Electronic Colloquium on Computational Complexity, Report No. 81 (2006)

Polynomial Algorithms for Approximating Nash Equilibria of Bimatrix Games

Spyros C. Kontogiannis[1,3], Panagiota N. Panagopoulou[2,3], and Paul G. Spirakis[2,3]

[1] Computer Science Department, University of Ioannina, 45110 Ioannina, Greece.
E-mail: kontog@cs.uoi.gr
[2] Department of Computer Engineering and Informatics, Patras University, Greece.
[3] Research Academic Computer Technology Institute
N. Kazantzaki Str., Patras University, Rion, Patras GR 26500, Greece.
E-mail: panagopp@cti.gr, spirakis@cti.gr

Abstract. We focus on the problem of computing an ϵ-Nash equilibrium of a bimatrix game, when ϵ is an absolute constant. We present a simple algorithm for computing a $\frac{3}{4}$-Nash equilibrium for any bimatrix game in strongly polynomial time and we next show how to extend this algorithm so as to obtain a (potentially stronger) parameterized approximation. Namely, we present an algorithm that computes a $\frac{2+\lambda+\epsilon}{4}$-Nash equilibrium for any ϵ, where λ is the minimum, among all Nash equilibria, expected payoff of either player. The suggested algorithm runs in time polynomial in $\frac{1}{\epsilon}$ and the number of strategies available to the players.

1 Introduction

Motivation, Framework and Overview. Non-cooperative game theory has been extensively used in understanding the phenomena observed when decision makers interact. A game consists of a set of players, and, for each player, a set of strategies available to her as well as a payoff function mapping each strategy profile (i.e. each combination of strategies, one for each player) to a real number that captures the preferences of the player over the possible outcomes of the game. The most important solution concept in non-cooperative game theory is the notion of *Nash equilibrium*: it is a strategy profile such that no player would have an incentive to unilaterally deviate from her strategy, i.e. no player could increase her payoff by choosing another strategy while the rest of the players persevered their strategies.

Despite of the certain existence of such equilibria, the problem of finding any Nash equilibrium even for games involving only two players has been recently proved to be computationally difficult. This fact emerged the computation of *approximate* Nash equilibria, also referred to as ϵ-Nash equilibria. An ϵ-Nash equilibrium is a strategy profile such that no deviating player could achieve a payoff higher than the one that the specific profile gives her, plus ϵ.

In this work, we focus on the problem of approximating Nash equilibria of 2-player games. We propose simple and efficient algorithms for computing ϵ-Nash equilibria of such games, for sufficiently small absolute constants ϵ.

Previous Work. Nash [6] introduced the concept of Nash equilibria in non-cooperative games and proved that any game possesses at least one such equilibrium; however, the

ISSN 1433-8092

computational complexity of finding a Nash equilibrium used to be a wide open problem for several years. Recently, Chen and Deng [1] proved that the problem is PPAD-complete for bimatrix games in which each player has n available pure strategies.

In [5] it was shown that, for any bimatrix game and for any *constant* $\epsilon > 0$, there exists an ϵ-Nash equilibrium with only logarithmic support (in the number n of available pure strategies). This result directly yields a quasi-polynomial ($n^{O(\ln n)}$) algorithm for computing such an approximate equilibrium.

In [2] it was shown that the problem of computing a $\frac{1}{n^{\Theta(1)}}$-Nash equilibrium is PPAD-complete, and that bimatrix games are unlikely to have a fully polynomial time approximation scheme (unless PPAD $\subseteq$ P). However, it was conjectured that it is unlikely that finding an ϵ-Nash equilibrium is PPAD-complete when ϵ is an absolute constant.

Our Results. In this work, we deal with the problem of computing an ϵ-Nash equilibrium of a bimatrix game, for some constant ϵ. We first present a simple algorithm for computing a $\frac{3}{4}$-Nash equilibrium for any bimatrix game in strongly polynomial time (Lemma 1).

Next we show how to extend this result so as to obtain a parameterized and potentially stronger approximation. More specifically, we present an algorithm that computes a $\frac{2+\lambda+\epsilon}{4}$-Nash equilibrium for any ϵ, where λ is the minimum, among all Nash equilibria, expected payoff of either player (Theorem 3). The suggested algorithm runs in time polynomial in $\frac{1}{\epsilon}$ and the number of strategies available to the players.

