_j(x) \notin M_i\}$. These subsets meet the size restrictions claimed in the lemma and are readily seen to be constructable in time $\mathrm{poly}(n, r)$.

Now, let $M \subseteq [n]$ be of size k_2, and suppose f_j is injective on M. Then $A_i^{(j)} \cap M = \{x \in M \mid f_j(x) \notin M_i\}$, and because f_j is a bijection between M and $[k_2]$, this set has size $k_2 - k_1$ for all $i \in [s]$.

Furthermore, if $M \subseteq [n]$ is of size k_1, then $|f_j(M)| \leq k_1$ for all $j \in [r]$. Thus for each j there is an i such that $f_j(M) \subseteq M_i$, which implies $A_i^{(j)} \cap M = \emptyset$. $\square$

Corollary 14. *If $k_2 = O(\log n)$ then both r and s are polynomial in n, and the families of subsets can be constructed in time polynomial in n.*

Our derandomized reduction now looks as follows:

Theorem 15. *For $0 < \eta < 0.1$ and $0 < c_1 < c_2$ and such that $\frac{c_2}{c_1} > \frac{1}{\eta}$, there is a polynomial-time Turing reduction from Gap-DBS to Gap-Minmax with a gap $(\eta, \xi - \eta/5)$.*

Proof. The reduction works as in the randomized case, the main difference being that instead of picking sets $S_i^{(r)}$ and $S_i^{(c)}$ at random, we construct (polynomially many) set families $A^{(1)}, \ldots, A^{(r)}$ using the construction in Lemma 13 with $k_{1/2} = c_{1/2} \ln n$. We then use each pair of such families to construct a game $G_A^{(j_1,j_2)}$ as in the proof of Theorem 9; using the family $A^{(j_1)}$ for the row strategies and $A^{(j_2)}$ for the column strategies. We show that

- if $\mathrm{d}(S,T) \geq 1-\eta$ for some sets S, T of size at least $c_2 \ln n$, then $\mathrm{minmax}_1 G_A^{(j_1,j_2)} \leq \eta$, for *some* j_1 and j_2, and
- if $\mathrm{d}(S,T) \leq \eta$ for all sets S, T of size at least $c_1 \ln n$, then $\mathrm{minmax}_1 G_A^{(j_1,j_2)} \geq \xi - \eta/5$, for *all* j_1, j_2.

The proof works as in the randomized case: For the first part, we note that by Lemma 13, for some j_1, j_2 and all i we have $\left|A_i^{(j_1)} \cap S\right| = k_2 - k_1 \geq (1-\eta)k_2$ and $\left|A_i^{(j_2)} \cap T\right| = k_2 - k_1 \geq (1-\eta)k_2$, and therefore $\mathrm{minmax}_1 G_A^{(j_1,j_2)} \leq \eta$ in this case. The second part is unchanged from the randomized case. □

3. Further Reductions

3.1. A Turing-Reduction from Gap-DBS to approximate Nash-Equilibrium

In [19], Minder and Vilenchik show how a polynomial time algorithm computing an ϵ-good ϵ-equilibrium in a two player game can be used to obtain a *randomised* algorithm which distinguishes, with high probability, between a random graph drawn from the distribution $G_{n,1/2}$ from a graph with a clique of size at least $(2 + 28\epsilon^{1/8})\log n$. Specifically, they obtain a randomised algorithm which, on input a random graph from $G_{n,1/2}$, rejects with high probability (over the random choices of the algorithm *and* over the input), and which will accept any particular graph containing a clique of size $(2 + 28\epsilon^{1/8})\log n$ with high probability over its internal random choices.

The distinguishing algorithm which Minder and Valenchik construct actually rejects with high probability any graph which does not contain a dense subgraph of size slightly larger than $2\log n$ but still smaller than the planted clique. With high probability, the random graph $G_{n,1/2}$ does not contain such a subgraph. In turns out that we may also introduce some slack in the planted clique and just plant a dense subgraph of appropriate size. Finally, to avoid reducing a decision problem to a search problem, we introduce the following promise-problem for Nash Equilibria:

