Approximating Nash Equilibria using Small-Support Strategies

Tomas Feder
268 Waverley St.
Palo Alto, CA ,USA
tomas@theory.stanford.edu

Hamid Nazerzadeh
Stanford University
Stanford, CA ,USA
hamidnz@stanford.edu

Amin Saberi
Stanford University
Stanford, CA ,USA
saberi@stanford.edu

ABSTRACT

We study the problem of finding approximate Nash equilibria of two player games. We show that for any $0 < \epsilon < 1$, there is no $\frac{1}{1+\epsilon}$-approximate equilibrium with strategies of support $O(\frac{\log n}{\epsilon^2})$.

Categories and Subject Descriptors

F.2 [**Theory of Computation**]: Analysis Of Algorithms and Problem Complexity; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms

Algorithm, Theory, Economics

Keywords

Nash equilibrium, Small-support Strategies, Probabilistic methods

1. INTRODUCTION

Nash equilibrium is a central solution concept in game theory. In a game involving two or more players, a Nash equilibrium is a set of strategies, one for each player, such that no player can improve his or her payoff by changing strategy while the other players keep theirs unchanged. In 1950, Nash showed that Nash equilibria must exist for all finite games with any number of players [8]. Since then, this concept has been widely used for predicting the result of the interaction of independent and selfish agents in a conflict.

However, there are still many questions left about the predictiveness of this equilibrium concept related to its complexity. For a given game, how hard is it for the players to reach an equilibrium? Are there any natural dynamics that reach an equilibrium quickly? Are there always equilibrium points that can be found easily, e.g. in polynomial time? Similar questions can be asked about approximate equilibria, in which the players have a small incentive to change their strategy.

In order to address the above problems, there has been a lot of interest in understanding the computational complexity of finding a Nash equilibrium of a game. In a series of beautiful results, [6, 2, 3] proved that it is not possible to find an equilibrium or give a fully polynomial time approximation scheme for finding the equilibria of even a two person game unless PPAD is in P.

In the light of the above results, the question of finding approximate equilibria emerges as an important open problem. The focus of this paper is on two player or bimatrix games. In these games, there is a row and a column player with payoff matrices R and C respectively. We assume that the elements of R and C are non-negative. Let n and m denote the number of rows and columns of these matrices and assume that $n \geq m$.

An α-approximate equilibrium is a pair of mixed strategies (X, Y) with payoffs (x, y) such that if the best response to Y has payoff x' and the best response to X has payoff y' then $x \geq \alpha x'$ and $y \geq \alpha y'$.

One can use an additive notion for the approximation [7, 4]. Our results hold also for the additive notion as well.

A simple polynomial-time approximation algorithm for computing Nash equilibrium is a linear-time algorithm that finds a $\frac{1}{2}$-approximate equilibrium by examining all strategies with support of size at most 2 [4]. The *support of a strategy* is the set of pure strategies used to construct it. A natural question is whether it is possible to improve the factor $\frac{1}{2}$ by searching over strategies with larger support. Althöfer [1] (for zero-sum games) and Lipton *et al.* [7] showed that for any $0 < \alpha < 1$, there is always an α-additive-approximate Nash equilibrium with support of size $O(\frac{\log n}{\alpha^2})$.

We will show that the above two results are *asymptotically optimum.* We prove that it is not possible to improve the factor $\frac{1}{2}$, if the support of both players are less than or equal to $\log_2 n - 2\log_2 \log n$. This improves the factor $\frac{1}{4}$ given in [1, 4].

Furthermore, we prove that for any $0 < \epsilon < 1$, it is not possible to find a $\frac{1}{1+\epsilon}$-approximate equilibrium using strategies of support $O(\frac{\log n}{\epsilon^2})$. We also show that it is not possible to find a $\frac{1}{2-\epsilon}$-equilibrium even if we limit the support of one of the players to $\frac{\log n}{1-\log \epsilon}$. All our negative results apply to symmetric games, zero-sum games and $0-1$ bimatrix games as well. We use a simple probabilistic method for proving these Theorems. We show that a zero-one matrix whose elements are chosen uniformly at random, has the desired properties with high probability.