Organization. In Section 2 we present the notation used throughout this paper, together with the definitions of bimatrix games, Nash equilibria and approximate Nash equilibria, as well as some previous results on the problem of approximating Nash equilibria.

Our first algorithm for computing a $\frac{3}{4}$-Nash equilibrium is described in Section 3, while in Section 4 we present an extension of this algorithm that can give a stronger approximation. We conclude, in Section 5, with a discussion of our results and suggestions for further research.

2 Background

2.1 Notation

For an integer n, let $[n] = \{1, 2, \ldots, n\}$. For a $n \times 1$ vector $\mathbf{x}$ we denote by $x_1, x_2, \ldots x_n$ the components of $\mathbf{x}$ and by $\mathbf{x}^T$ the transpose of $\mathbf{x}$. For an $n \times m$ matrix A, we denote $a_{i,j}$ the element in the i-th row and j-th column of A. Let $\mathbb{P}^n$ be the set of all probability vectors in n dimensions, i.e.

$$\mathbb{P}^n \equiv \left\{ \mathbf{x} \in \mathbb{R}^n \ : \ \sum_{i=1}^{n} x_i = 1 \ \text{ and } \ x_i \geq 0 \ \text{ for all } \ i \in [n] \right\} \ .$$

Denote $\mathbb{R}_{[0:1]}^{n \times m}$ the set of all $n \times m$ matrices with real entries between 0 and 1, i.e.

$$\mathbb{R}_{[0:1]}^{n \times m} \equiv \{ A \in \mathbb{R}^{n \times m} \ : \ 0 \leq a_{i,j} \leq 1 \ \text{ for all } \ i \in [n], j \in [m] \} \ .$$

2.2 Bimatrix games

A *noncooperative game* $\varGamma = \langle N, (S_i)_{i\in N}, (u_i)_{i\in N}\rangle$ comprises (i) a finite set of *players* N, (ii) a nonempty finite set of *pure strategies* S_i for each player $i \in N$ and (iii) a *payoff function* $u_i : \times_{i\in N} S_i \rightarrow \mathbb{R}$ for each player $i \in N$.

Bimatrix games [3, 4] are a special case of 2-player games (i.e. $|N| = 2$) such that the payoff functions can be described by two real $n \times m$ matrices A and B, where $n = |S_1|$ and $m = |S_2|$. More specifically, the n rows of A, B represent the pure strategies of the first player (the *row player*) and the m columns represent the pure strategies of the second player (the *column player*). Then, when the row player chooses strategy i and the column player chooses strategy j, the former gets payoff $a_{i,j}$ while the latter gets payoff $b_{i,j}$. Based on this observation, bimatrix games are denoted by $\varGamma = \langle A, B\rangle$.

A *mixed strategy* for player $i \in N$ is a probability distribution on the set of her pure strategies S_i. In a bimatrix game $\varGamma = \langle A, B\rangle$, a mixed strategy for the row player can be expressed as a probability vector $\mathbf{x} \in \mathbb{P}^n$ while a mixed strategy for the column player can be expressed as a probability vector $\mathbf{y} \in \mathbb{P}^m$. When the row player chooses mixed strategy $\mathbf{x}$ and the column player chooses $\mathbf{y}$, then the players get *expected payoffs* $\mathbf{x}^T A\mathbf{y}$ (row player) and $\mathbf{x}^T B\mathbf{y}$ (column player). The *support* of a mixed strategy is the set of pure strategies that are assigned non-zero probability.

2.3 Nash equilibria and ϵ-Nash equilibria

A Nash equilibrium [6] for a game $\varGamma$ is a combination of (pure or mixed) strategies, one for each player, such that no player could increase her payoff by unilaterally changing her strategy. We formally give the definition of a Nash equilibrium and an ϵ-Nash equilibrium for a bimatrix game.

Definition 1 (Nash equilibrium). *A pair of strategies* $(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$ *is a* Nash equilibrium *for the bimatrix game* $\varGamma = \langle A, B\rangle$ *if*

(i) For every (mixed) strategy $\mathbf{x}$ *of the row player,* $\mathbf{x}^T A\tilde{\mathbf{y}} \leq \tilde{\mathbf{x}}^T A\tilde{\mathbf{y}}$ *and*
(ii) For every (mixed) strategy $\mathbf{y}$ *of the column player,* $\tilde{\mathbf{x}}^T B\mathbf{y} \leq \tilde{\mathbf{x}}^T B\tilde{\mathbf{y}}$.