Gap Approximate Nash Equilibrium (Gap-ANE)	
Input:	a bimatrix game represented by two $(n \times n)$-payoff-matrices with payoffs in $[-2, 2]$
Promise:	Either (i) there is a δ-approximate Nash Equilibrium with average payoff $\geq 1 - \delta$ or (ii) there is no δ-approximate Nash Equilibrium with average payoff $\geq 1 - 2\delta$.
Problem:	Decide which of these is the case

With these preliminaries, we can now state our derandomisation:

Theorem 16. *For sufficiently small $\eta > 0$ there are $c_1 < c_2$ and $\delta > 0$ such that there is a (deterministic) polynomial time Turing reduction from Gap-DBS to Gap-ANE* *with these parameters. In particular, one possible choice of c_1 and c_2 is*

$$c_1 := 2 + 6\sqrt{2(2\eta)^{1/4}} \qquad \text{and} \qquad c_2 := 2 + 7\sqrt{2(2\eta)^{1/4}}.$$

The payoff matrices of the games which our reduction produces contain only three distinct values, which only depend on η (and not on n). Therefore we may assume these values to be specified up to some arbitrary accuracy.

Proof. Let A be the adjacency matrix of a bipartite graph satisfying the promise of Gap-DBS. We proceed as in section 3 of [19]. Let

$$\begin{aligned}
\delta &:= 2\eta \\
\beta &:= 2\delta^{1/4}, \\
k_1 &:= (2 + 6\sqrt{\beta}) \log n, \\
k_2 &:= \left\lceil (2 + 7\sqrt{\beta}) \right\rceil \log n.
\end{aligned}$$

By lemma 13 there are families $A^{(1)}, \ldots, A^{(r)}$ of subsets of $[n]$, with s sets in each family, such that

1. For every subset $S \subseteq [n]$ of size k_2, there is an index $i \in [r]$ such that

$$\left| S \cap A_j^{(i)} \right| = k_2 - k_1$$

   for all $j \in [s]$,

2. For every subset $S \subset [n]$ of size k_1 and every $i \in [r]$ there is a $j \in [s]$ such that

$$S \cap A_j^{(i)} = \emptyset.$$

Furthermore, $r, s \leq n^{O(1)}$ and these set families can be constructed deterministically in polynomial time. Our reduction outputs one instance $\mathcal{G}_{i,j}$ of Gap-ANE for each pair of indices $i, j \in [r]$ such that

- if there are no sets S and T with $|S|, |T| \geq k_1$ such that

$$d(S,T) > 1 - \beta = 1 - 2(2\eta)^{1/4}$$

  (i.e., A has no dense subgraph), then none of the games $\mathcal{G}_{i,j}$ has a δ-approximate Nash Equilibrium of value $\geq 1 - 2\delta$ and

- if there are sets S and T with $|S|, |T| \geq k_2$ such that

$$d(S,T) > 1 - \eta$$

  (i.e., A does have a dense subgraph), then at least one of the games $\mathcal{G}_{i,j}$ has a δ-approximate Nash Equilibrium of value $\geq 1 - \delta$.

Note that, in particular, in the first case all games $\mathcal{G}_{i,j}$ satisfy the promise of Gap-ANE. In the second case, this is not necessarily true, but there is guaranteed to be at least one pair (i,j) for which $\mathcal{G}_{i,j}$ is a yes-instance satisfying the promise.

We define the payoff-matrices for the row and column player in the game $\mathcal{G}_{i,j}$ as follows:

$$R_{i,j} := \begin{pmatrix} A & -B_j^{\mathrm{T}} \\ B_i & 0 \end{pmatrix} \qquad C_{i,j} := \begin{pmatrix} A & B_j^{\mathrm{T}} \\ -B_i & 0 \end{pmatrix}.$$

is zero-sum outside the upper left block corresponding to the matrix A. Each matrix B_i is an $(s \times n)$-matrix with entries

$$(B_i)_{a,b} := \begin{cases} 0 & \text{if } b \in A_a^{(i)} \\ 1 + 18\beta & \text{if } b \notin A_a^{(i)} \end{cases}$$

First, assume there are subsets S, T of size k_2 such that $d(S,T) > 1 - \eta$. Then playing the uniform strategies on S and T will give each of the players a payoff of at least $1 - \eta$. Furthermore, for at least one combination of indices i and j this is an η-approximate Nash Equilibrium. In fact, there is a tuple (i,j) such that no player can increase his profit to more than

$$\frac{k_1}{k_2}(1 + 18\beta),$$

and, setting $\alpha := \sqrt{\beta}$, this is at most one if

$$\begin{aligned} & (1 + 18\alpha^2)(2 + 6\alpha) \leq 2 + 7\alpha \\ \Leftrightarrow \qquad & -\alpha + O(\alpha^2) < 0, \end{aligned}$$

which is the case for small enough values of $\alpha = \sqrt{2}(2\eta)^{1/8}$. This proves the first part of our claim.

