

Progress in Approximate Nash Equilibria

Constantinos Daskalakis^{*}
Computer Science
UC Berkeley
Berkeley, CA
costis@cs.berkeley.edu

Aranyak Mehta
IBM Almaden Research
Center
San Jose, CA
mehtaa@us.ibm.com

Christos Papadimitriou
Computer Science
UC Berkeley
Berkeley, CA
christos@cs.berkeley.edu

ABSTRACT

It is known [5] that an additively ϵ -approximate Nash equilibrium (with supports of size at most two) can be computed in polynomial time in any 2-player game with $\epsilon = .5$. It is also known that no approximation better than .5 is possible unless equilibria with support larger than $\log n$ are considered, where n is the number of strategies per player. We give a polynomial algorithm for computing an ϵ -approximate Nash equilibrium in 2-player games with $\epsilon \approx .38$; our algorithm computes equilibria with arbitrarily large supports.

Categories and Subject Descriptors

F.2.0 [Theory of Computation]: analysis of algorithms and problem complexity—*general*

General Terms

Algorithms, Economics, Theory

Keywords

Nash Equilibrium, Approximate Nash, Algorithm

1. INTRODUCTION

It was recently shown that finding a Nash equilibrium is PPAD-complete [4], even for 2-player games [2]. As a consequence, finding approximate Nash equilibria has emerged as the main remaining open question in the area of equilibrium computation. The most commonly studied form of approximation is *additive approximation* in games in which all utilities have been normalized to be between 0 and 1 (this is a common assumption, since scaling the utilities of a player by any positive factor, and applying any additive constant, results in an equivalent game). A set of mixed strategies is called an ϵ -approximate Nash equilibrium, where $\epsilon > 0$, if

^{*}The first and third authors were supported through NSF grant CCF - 0635319, a gift from Yahoo! Research and a MICRO grant.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'07, June 13–16, 2007, San Diego, California, USA.

Copyright 2007 ACM 978-1-59593-653-0/07/0006 ...\$5.00.

for each player all strategies have expected payoff that is at most ϵ more than the expected payoff of the given strategy. Clearly, any mixed strategy combination is a 1-approximate Nash equilibrium, and it is quite straightforward to find a $\frac{3}{4}$ -approximate Nash equilibrium in 2-player games by examining all supports of size two; see [8] for a slightly improved result. In [9] it was shown that an ϵ -approximate Nash equilibrium can be found in time $O(n^{\frac{\log n}{\epsilon^2}})$ by examining all supports of size $\frac{\log n}{\epsilon^2}$. It was pointed out in [1] that no algorithm that examines supports smaller than about $\log n$ can achieve an approximation better than $\frac{1}{4}$, even for zero-sum games. In fact [6] have sharpened the $\frac{1}{4}$ to a $\frac{1}{2}$.

A very simple algorithm achieving $\epsilon = \frac{1}{2}$ in 2-person games was pointed out in [5]: For any strategy i of player I let j be the best response of player II, and let k be the best response of player I to j . Then player I plays an equal mixture of i and k , while player II plays j . The proof of $\frac{1}{2}$ -approximation is not very hard.

In this paper we give an algorithm which breaks the barrier of $\frac{1}{2}$ by considering supports of arbitrarily large cardinality. It achieves an approximation ratio $\frac{3-\sqrt{5}}{2} + \epsilon \approx .38 + \epsilon$, for any $\epsilon > 0$, in time $n^{O(\frac{1}{\epsilon^2})}$. It is based on the following two ideas: (a) If the values (u, v) of a Nash equilibrium to the two players were known, then we would be able to find a $\max\{u, v\}$ -approximate Nash equilibrium by solving a set of linear inequalities; and (b) for every Nash equilibrium, there is a pair of mixed strategies with support size $O(\frac{1}{\epsilon^2})$ which approximates within ϵ the true values of that Nash equilibrium [9]. Our technique uses both these ideas and also takes care of the interaction between them. The linear program in (a) will, of course, return in general mixed strategies of arbitrarily large support.