Definition 2 (ϵ-Nash equilibrium). *For any* $\epsilon > 0$ *a pair of strategies* $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ *is an* ϵ-Nash equilibrium *for the bimatrix game* $\varGamma = \langle A, B\rangle$ *if*

(i) For every (mixed) strategy $\mathbf{x}$ *of the row player,* $\mathbf{x}^T A\hat{\mathbf{y}} \leq \hat{\mathbf{x}}^T A\hat{\mathbf{y}} + \epsilon$, *and*
(ii) For every (mixed) strategy $\mathbf{y}$ *of the column player,* $\hat{\mathbf{x}}^T B\mathbf{y} \leq \hat{\mathbf{x}}^T B\hat{\mathbf{y}} + \epsilon$.

Positively normalized bimatrix games. As pointed out in [2], since the notion of ϵ-Nash equilibria is defined in the additive fashion, it is important to consider bimatrix games with normalized matrices so as to study their complexity. That is, the absolute value of each entry in the matrices is bounded, for example by 1. [5] also used a similar normalization, which we adopt in this paper and describe it below.

Consider the $n \times m$ bimatrix game $\Gamma = \langle A, B \rangle$ and let c, d be two arbitrary positive real constants. Suppose that $(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$ is a Nash equilibrium for Γ and $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ is an ϵ-Nash equilibrium for Γ. Let $\mathbf{x}$ and $\mathbf{y}$ be any strategy of the row and column player respectively. Now consider the game $\Gamma' = \langle cA, dB \rangle$. Then it holds that

$$\mathbf{x}^T (cA)\tilde{\mathbf{y}} = c\mathbf{x}^T A\tilde{\mathbf{y}} \le c\tilde{\mathbf{x}} A\tilde{\mathbf{y}} = \tilde{\mathbf{x}}^T (cA)\tilde{\mathbf{y}}$$

and, similarly,

$$\tilde{\mathbf{x}}^T (dB)\mathbf{y} \le \tilde{\mathbf{x}}^T (dB)\tilde{\mathbf{y}} \ .$$

Moreover,

$$\mathbf{x}^T (cA)\hat{\mathbf{y}} \le \hat{\mathbf{x}}^T (cA)\hat{\mathbf{y}} + c\epsilon$$

and

$$\hat{\mathbf{x}}^T (dB)\mathbf{y} \le \hat{\mathbf{x}}^T (dB)\hat{\mathbf{y}} + d\epsilon \ .$$

Hence Γ and Γ' have precisely the same set of Nash equilibria; furthermore, any ϵ-Nash equilibrium for Γ is a $\ell\epsilon$-Nash equilibrium for Γ' (where $\ell = \max\{c, d\}$) and vice versa.

Now let C be an $n \times m$ matrix such that, for all (columns) $j \in [m]$, $c_{i,j} = c_j \in \mathbb{R}$ for all $i \in [n]$. Similarly, let D be an $n \times m$ matrix such that, for all (rows) $i \in [m]$, $d_{i,j} = d_i \in \mathbb{R}$ for all $j \in [m]$. Note that, for every pair $\mathbf{x} \in \mathbb{P}^n$ and $\mathbf{y} \in \mathbb{P}^m$,

$$\mathbf{x}^T C\mathbf{y} = \sum_{j=1}^{m}\sum_{i=1}^{n} c_{i,j} x_i y_j = \sum_{j=1}^{m} y_j \sum_{i=1}^{n} c_j x_i = \sum_{j=1}^{m} c_j y_j$$

and

$$\mathbf{x}^T D\mathbf{y} = \sum_{i=1}^{n}\sum_{j=1}^{m} d_{i,j} x_i y_j = \sum_{i=1}^{n} x_i \sum_{j=1}^{m} d_i y_j = \sum_{i=1}^{n} d_i x_i \ .$$