For the second part we show that if there is a δ-equilibrium with payoff least $1 - 2\delta$, then there are subsets S and T of size at least k_1 such that $d(S,T) > \eta$. Our arguments are the same as those in Propositions 3–6 of [19].

We write a strategy x as $x_A + x_{\bar{A}}$, where x_A is the part of the strategy played on A and $x_{\bar{A}}$ is the rest. In particular, $\|x_A\|_1$ is the probability that the player choses a strategy from A.

Claim 17. If (x,y) is a pair of strategies with average payoff $\geq 1 - 2\delta$, then each of the players puts probability at least $1 - 2\delta$ on A.

Proof. We have

$$\begin{aligned} 1 - 2\delta &\leq \frac{1}{2} x^{\mathsf{T}} (R + C) y \\ &= x^{\mathsf{T}} \begin{pmatrix} A & 0 \\ 0 & 0 \end{pmatrix} y \\ &\leq x^{\mathsf{T}} \begin{pmatrix} J & 0 \\ 0 & 0 \end{pmatrix} y \\ &= \|x_A\|_1 \cdot \|y_A\|_1 , \end{aligned}$$

and the claim follows because $\|x_A\|_1, \|y_A\|_1 \in [0,1]$. $\square$

Claim 18. If (x,y) is a δ-approximate Nash Equilibrium with average payoff $\geq 1-2\delta$, then

$$\tilde{x} := \frac{1}{\|x_A\|_1} x_A \qquad \text{and} \qquad \tilde{y} := \frac{1}{\|y_A\|_1} y_A$$

define a 16δ-approximate Nash Equilibrium with average payoff $\geq 1-2\delta$.

Proof. For the average payoff, we note that

$$\begin{aligned}
1 - 2\delta &\leq \frac{1}{2} x^{\mathsf{T}}(R+C)y \\
&= \frac{1}{2} x_A^{\mathsf{T}}(R+C)y_A \\
&\leq \frac{1}{2\|x_A\|_1 \|y_A\|_1} x_A^{\mathsf{T}}(R+C)y_A \\
&= \frac{1}{2}\tilde{x}^{\mathsf{T}}(R+C)\tilde{y}.
\end{aligned}$$

We now show that $(\tilde{x},\tilde{y})$ is a 16δ-approximate Nash Equilibrium. In fact for every pure strategy e_i,

$$\begin{aligned}
e_i^{\mathsf{T}} R\tilde{y} &= \frac{1}{\|y_A\|_1} e_i^{\mathsf{T}} R y_A \\
&\leq \frac{1}{\|y_A\|_1}\left(e_i^{\mathsf{T}} R y + 4\delta\right) && \text{because } \|y_{\bar{A}}\|_1 < 2 \\
&\leq \frac{1}{\|y_A\|_1}\left(x^{\mathsf{T}} R y + 5\delta\right) \\
&= \frac{1}{\|y_A\|_1}\left(\|x_A\|_1 \|y_A\|_1 \tilde{x}^{\mathsf{T}} R\tilde{y} + x_{\bar{A}}^{\mathsf{T}} B y_A - x_A^{\mathsf{T}} B^{\mathsf{T}} y_{\bar{A}} + 5\delta\right) \\
&\leq (1+2\delta)(\tilde{x}^{\mathsf{T}} R\tilde{y} + 13\delta) \\
&\leq \tilde{x}^{\mathsf{T}} R\tilde{y} + 16\delta
\end{aligned}$$

The proof for the column player is similar. $\square$

Claim 19. Let $\mathcal{G} = \mathcal{G}_{i,j}$ be any of the games generated by our reduction from the matrix A, and assume that $\beta \in [\delta, 1/18]$. Let (x,y) by a 16δ-approximate Nash Equilibrium of $\mathcal{G}$ played entirely on A, i.e., $x_{\bar{A}} = y_{\bar{A}} = 0$. Then if M is a set of rows such that $\Pr_x(M) \geq 1-\beta$, then $|M| > k_1$, and similarly for y.