2. PRELIMINARIES

We consider normal form games between two players, the *row player* and the *column player*, each with n strategies at his disposal. The game is defined by two $n \times n$ payoff matrices, R for the row player, and C for the column player. The pure strategies of the row player correspond to the n rows and the pure strategies of the column player correspond to the n columns. If the row player plays row i and the column player plays column j , then the row player receives a payoff of R_{ij} and the column player gets C_{ij} . Payoffs are extended linearly to pairs of mixed strategies — if the row player plays a probability distribution x over the rows and column player plays a distribution y over the columns, then

the row player gets a payoff of $x^T R y$ and the column player gets a payoff of $x^T C y$.

A *Nash equilibrium* in this setting is a pair of mixed strategies, x^* for the row player and y^* for the column player, such that neither player has an incentive to unilaterally defect. Note that, by linearity, the best defection is to a pure strategy. Let e_i denote the vector with a 1 at the i th coordinate and 0 elsewhere. A pair of mixed strategies (x^*, y^*) is a Nash equilibrium if

$$\forall i = 1..n, \quad e_i^T R y^* \leq x^{*T} R y^*$$

$$\forall i = 1..n, \quad x^{*T} C e_i \leq x^{*T} C y^*$$

It can be easily shown that every pair of equilibrium strategies of a game does not change upon multiplying all the entries of a payoff matrix by a constant, and upon adding the same constant to each entry. We shall therefore assume that the entries of both payoff matrices R and C are between 0 and 1.

For $\epsilon > 0$, we define an ϵ -approximate Nash equilibrium to be a pair of mixed strategies x^* for the row player and y^* for the column player, so that the incentive to unilaterally deviate is at most ϵ :

$$\forall i = 1..n, \quad e_i^T R y^* \leq x^{*T} R y^* + \epsilon$$

$$\forall i = 1..n, \quad x^{*T} C e_i \leq x^{*T} C y^* + \epsilon$$

A stronger notion of approximately equilibrium strategies was introduced in [7, 4]: For $\epsilon > 0$, a *well-supported ϵ -approximate Nash equilibrium*, or an *ϵ -well-supported Nash equilibrium*, is a pair of mixed strategies, x^* for the row player and y^* for the column player, so that a player plays only approximately best-response pure strategies with non-zero probability:

$$\forall i : x_i^* > 0 \Rightarrow e_i^T R y^* \geq e_j^T R y^* - \epsilon, \quad \forall j$$

$$\forall i : y_i^* > 0 \Rightarrow x^{*T} C e_i \geq x^{*T} C e_j - \epsilon, \quad \forall j$$

This is indeed a stronger definition, in the sense that every ϵ -well supported Nash equilibrium is also an ϵ -approximate Nash equilibrium, but the converse need not be true. However, the following lemma from [3] shows that there does exist a polynomial relationship between the two:

LEMMA 2.1. [3] *For every 2 player normal form game, for every $\epsilon > 0$, given an $\frac{\epsilon^2}{8n}$ -approximate equilibrium we can compute in polynomial time an ϵ -well-supported equilibrium.*

It is known, see [4, 2, 3], that computing an $\frac{1}{n^\alpha}$ -well supported Nash equilibrium is PPAD-complete, for any constant $\alpha > 0$, and, by the above lemma, so is computing a $\frac{1}{n^{\alpha+1}}$ -approximate equilibrium. Therefore, a FPTAS for the problem is unlikely. In [5] we provide an algorithm for computing a $\frac{1}{2}$ -approximate Nash equilibrium via a simple algorithm which considers strategies of support 2 and an algorithm for computing a $\frac{2}{3}$ -well supported equilibrium conditional on some graph theoretic conjecture. Another attempt towards computing approximate equilibria (although with inferior approximation guarantees) can be found in [8].

3. COMPUTATION OF APPROXIMATE NASH EQUILIBRIA

3.1 An Existence Proof

Using the properties of Nash equilibrium and Hoeffding-Chernoff bounds we derive the following.