Consider now the game $\Gamma'' = \langle C + A, D + B \rangle$. Then

$$\mathbf{x}^T (C + A)\tilde{\mathbf{y}} = \mathbf{x}^T C\tilde{\mathbf{y}} + \mathbf{x}^T A\tilde{\mathbf{y}} \le \sum_{j=1}^{m} c_j \tilde{y}_j + \tilde{\mathbf{x}}^T A\tilde{\mathbf{y}} = \tilde{\mathbf{x}}^T (C + A)\tilde{\mathbf{y}}$$

and similarly

$$\tilde{\mathbf{x}}^T (D + B)\mathbf{y} \le \tilde{\mathbf{x}}^T (D + B)\tilde{\mathbf{y}} \ .$$

Also, it holds that

$$\mathbf{x}^T (C + A)\hat{\mathbf{y}} = \mathbf{x}^T C\hat{\mathbf{y}} + \mathbf{x}^T A\hat{\mathbf{y}} \le \sum_{j=1}^{m} c_j \hat{y}_j + \hat{\mathbf{x}}^T A\hat{\mathbf{y}} + \epsilon = \hat{\mathbf{x}}^T (C + A)\hat{\mathbf{y}} + \epsilon$$

and similarly

$$\hat{\mathbf{x}}^T (D + B)\mathbf{y} \le \hat{\mathbf{x}}^T (D + B)\hat{\mathbf{y}} + \epsilon \ .$$

Thus Γ and Γ'' are equivalent as regards their sets of Nash equilibria, as well as their sets of ϵ-Nash equilibria.

This equivalence allows us to focus only on bimatrix games where the payoffs are between 0 and 1, i.e. on games $\langle A, B \rangle$ where $A, B \in \mathbb{R}^{m \times n}_{[0:1]}$. Such games are referred to as *positively normalized* [2].

2.4 Existence and tractability of ϵ-Nash equilibria

Consider a bimatrix game $\Gamma = \langle A, B \rangle$ and let $(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$ be a Nash equilibrium for Γ. Fix a positive integer k and assume that we form a multiset S_1 by sampling k times from the set of pure strategies of the row player, independently at random according to the distribution $\tilde{\mathbf{x}}$. Similarly, assume we form a multiset S_2 by sampling k times from set of pure strategies of the column player, independently at random according to the distribution $\tilde{\mathbf{y}}$. Let $\hat{\mathbf{x}}$ be the mixed strategy for the row player that assigns probability $1/k$ to each member of S_1 and 0 to all other pure strategies, and let $\hat{\mathbf{y}}$ be the mixed strategy for the column player that assigns probability $1/k$ to each member of S_2 and 0 to all other pure strategies. Clearly, if a pure strategy occurs α times in the multiset, then it is assigned probability α/k. Then $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ are called *k-uniform* [5] and the following holds:

Theorem 1 ([5]). *For any Nash equilibrium $(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$ of a positively normalized $n \times n$ bimatrix game and for every $\epsilon > 0$, there exists, for every $k \geq \frac{12 \ln n}{\epsilon^2}$, a pair of k-uniform strategies $\hat{\mathbf{x}}, \hat{\mathbf{y}}$ such that $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ is an ϵ-Nash equilibrium.*

However,

Theorem 2 ([2]). *The problem of computing a $\frac{1}{n^{\Theta(1)}}$-Nash equilibrium of a positively normalized $n \times n$ bimatrix game is* PPAD*-complete.*

Theorem 2 asserts that, unless PPAD $\subseteq$ P, there exists no fully polynomial time approximation scheme for computing equilibria in bimatrix games. However, this does not rule out the existence of a polynomial approximation scheme for computing an ϵ-Nash equilibrium when ϵ is an absolute constant, or even when $\epsilon = \Theta\left(\frac{1}{poly(\ln n)}\right)$. Furthermore, as observed in [2], if the problem of finding an ϵ-Nash equilibrium were PPAD-complete when ϵ is an absolute constant, then, due to Theorem 1, all PPAD problems would be solved in quasi-polynomial time, which is unlikely to be the case.

3 A $\frac{3}{4}$-Nash equilibrium

In this section we present a straightforward method for computing a $\frac{3}{4}$-Nash equilibrium for any positively normalized bimatrix game.