Proof. Assume that M is a subset of rows such that $|M| \leq k_1$. Me may assume that M contains only rows correspdonding to strategies in A. By our construction of B there is a column of $C_{i,j}$ such that all entries of this column corresponding to rows in M have value $1+18\beta$. By defecting to this column, the column player will receive a payoff of at least

$$(1+18\beta)(1-\beta) = 1 + 17\beta - 18\beta^2 \geq 1 + 16\delta$$

by our assumption on β. Because the column player's expected payoff when playing y can not exceed 1, the pair (x, y) can not be a Nash Equilibrium. $\square$

Claim 20. Given a 16δ-approximate Nash Equilibrium (x, y) of average payoff $\geq 1-2\delta$ which is played entirely on A in one of the games $\mathcal{G}_{i,j}$ generated by our reduction, one can efficiently find two sets of vertices S and T such that

- $|S|, |T| > k_1$ and
- $d(S, T) > 1 - \beta$.

Proof. Since we assume that both players play entirely on A and each player's maximum payoff on A is 1, we note that each of the players must receive a payoff of at least $1 - 4\delta$. Let T be the set of actions on A for which the column player receives payoff at least $1 - 4\sqrt{\delta}$, i.e.,

$$T := \left\{ i \mid x^{\mathsf{T}} A e_i \geq 1 - 4\sqrt{\delta} \right\}.$$

Let $p := \Pr_y(T)$. Then

$$\begin{aligned} 1 - 4\delta &\leq x^{\mathsf{T}} A y \\ &\leq p + (1-p)\left(1 - 4\sqrt{\delta}\right) \\ &= 1 - 4\sqrt{\delta}(1-p), \end{aligned}$$

and $p \geq 1 - \sqrt{\delta}$. Let u_T be the uniform strategy for the column player on the columns in T, and define

$$S := \left\{ i \mid e_i^{\mathsf{T}} A u_T \geq 1 - \beta \right\}.$$

Let $q := \Pr_x(T)$. Then

$$\begin{aligned} 1 - 4\sqrt{\delta} &\leq x^{\mathsf{T}} A u_T \\ &\leq q + (1-q)(1-\beta) \\ &= 1 - \beta(1-q), \end{aligned}$$

and, recalling that $\beta = 2\delta^{1/4}$,

$$q \geq 1 - \frac{4\sqrt{\delta}}{\beta} = 1 - \beta.$$

In particular,

$$d(S, T) = u_S^{\mathsf{T}} A u_T \geq 1 - \beta$$

and $\Pr_x(S), \Pr_y(T) \geq 1 - \beta$, so $|S|, |T| > k_1$ by the previous claim. $\square$

This concludes the proof of Theorem 16. Note that the condition $\beta \geq \delta$ in Claim 19 is always satisfied by our choice of β, while the upper bound $\beta \geq 1/18$ is satisfied if $\eta < (1/36)^4/2 \approx 3 \cdot 10^{-7}$. $\square$

It is possible to derandomise the other algorithms in [13] and [19] using Lemma 13 as well, but we do not give details here.

3.2. A reduction from optimal NE to Minmax

The following reduction gives evidence to the fact that computing the minmax-value in three player games is at least as hard as finding ϵ-Nash equilibria with high average payoff (note that the statement involves both ϵ-Nash equilibria and 2ϵ-Nash equilibria).

Theorem 21. *There is a polynomial time reduction which, given payoff-matrices $R, C \in [0,1]^{m\times n}$ specifying a game $\mathcal{G}$ in which the players have m and n strategies respectively, and $\alpha \in [0,1]$, $\epsilon > 0$, outputs payoff matrices for player 1 in a three player game $\mathcal{H}$ such that:*

- *If $\mathcal{G}$ has an ϵ-Nash equilibrium with average payoff $> 1-\alpha$, then* $\text{minmax}_1\mathcal{H} \leq \alpha$.
- *If $\mathcal{G}$ has no 2ϵ-Nash equilibrium with average payoff $> 1-\alpha-\epsilon$, then* $\text{minmax}_1\mathcal{H} > \alpha + \epsilon$.