LEMMA 3.1. *For any $\epsilon > 0$ and for any two-player game $\mathcal{G} = (R, C)$, where R, C are $n \times n$ matrices with entries in $[0, 1]$, there exist $v_R, v_C \in [0, 1]$, mixed strategies x, α for the row player, y, β for the column player, with $|\text{supp}(\alpha)| \leq 4/\epsilon^2$, $|\text{supp}(\beta)| \leq 4/\epsilon^2$, such that the following are satisfied*

$$|\alpha^T R \beta - v_R| \leq \epsilon \quad (1)$$

$$\forall i \in \text{supp}(\alpha) : e_i^T R y = v_R \quad (2)$$

$$|x^T R \beta - v_R| \leq \epsilon \quad (3)$$

$$\forall i : e_i^T R y \leq v_R \quad (4)$$

$$|\alpha^T C \beta - v_C| \leq \epsilon \quad (5)$$

$$\forall j \in \text{supp}(\beta) : x^T C e_j = v_C \quad (6)$$

$$|\alpha^T C y - v_C| \leq \epsilon \quad (7)$$

$$\forall j : x^T C e_j \leq v_C \quad (8)$$

PROOF. Let (x^*, y^*) be a Nash equilibrium of game \mathcal{G} and let $v_R = x^{*T} R y^*$, $v_C = x^{*T} C y^*$ be the values obtained by the two players in the Nash equilibrium. Let us choose $x = x^*$ and $y = y^*$. We will argue that there exists a pair of mixed strategies α, β for the row and column player respectively so that the required properties hold.

To show that such strategies exist we apply the probabilistic method. Suppose that we take t independent samples from the strategy space of the row player according to the distribution induced by x^* and let us denote by A the resulting multiset of pure strategies of the row player. Similarly, let us take t samples according to y^* and denote by B the resulting multiset for the column player. Let then α be the mixed strategy corresponding to the uniform distribution on multiset A and β the mixed strategy corresponding to the uniform distribution on B .

We are interested in the probability that x, y, α and β satisfy the properties stated above. Clearly, Properties (2), (4), (6), (8) are satisfied with probability 1 by the definition of Nash equilibrium and since every strategy in A is in the support of x and every strategy in B in the support of y . So we only need to worry about Properties (1), (3), (5) and (7).

Let X and Y be independent random variables such that

$$X = e_i, \text{ with probability } x(i), \text{ for all } i \in \{1, \dots, n\},$$

$$Y = e_i, \text{ with probability } y(i), \text{ for all } i \in \{1, \dots, n\}.$$

Let then X_1, \dots, X_t be t copies of variable X and Y_1, \dots, Y_t be t copies of variable Y , where variables $X_1, \dots, X_t, Y_1, \dots, Y_t$ are mutually independent. Denoting $\mathcal{A} = \frac{1}{t} \sum_{i=1}^t X_i$ and $\mathcal{B} = \frac{1}{t} \sum_{i=1}^t Y_i$ we will argue that with positive probability the random variables \mathcal{A} and \mathcal{B} satisfy Properties (1), (3), (5) and (7), i.e.

$$\mathbb{P}[(|\mathcal{A}^T R \mathcal{B} - v_R| \leq \epsilon) \wedge (|x^T R \mathcal{B} - v_R| \leq \epsilon) \wedge (|\mathcal{A}^T C \mathcal{B} - v_C| \leq \epsilon) \wedge (|\mathcal{A}^T C y - v_C| \leq \epsilon)] > 0. \quad (9)$$

Clearly, $\mathbb{E}[x^T RY_i] = \sum_{i,j} R_{ij}x(i)\mathbb{E}[\mathbf{1}_{\{Y=e_j\}}] = \sum_{i,j} R_{ij}x(i)y(j) = v_R$. Therefore, a Chernoff bound on the sequence of random variables $Z_i = x^T RY_i$ gives

$$\mathbb{P}[|x^T R\mathcal{B} - v_R| > \epsilon] \leq e^{-2t\epsilon^2}. \quad (10)$$