Lemma 1. *Consider any positively normalized $n \times m$ bimatrix game $\Gamma = \langle A, B \rangle$ and let $a_{i_1,j_1} = \max_{i,j} a_{i,j}$ and $b_{i_2,j_2} = \max_{i,j} b_{i,j}$. Then the pair of strategies $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ where $\hat{x}_{i_1} = \hat{x}_{i_2} = \hat{y}_{j_1} = \hat{y}_{j_2} = \frac{1}{2}$ is a $\frac{3}{4}$-Nash equilibrium for Γ.*

Proof. First observe that

$$\begin{aligned}
\hat{\mathbf{x}}^T A \hat{\mathbf{y}} &= \sum_{i=1}^{n} \sum_{j=1}^{m} \hat{x}_i \hat{y}_j a_{i,j} \\
&= \hat{x}_{i_1} \hat{y}_{j_1} a_{i_1,j_1} + \hat{x}_{i_1} \hat{y}_{j_2} a_{i_1,j_2} + \hat{x}_{i_2} \hat{y}_{j_1} a_{i_2,j_1} + \hat{x}_{j_1} \hat{y}_{j_1} a_{i_2,j_2} \\
&= \frac{1}{4} \left(a_{i_1,j_1} + a_{i_1,j_2} + a_{i_2,j_1} + a_{i_2,j_2} \right) \\
&\geq \frac{1}{4} a_{i_1,j_1} \quad .
\end{aligned}$$

Similarly,

$$\begin{aligned}\hat{\mathbf{x}}^T B\hat{\mathbf{y}} &= \sum_{i=1}^{n}\sum_{j=1}^{m} \hat{x}_i \hat{y}_j b_{i,j} \\ &= \hat{x}_{i_1}\hat{y}_{j_1} b_{i_1,j_1} + \hat{x}_{i_1}\hat{y}_{j_2} b_{i_1,j_2} + \hat{x}_{i_2}\hat{y}_{j_1} b_{i_2,j_1} + \hat{x}_{j_1}\hat{y}_{j_1} b_{i_2,j_2} \\ &= \frac{1}{4}\left(b_{i_1,j_1} + b_{i_1,j_2} + b_{i_2,j_1} + b_{i_2,j_2}\right) \\ &\geq \frac{1}{4} b_{i_2,j_2} \ .\end{aligned}$$

Now observe that, for any (mixed) strategies $\mathbf{x}$ and $\mathbf{y}$ of the row and column player respectively,

$$\mathbf{x}^T A\hat{\mathbf{y}} \leq a_{i_1,j_1} \quad \text{and} \quad \hat{\mathbf{x}}^T B\mathbf{y} \leq b_{i_2,j_2}$$

and recall that $a_{i,j}, b_{i,j} \in [0,1]$ for all $i \in N$, $j \in M$. Hence

$$\mathbf{x}^T A\hat{\mathbf{y}} \leq a_{i_1,j_1} = \frac{1}{4}a_{i_1,j_1} + \frac{3}{4}a_{i_1,j_1} \leq \hat{\mathbf{x}}^T A\hat{\mathbf{y}} + \frac{3}{4}$$

and

$$\hat{\mathbf{x}}^T B\mathbf{y} \leq b_{i_2,j_2} = \frac{1}{4}b_{i_2,j_2} + \frac{3}{4}a_{i_2,j_2} \leq \hat{\mathbf{x}}^T B\hat{\mathbf{y}} + \frac{3}{4} \ .$$

Thus $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ is a $\frac{3}{4}$-Nash equilibrium for Γ. □

4 A Parameterized Approximation

We now proceed in extending the technique used in the proof of Lemma 1 so as to obtain a parameterized, stronger approximation.

Theorem 3. *Consider a positively normalized $n \times m$ bimatrix game $\Gamma = \langle A, B\rangle$. Let λ_1^* (λ_2^*) be the minimum, among all Nash equilibria of Γ, expected payoff for the row (column) player and let $\lambda = \min\{\lambda_1^*, \lambda_2^*\}$. Then, for any $0 < \epsilon < 1$, there exists a $\frac{2+\lambda+\epsilon}{4}$-Nash equilibrium that can be computed in time polynomial in $\frac{1}{\epsilon}$, n and m.*

Proof. Observe that, for any pair of strategies $\mathbf{x}, \mathbf{y}$ of the row and column player respectively, it holds that $\mathbf{x}^T A\mathbf{y} \in [0,1]$ and $\mathbf{x}^T B\mathbf{y} \in [0,1]$. Given $0 < \epsilon < 1$, consider a partition of the interval $[0,1]$ into $\kappa = \frac{1}{\epsilon}$ subintervals $\delta_1 = [0,\epsilon], \delta_2 = (\epsilon, 2\epsilon], \ldots, \delta_\kappa = (1-\epsilon, 1]$.