Proof. Player one has $m + n + 1$ strategies, player two has m strategies and player three has n strategies. We group player one's strategies into three categories:

1. one strategy called v which has payoff-matrix
$$1 - \frac{1}{2}(R + C)$$
2. for each $\tilde{\imath} \in [m]$ a strategy $a_{\tilde{\imath}}$ with payoff-matrix
$$\alpha - \epsilon + (R_{\tilde{\imath}j} - R_{ij})_{i,j}$$
3. for each $\tilde{\jmath} \in [n]$ a strategy $b_{\tilde{\jmath}}$ with payoff-matrix
$$\alpha - \epsilon + (C_{i\tilde{\jmath}} - C_{ij})_{i,j}$$

Let $\sigma_2 \in \Delta_m$ and $\sigma_3 \in \Delta_n$ be mixed strategies for players 2 and 3. Then

1. the expected payoff for player one when playing strategy v is one minus the social welfare of the game specified by R and C if players two and three play the strategy profile (σ_2, σ_3),
2. the expected payoff when playing $a_{\tilde{\imath}}$ is $\alpha - \epsilon$ plus player 2's gain when defecting to strategy $\tilde{\imath}$,
3. the expected payoff when playing $b_{\tilde{\jmath}}$ is $\alpha - \epsilon$ plus player 3's gain when defecting to strategy $\tilde{\jmath}$.

In particular, if $\sigma_2 \in \Delta_m$ and $\sigma_3 \in \Delta_n$ are an ϵ-Nash equilibrium with average payoff $> 1 - \alpha$, then no strategy for player 1 in $\mathcal{H}$ will have expected payoff $> \alpha$, if players 2 and 3 play according to σ_2 and σ_3. Thus, $\text{minmax}_1\mathcal{H} \leq \alpha$ in this case.

On the other hand, suppose that $\mathcal{G}$ has no 2ϵ-Nash equilibrium with average payoff $> 1 - \alpha - \epsilon$. Let σ_2 and σ_3 be strategies for player 2 and 3 in $\mathcal{H}$. If player 1 receives payoff $< \alpha + \epsilon$ when responding to σ_2 and σ_3 with strategy v, then the average payoff of (σ_2, σ_3), as a pair of strategies in $\mathcal{G}$, will be at least $1 - \alpha - \epsilon$. By our assumption on $\mathcal{G}$, (σ_2, σ_3) can not be an 2ϵ-Nash equilibrium, i.e., one of the players can gain more than 2ϵ by deviating. But then one of the strategies $a_{\tilde{\imath}}, b_{\tilde{\jmath}}$ will give player 1 an expected payoff of at least $\alpha + \epsilon$ in $\mathcal{H}$. Therefore $\text{minmax}_1\mathcal{H} > \alpha + \epsilon$ in this case. $\square$

4. Conclusion

We have considered a promise graph problem, which is hard assuming standard hardness assumptions on detecting planted cliques in random graphs. We have shown that the problem of approximating the minmax value in 3-player games with 0-1 payoffs is at least as hard as this promise graph problem. To this end we have given both a randomized many-one reduction and a deterministic Turing reduction. In doing this we believe we have given a satisfactory answer (in the negative) to the question of whether the minmax value in 3-player games can be approximated in polynomial time within any additive error $\epsilon > 0$. We leave open the problem of whether the minmax value of 3-player games can be approximated within *some* nontrivial additive error $0 < \epsilon < 1/2$ in polynomial time.

References

[1] Alon, N., Krivelevich, M., Sudakov, B.: Finding a large hidden clique in a random graph. Random Structures & Algorithms 13(3-4), 457–466 (1998)

[2] Alon, N., Yuster, R., Zwick, U.: Color-coding. J. ACM 42(4), 844–856 (1995)

[3] Austrin, P., Braverman, M., Chlamtáč, E.: Inapproximability of NP-complete variants of nash equilibrium. In: APPROX and RANDOM 2011. Lecture Notes in Computer Science, vol. 6845, pp. 13–25. Springer (2011)

[4] Bollobas, B.: Random Graphs. Cambridge University Press (2001)

[5] Borgs, C., Chayes, J., Immorlica, N., Kalai, A.T., Mirrokni, V., Papadimitriou, C.: The myth of the folk theorem. Games and Economic Behavior 70(1), 34 – 43 (2010), special Issue In Honor of Ehud Kalai