Via similar arguments

$$\mathbb{P}[|\mathcal{A}^T C y - v_C| > \epsilon] \leq e^{-2t\epsilon^2}.$$

To bound the probability of the event $|\mathcal{A}^T R\mathcal{B} - v_R| > \epsilon$ we note that

$$|\mathcal{A}^T R\mathcal{B} - v_R| \leq |\mathcal{A}^T R\mathcal{B} - x^T R\mathcal{B}| + |x^T R\mathcal{B} - v_R|.$$

Therefore,

$$\begin{aligned} \mathbb{P}[|\mathcal{A}^T R\mathcal{B} - v_R| > \epsilon] \\ &\leq \mathbb{P}[|\mathcal{A}^T R\mathcal{B} - x^T R\mathcal{B}| > \epsilon/2 \vee |x^T R\mathcal{B} - v_R| > \epsilon/2] \\ &\leq \mathbb{P}[|\mathcal{A}^T R\mathcal{B} - x^T R\mathcal{B}| > \epsilon/2] + \mathbb{P}[|x^T R\mathcal{B} - v_R| > \epsilon/2] \end{aligned}$$

The second term of the latter expression can be bounded as follows from (10)

$$\mathbb{P}[|x^T R\mathcal{B} - v_R| > \epsilon/2] \leq e^{-t\epsilon^2/2}.$$

To bound the first term we note that $\mathbb{E}[X_i^T R\mathcal{B}| \mathcal{B}] = \sum_{i,j} R_{ij}\mathbb{E}[\mathbf{1}_{\{X=e_i\}}|\mathcal{B}(j)] = \sum_{i,j} R_{ij}x(i)\mathcal{B}(j) = x^T R\mathcal{B}$. Therefore, conditioned on the value of \mathcal{B} , a Chernoff bound on the sequence of random variables $Z'_i = X_i^T R\mathcal{B}$ gives

$$\mathbb{P}[|\mathcal{A}^T R\mathcal{B} - x^T R\mathcal{B}| > \epsilon/2|\mathcal{B}] \leq e^{-t\epsilon^2/2}$$

which implies

$$\mathbb{P}[|\mathcal{A}^T R\mathcal{B} - x^T R\mathcal{B}| > \epsilon/2] \leq e^{-t\epsilon^2/2}$$

Putting everything together we get

$$\mathbb{P}[|\mathcal{A}^T R\mathcal{B} - v_R| > \epsilon] \leq 2e^{-t\epsilon^2/2}$$

Similarly,

$$\mathbb{P}[|\mathcal{A}^T C\mathcal{B} - v_C| > \epsilon] \leq 2e^{-t\epsilon^2/2}$$

Choosing $t = \frac{4}{\epsilon^2}$ we get

$$\begin{aligned} \mathbb{P}[|x^T R\mathcal{B} - v_R| > \epsilon] &\leq e^{-8}, \\ \mathbb{P}[|\mathcal{A}^T C y - v_C| > \epsilon] &\leq e^{-8}, \\ \mathbb{P}[|\mathcal{A}^T R\mathcal{B} - v_R| > \epsilon] &\leq 2e^{-2} \\ \mathbb{P}[|\mathcal{A}^T C\mathcal{B} - v_C| > \epsilon] &\leq 2e^{-2} \end{aligned}$$

Therefore, by union bound

$$\begin{aligned} \mathbb{P}[(|\mathcal{A}^T R\mathcal{B} - v_R| > \epsilon) \vee (|x^T R\mathcal{B} - v_R| > \epsilon) \\ \vee (|\mathcal{A}^T C\mathcal{B} - v_C| > \epsilon) \vee (|\mathcal{A}^T C y - v_C| > \epsilon)] < 0.55. \end{aligned}$$

The latter implies (9). This completes the proof. We note that (1) and (5) were also proved to be true in [9] using a similar method. \square

3.2 An Algorithm

Let $\epsilon > 0$. The following algorithm computes an $\frac{3-\sqrt{5}}{2} + \epsilon$ approximate equilibrium.