Starting from interval δ_1, we seek until we find an interval $\delta_t = ((t-1)\epsilon, t\epsilon]$ for which there exists a probability vector $\mathbf{y}$ of the column player such that

$$\begin{aligned}\sum_{j=1}^{m} a_{i,j} y_j &\leq t\epsilon \quad \forall i \in [n] \\ \sum_{j=1}^{m} y_j &= 1 \ .\end{aligned}$$

We argue that there exists some t for which there exists a feasible solution to the preceding linear constraints. Indeed, let λ_1^* be the minimum, over all Nash equilibria, expected payoff of the row player. Clearly $\lambda_1^* \in \delta_t$ for some t. Then, the probability vector of the column player that corresponds to the specific Nash equilibrium satisfies the constraints.

Furthermore, note that, once we find a feasible solution $\mathbf{y}$ for an interval δ_t, then there exists at least one row $r \in [n]$ such that $\sum_j a_{r,j} y_j > (t-1)\epsilon$; otherwise $\mathbf{y}$ would be a feasible solution for the interval δ_{t-1}.

In a similar manner, we (independently) seek until we find an interval $\delta_s = ((s-1)\epsilon, s\epsilon]$ for which there exists a probability vector $\mathbf{x}$ of the row player such that

$$\sum_{i=1}^{n} b_{i,j} x_i \leq s\epsilon \quad \forall j \in [m]$$
$$\sum_{i=1}^{n} x_i = 1 \ .$$

Again, there exists some s for which there exists a feasible solution to the preceding linear constraints: let λ_2^* be the minimum, over all Nash equilibria, expected payoff of the column player. Clearly $\lambda_2^* \in \delta_s$ for some s. Then, the probability vector of the row player that corresponds to the specific Nash equilibrium satisfies the constraints.

As in the previous case, once we find a feasible solution $\mathbf{x}$ for an interval δ_s, then there exists at least one column $c \in [m]$ such that $\sum_i a_{i,c} x_i > (s-1)\epsilon$; otherwise $\mathbf{x}$ would be a feasible solution for the interval δ_{s-1}.

Suppose now that $\mathbf{y}$ and $\mathbf{x}$ are the feasible solutions found and consider the pair of strategies $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ for the row and column player respectively defined as follows:

$$\begin{aligned}
\hat{x}_i &= \frac{x_i}{2} \quad \forall i \in [n] - \{r\} \\
\hat{x}_r &= \frac{x_r}{2} + \frac{1}{2} \\
\hat{y}_j &= \frac{y_j}{2} \quad \forall j \in [m] - \{c\} \\
\hat{y}_c &= \frac{y_c}{2} + \frac{1}{2} \ .
\end{aligned}$$

Observe that

$$\begin{aligned}
\hat{\mathbf{x}}^T A \hat{\mathbf{y}} &= \sum_{i=1}^{n} \hat{x}_i \sum_{j=1}^{m} \hat{y}_j a_{i,j} \\
&= \sum_{i \neq r} \frac{x_i}{2} \sum_{j \neq c} \frac{y_j}{2} a_{i,j} + \sum_{i \neq r} \frac{x_i}{2} \left(\frac{y_c}{2} + \frac{1}{2} \right) a_{i,c} \\
&\quad + \left(\frac{x_r}{2} + \frac{1}{2} \right) \sum_{j \neq c} \frac{y_j}{2} a_{r,j} + \left(\frac{x_r}{2} + \frac{1}{2} \right) \left(\frac{y_c}{2} + \frac{1}{2} \right) a_{r,c} \\
&\geq \frac{1}{4} \sum_{j=1}^{m} a_{r,j} y_j
\end{aligned}$$

$$> \frac{(t-1)\epsilon}{4} \quad .$$

Furthermore, for each row $i \in [n]$,

$$\begin{aligned}
\sum_{j=1}^{m} \hat{y}_j a_{i,j} &= \sum_{j=1}^{m} \frac{y_j}{2} a_{i,j} + \frac{1}{2} a_{i,c} \\
&\le \frac{t\epsilon}{2} + \frac{1}{2} \\
&= \frac{(t-1)\epsilon}{4} + \frac{t\epsilon}{4} + \frac{\epsilon}{4} + \frac{1}{2} \\
&< \hat{\mathbf{x}}^T A \hat{\mathbf{y}} + \frac{2 + \lambda_1^* + \epsilon}{4} \quad .
\end{aligned}$$