[6] Chen, X., Deng, X., Teng, S.H.: Settling the complexity of computing two-player Nash equilibria. J. ACM 56, 14:1–14:57 (2009)

[7] Conitzer, V., Sandholm, T.: Complexity results about nash equilibria. In: IJCAI 2003. pp. 765–771. Morgan Kaufmann (2003)

[8] Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The complexity of computing a Nash equilibrium. SIAM Journal on Computing 39(1), 195–259 (2009)

[9] Daskalakis, C., Mehta, A., Papadimitriou, C.H.: Progress in approximate nash equilibria. In: EC 2007. pp. 355–358. ACM (2007)

[10] Feige, U., Krauthgamer, R.: Finding and certifying a large hidden clique in a semirandom graph. Random Structures & Algorithms 16(2), 195–208 (2000)

[11] Gilboa, I., Zemel, E.: Nash and correlated equilibria: Some complexity considerations. Games and Economic Behavior 1(1), 80–93 (1989)

[12] Hansen, K.A., Hansen, T.D., Miltersen, P.B., Sørensen, T.B.: Approximability and parameterized complexity of minmax values. In: WINE 2008. Lecture Notes in Computer Science, vol. 5385, pp. 684–695. Springer (2008)

[13] Hazan, E., Krauthgamer, R.: How hard is it to approximate the best nash equilibrium? SIAM Journal on Computing 40(1), 79–91 (2011)

[14] Jerrum, M.: Large cliques elude the metropolis process. Random Structures & Algorithms 3(4), 347–359 (1992)

[15] Juels, A., Peinado, M.: Hiding cliques for cryptographic security. Designs, Codes and Cryptography 20, 269–280 (2000)

[16] Kučera, L.: Expected complexity of graph partitioning problems. Discrete Appl. Math. 57, 193–212 (1995)

[17] Lipton, R.J., Markakis, E., Mehta, A.: Playing large games using simple strategies. In: EC 2003. pp. 36–41. ACM (2003)

[18] Lipton, R.J., Young, N.E.: Simple strategies for large zero-sum games with applications to complexity theory. In: STOC 1994. pp. 734–740. ACM Press (1994)

[19] Minder, L., Vilenchik, D.: Small clique detection and approximate nash equilibria. In: APPROX and RANDOM 2009. Lecture Notes in Computer Science, vol. 5687, pp. 673–685. Springer (2009)

[20] Tsaknakis, H., Spirakis, P.G.: An optimization approach for approximate nash equilibria. Internet Mathematics 5(4), 365–382 (2008)

A. Proof of Lemma 4

In this section we give a proof of Lemma 4. Here we do not attempt to optimize parameters. We let $p = \eta/2e$, and let c_1 be a constant to be specified later.

Consider fixed sets $S, T \subseteq V$ of size $|S| = |T| = c_1 \ln n$. We will estimate the probability that $u_S^\intercal A u_T \le \eta$, and after that take a union bound over all such sets S and T.

The number of potential edges between S and T in G is exactly given by $\ell = |S|\,|T| - \binom{|S\cap T|}{2} - |S \cap T| \ge (c_1 \ln n)^2/3$ for large enough n. Letting X be a random variable denoting the number of edges between S and T we have $\mathrm{E}[X] = p\ell$. Note that if $X \le \eta/2 \cdot \ell$ then we have

$$u_S^\intercal A u_T \le \eta \cdot \frac{l}{|S|\,|T|} \le \eta,$$

because each edge is counted at most twice on the left hand side. By the Chernoff bound for the upper tail we have

$$\begin{aligned} \Pr[X \ge \eta/2 \cdot \ell] &= \Pr[X \ge e\,\mathrm{E}[X]] \le \left(\frac{e^{e-1}}{e^e}\right)^{\mathrm{E}[X]} \\ &= \exp(-E[X]) = \exp(-p\ell) \le \exp(-\frac{p}{3} \cdot (c_1 \ln n)^2) \ . \end{aligned}$$

We then wish to take a union over all choices of sets S and T. We have at most $\binom{n}{c_1 \ln n}^2 \le \exp(2c_1(\ln n)^2)$ such sets. We can thus obtain the statement of the lemma, by letting $c_1 > 6/p = 12e/\eta$.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de