1. Discretize $[0, 1]$ into the set $\mathcal{V} = \{0, \epsilon, 2\epsilon, \dots, k\epsilon\}$, where $k = k(\epsilon)$ is such that $k\epsilon \leq 1$ and $(k+1)\epsilon > 1$.

2. Guess a pair of values $v_R, v_C \in \mathcal{V}$ that are ϵ -close to the values of a Nash equilibrium to the two players (this is, of course, done by trying all $O(\frac{1}{\epsilon^2})$ pairs until the A, B, x, y sought by the algorithm below are found).

(a) Let $v_{max} = \max\{v_R, v_C\}$.

- (b) Find a multiset A of row player's pure strategies, and a multiset B of column player's pure strategies, both multisets of size $4/\epsilon^2$ so that the following is satisfied.

$$\alpha^T R\beta \geq v_R - 3\epsilon/2 \quad (11)$$

$$\alpha^T C\beta \geq v_C - 3\epsilon/2 \quad (12)$$

where α , respectively β , is the uniform distribution on the elements of the multiset A , respectively B .

- (c) Find x , a mixed strategy for the row player, and y , a mixed strategy for the column player, as any solution of the following linear program

$$\mathbf{L}(v_R, v_C, \alpha, \beta) : \quad \alpha^T R y \geq v_R - 3\epsilon/2 \quad (13)$$

$$\forall i : \quad e_i^T R y \leq v_R + \epsilon/2 \quad (14)$$

$$x^T R\beta \geq v_R - 3\epsilon/2 \quad (15)$$

$$\alpha^T C y \geq v_C - 3\epsilon/2 \quad (16)$$

$$\forall j : \quad x^T C e_j \leq v_C + \epsilon/2 \quad (17)$$

$$x^T C\beta \geq v_C - 3\epsilon/2 \quad (18)$$

- (d) If $v_{max} \geq 1/3$ then return the pair of strategies:

$$(\delta\alpha + (1-\delta)x, \delta\beta + (1-\delta)y)$$

$$\text{where } \delta = \delta(v_R, v_C) = \frac{3}{2} - \frac{1}{2v_{max}}.$$

- (e) If $v_{max} < 1/3$, then return the pair of strategies (x, y) .

THEOREM 3.2. *For every $\epsilon > 0$, the Algorithm returns an $\frac{3-\sqrt{5}}{2} + \epsilon$ approximate equilibrium, in time $n^{O(1/\epsilon^2)}$*

PROOF. The algorithm exhaustively searches through its guesses of v_R, v_C, α and β , and for each guess it checks if the system of equations $\mathbf{L}(v_R, v_C, \alpha, \beta)$ has a feasible solution. Since there are only $(1/\epsilon)^2 \times \binom{n}{1/\epsilon^2}^2$, i.e., $n^{O(1/\epsilon^2)}$ different choices for the tuple $(v_R, v_C, \alpha, \beta)$, so it is clear that, provided that the algorithm finds a solution, it terminates in time $n^{O(1/\epsilon^2)}$. Now, to prove that the algorithm indeed finds a solution, we have to show that there is some guess of $(v_R, v_C, \alpha, \beta)$ for which $\mathbf{L}(v_R, v_C, \alpha, \beta)$ has a feasible solution.

By Lemma 3.1, we know that there exists a tuple $(v_R^*, v_C^*, x, \alpha, y, \beta)$, with $|\text{supp}(\alpha)|, |\text{supp}(\beta)| \leq 4/\epsilon^2$, s.t. (1) – (8) hold. Let $v_R, v_C \in \mathcal{V}$ be s.t. $|v_R - v_R^*| \leq \epsilon/2$ and $|v_C - v_C^*| \leq \epsilon/2$ (these exist because of the discretization in \mathcal{V}). By (1) $\alpha^T R\beta \geq v_R^* - \epsilon \geq v_R - 3\epsilon/2$, proving (11). Similarly (12) is true as well. Hence $(v_R, v_C, \alpha, \beta)$ is one of the possible guesses of the algorithm. We will now show that $\mathbf{L}(v_R, v_C, \alpha, \beta)$ has a feasible solution.