Similarly,

$$\begin{aligned}
\hat{\mathbf{x}}^T B \hat{\mathbf{y}} &= \sum_{j=1}^{m} \hat{y}_j \sum_{i=1}^{n} \hat{x}_i b_{i,j} \\
&= \sum_{j \neq c} \frac{y_j}{2} \sum_{i \neq r} \frac{x_i}{2} b_{i,j} + \sum_{j \neq c} \frac{y_j}{2} \left(\frac{x_r}{2} + \frac{1}{2} \right) b_{r,j} \\
&\quad + \left(\frac{y_c}{2} + \frac{1}{2} \right) \sum_{i \neq r} \frac{x_i}{2} b_{i,c} + \left(\frac{y_c}{2} + \frac{1}{2} \right) \left(\frac{x_r}{2} + \frac{1}{2} \right) b_{r,c} \\
&\ge \frac{1}{4} \sum_{i=1}^{n} b_{i,c} x_i \\
&> \frac{(s-1)\epsilon}{4}
\end{aligned}$$

and, for each column $j \in [m]$,

$$\begin{aligned}
\sum_{i=1}^{n} \hat{x}_i b_{i,j} &= \sum_{i=1}^{n} \frac{x_i}{2} b_{i,j} + \frac{1}{2} b_{r,j} \\
&\le \frac{s\epsilon}{2} + \frac{1}{2} \\
&= \frac{(s-1)\epsilon}{4} + \frac{s\epsilon}{4} + \frac{\epsilon}{4} + \frac{1}{2} \\
&< \hat{\mathbf{x}}^T B \hat{\mathbf{y}} + \frac{2 + \lambda_2^* + \epsilon}{4} \quad .
\end{aligned}$$

Thus, for any $0 < \epsilon < 1$, we can compute a $\frac{2+\max\{\lambda_1^*, \lambda_2^*\}+\epsilon}{4}$-Nash equilibrium in polynomial time in $\frac{1}{\epsilon}, n$ and m. $\square$

Note that, for any bimatrix game $\Gamma = \langle A, B \rangle$, we can check in polynomial time whether there exists a Nash equilibrium in which each player chooses with probability 1 one of her pure strategies (i.e. a *pure Nash equilibrium*). If there exists such an equilibrium, then we can find it in polynomial time and there is no point in searching for ϵ-Nash equilibria. On the other hand, if all Nash equilibria are not pure, then the payoff of either player is strictly less than 1, hence $\lambda = \max\{\lambda_1^*, \lambda_2^*\} < 1$. Thus there exists some $0 < \bar{\epsilon} < 1$ such that $\frac{2+\lambda+\bar{\epsilon}}{4} < \frac{3}{4}$, assuring that the the algorithm described in the above proof can yield a stronger approximation than the one presented in Section 3.

An Application. The approximation factor achieved by the algorithm we just described depends on λ_1^* and λ_2^*. We believe that, in most situations, there exists a Nash equilibrium such that the payoff of the row player is small, and that there exists a (possibly different) Nash equilibrium such that the payoff of the column player is small, and thus the approximation achieved is close to $\frac{1}{2}$.

As an example, consider the $n \times n$ *generalized matching pennies* game $\Gamma = \langle A, B \rangle$ where A and B are described as follows:

$$a_{i,j} = \begin{cases} 1 \text{ if } i = j \\ 0 \text{ else} \end{cases}$$

$$b_{i,j} = \begin{cases} 1 \text{ if } j = i+1 \text{ or if } i = n \text{ and } j = 1 \\ 0 \text{ else} \end{cases} .$$

Observe that the pair of strategies $(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$ where $\tilde{x}_i = \tilde{y}_i = \frac{1}{n}$ for all $i \in [n]$ is a Nash equilibrium of the generalized matching pennies game. Indeed, for any $\mathbf{x}, \mathbf{y} \in \mathbb{P}^n$,

$$\mathbf{x}^T A \tilde{\mathbf{y}} = \frac{1}{n^2} \sum_{i=1}^{n} \sum_{j=1}^{n} a_{i,j} = \frac{1}{n^2} n = \frac{1}{n} = \tilde{\mathbf{x}}^T A \tilde{\mathbf{y}}$$

$$\tilde{\mathbf{x}}^T B \mathbf{y} = \frac{1}{n^2} \sum_{i=1}^{n} \sum_{j=1}^{n} b_{i,j} = \frac{1}{n^2} n = \frac{1}{n} = \tilde{\mathbf{x}}^T B \tilde{\mathbf{y}} \; .$$

Thus $(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$ is a Nash equilibrium[4] that gives each player a payoff equal to $\frac{1}{n}$. By applying Theorem 3, we can obtain for each $0 < \epsilon < 1$ a $\frac{1+1/n+\epsilon}{2}$-Nash equilibrium. Thus we can guarantee an approximation factor that tends to $\frac{1+\epsilon}{2}$ as $n \to \infty$.

5 Conclusions

In this paper we tried to approximate, within a constant additive factor, the problem of computing a Nash equilibrium in an arbitrary $n \times m$ bimatrix game.

The (additive) approximation parameter achieved by the algorithm described in the above proof of Theorem 3 depends on λ_1^* and λ_2^*, i.e. the minimum payoff, over all Nash equilibria, for the row and column player respectively. Observe that, so long as not *all*

[4] In fact it can be proved that $(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$ is the *unique* Nash equilibrium of the generalized matching pennies game.

Nash equilibria of the game give payoffs very close to 1 for either player, the algorithm gives an approximation very close to $\frac{1}{2}$. In other words, it suffices that there exists a Nash equilibrium that gives row player a payoff close to 0 and a Nash equilibrium (not necessarily the same!) that gives column player a payoff close to 0 so that the approximation achieved can be assured to be close to $\frac{1}{2}$. Furthermore, this is just a *sufficient* and not a necessary condition: recall that we only used λ_1^* and λ_2^* so as to prove the existence of feasible solutions to some linear constraints.

Furthermore, for both Lemma 1 and Theorem 3, we used a factor of $\frac{1}{2}$ to deal with the underlying Linear Complementarity Problem. More specifically, we tried to compute independently for each player a strategy that guarantees her a sufficiently large payoff, and then we "merged" in an equivalent way the strategies found with the ones needed by the other player so as to approximate a Nash equilibrium. We observed that, for the specific algorithms presented in these results, this factor of $\frac{1}{2}$ is optimal.

Albeit simple, we believe that the techniques described here are a first step towards establishing whether there exists any approximation scheme for computing an ϵ-Nash equilibrium and that our methods can be extended in order to achieve stronger approximations to the problem of finding Nash equilibria of bimatrix games.

References

1. Xi Chen and Xiaotie Deng: "Settling the Complexity of 2-Player Nash-Equilibrium". *Electronic Colloquium on Computational Complexity (ECCC)*, TR05-140, 2005.
2. Xi Chen, Xiaotie Deng and Shang-Hua Teng: "Computing Nash Equilibria: Approximation and Smoothed Complexity". *Electronic Colloquium on Computational Complexity (ECCC)*, TR06-023, 2006.
3. C. E. Lemke: "Bimatrix Equilibrium Points and Mathematical Programming". *Management Science*, Vol. 11, pp. 681–689, 1965.
4. C. E. Lemke and J. T. Howson: "Equilibrium Points of Bimatrix Games". *J. Soc. Indust. Appl. Math.*, Vol. 12, pp. 413–423, 1964.
5. Richard J. Lipton, Evangelos Markakis and Aranyak Mehta: "Playing Large Games using Simple Startegies". In *EC '03: Proceedings of the 4th ACM Conference on Electronic Commerce*, pp. 36–41, 2003.
6. J. Nash: "Noncooperative Games". *Annals of Mathematics*, Vol. 54, pp. 289–295, 1951.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de/