Now (2) implies $\forall i \in \text{supp}(\alpha) : e_i^T R y = v_R^* \geq v_R - \epsilon/2$ which proves (13). Similarly (4) and (3) give (14) and (15) respectively. Similarly the column side inequalities also hold. Thus (x, y) is a feasible solution to $\mathbf{L}(v_R, v_C, \alpha, \beta)$, as required.

So it remains to determine the value of the approximation. Define val_R, val_C as the values obtained by the row and column player (resp.) in the solution output by the algorithm. Define def_R, def_C as the maximum value that the player can get by defecting to some pure strategy. Thus our algorithm outputs a $\max\{def_R - val_R, def_C - val_C\}$ approximate equilibrium. The analysis is in two cases.

In the easy case, when $v_{max} < 1/3$, the algorithm outputs (x, y) . By (14), we see that $def_R \leq v_R + \epsilon/2 < 1/3 + \epsilon/2$. Similarly $def_C < 1/3 + \epsilon/2$. Since $val_R, val_C \geq 0$, we have a $1/3 + \epsilon/2$ approximate equilibrium.

In the case when $v_{max} \geq 1/3$, we have, from (11), (13) and (15) that

$$\begin{aligned} val_R &\geq \delta^2(v_R - 3\epsilon/2) + 2\delta(1 - \delta)(v_R - 3\epsilon/2) \\ &= (v_R - 3\epsilon/2)(\delta^2 + 2\delta(1 - \delta)) \end{aligned}$$

From (14), and the fact that all the entries are at most 1, we get:

$$def_R \leq \delta + (1 - \delta)(v_R + \epsilon/2)$$

Hence

$$def_R - val_R \leq v_R g(\delta) + \epsilon f(\delta)$$

where $g(\delta) = \delta + (1 - \delta) - \delta^2 - 2\delta(1 - \delta)$ and $f(\delta) = (1 - \delta)/2 + 3/2(\delta^2 + 2\delta(1 - \delta))$

Similarly, we have

$$def_C - val_C \leq v_C g(\delta) + \epsilon f(\delta)$$

Assume, w.l.o.g., that $v_{max} = v_C$. Then, from above we see that $def_R - val_R \leq def_C - val_C$. Plugging in the choice of δ in the algorithm, we get $def_C - val_C \leq \frac{3}{2} - \frac{1}{4}(5v_{max} + \frac{1}{v_{max}}) + \epsilon f(\delta)$. This function has a maximum at $v_{max} = 1/\sqrt{5}$, where it equals $(3 - \sqrt{5})/2 + \epsilon f(\delta)$. This means that for every $\epsilon' > 0$, the algorithm returns an $(3 - \sqrt{5})/2 + \epsilon'$ approximate equilibrium. \square

4. REFERENCES

- [1] I. Althofer. On sparse approximations to randomized strategies and convex combinations. *Linear Algebra and its Applications*, 1994.
- [2] X. Chen and X. Deng. Settling the complexity of two-player nash equilibrium. *FOCS*, 2006.
- [3] X. Chen, X. Deng, and S.-H. Teng. Computing nash equilibria: approximation and smoothed complexity. *FOCS*, 2006.
- [4] C. Daskalakis, P. Goldberg, and C. Papadimitriou. The complexity of computing a nash equilibrium. *STOC*, 2006.
- [5] C. Daskalakis, A. Mehta, and C. Papadimitriou. A note on approximate nash equilibria. *WINE*, 2006.
- [6] T. Feder, H. Nazerzadeh, and A. Saberi. Personal communication, 2006.
- [7] P. Goldberg and C. Papadimitriou. Reducibility among equilibrium problems. *STOC*, 2006.
- [8] S. C. Kontogiannis, P. N. Panagopoulou, and P. G. Spirakis. Polynomial algorithms for approximating nash equilibria of bimatrix games. *WINE*, 2006.
- [9] R. Lipton, E. Markakis, and A. Mehta. Playing large games using simple strategies. *Electronic Commerce*, 2003.