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

On the flip side, we give an algorithm that finds an ϵ-approximate Nash equilibrium for ϵ slightly bigger than $\frac{1}{2}$. In the solution of our algorithm, one of the players plays a pure strategy but the support of the strategy of the other player may be arbitrarily large. More recently, Daskalakis et al. [5], proposed an algorithm which computes a $0.38+\alpha$-additive-approximate equilibrium for a bimatrix game. The size of the support of their solution can be as big as $O(\frac{n}{\alpha^2})$.

2. RESULTS

THEOREM 1. *For any large enough n and $0 \le \epsilon \le 1$, there exists a constant-sum $0/1$ game of size n such that if we limit the column player to strategies with support of size less than $\frac{\log n}{1-\log\epsilon}$, it is not possible to achieve an α-approximate Nash equilibrium for $\alpha \ge \frac{1}{2-\epsilon}$.*

PROOF. Let $s = \lfloor \frac{\log n}{1-\log\epsilon} \rfloor$ and $m = \lfloor \frac{s}{\epsilon} \rfloor$. For such values for s and m we have $\binom{m}{s} < (\frac{en}{s})^s \le n$. Hence, we can generate an $n \times m$ matrix R having all possible rows consisting of s 1's and $m-s$ 0's. Obtain C from R by exchanging 0's and 1's. Suppose the column player plays at most s rows. Then some row has only 1's in these columns in R. Thus, the best response for column player has payoff 1. Furthermore, column player can get payoff $1-\frac{s}{m}$ by playing uniformly on all of the m columns, independent of the strategy of the column player. So the sum of the payoffs of the best responses of the players is $2-\frac{s}{m}$, while sum of their payoffs is 1. Therefore, $\frac{1}{2-\frac{s}{m}}$ approximation is not possible. $\square$

THEOREM 2. *Consider a zero-sum game where R is an n by n matrix with entries chosen uniformly at random from $\{0,1\}$ and C is the matrix obtained from R by exchanging 0 and 1. Then, with high probability, for any $\alpha > \frac{1}{2}$, no pair of strategies with supports of size smaller than $\log_2 n - 2\log_2 \log n$ has an α-approximate Nash equilibrium. Furthermore, for any $0 < \epsilon < 1$, with high probability, no pair of strategies with supports of size smaller than $O(\frac{\log n}{\epsilon^2})$ has a $\frac{1}{1+\epsilon}$-approximate Nash equilibrium.*

PROOF. Let $k = \log_2 n - 2\log_2 \log n$. We show with high probability, for any choice of k columns of R, there is a row of R which has all 1's in these k columns. Similarly, for any choice of k rows of C, some row of C has all 1's in these k rows. Therefore with high probability, the best response of each player has payoff 1. But the payoffs of the two players sum up to 1, so one of the players will have payoff at most $\frac{1}{2}$. Hence, there is no α-approximate Nash equilibrium with $\alpha > \frac{1}{2}$.

The probability that in some row of R, some k chosen columns will not be all 1 in that row is $1-2^{-k}$, so the probability that this will be the case in all n rows of R is $(1-2^{-k})^n$. Since $n2^{-k} > k\log n$, the probability that this will happen in some choice of k columns is $(1-2^{-k})^n n^k <<$ 1. This proves the first part of the Theorem.

For the second part, let S be the set of the rows in the support of row player. We first prove the theorem for the case that the distributions of the mixed strategies of the players are uniform over their supports and the size of the supports is k. For column j, let s_j be the sum of the entries of column j that are in S, e.g. $s_j = \sum_{i\in S} C_{ij}$. We claim with high probability, there exists a column j such that $s_j \ge (1+\epsilon)\frac{k}{2}$. Hence, the payoff of column player by choosing column j is at least $\frac{(1+\epsilon)}{2}$. Similarly for row player, with high probability, the best response has payoff at least $\frac{1+\epsilon}{2}$. But the sum of the payoffs of both players is 1. Therefore, there is no α-approximate Nash equilibrium.

To prove the claim, consider an arbitrary column j. Let $s = \lceil (1+\epsilon)k/2 \rceil$. By Sterling's formula, there exist constants c_1 and c_2 such that:

$$\begin{aligned} P[s_j \ge s] &\ge P[s_j = s] = \frac{1}{2^k}\binom{k}{s} \\ &\ge \frac{c_1}{\sqrt{k}} e^{-\frac{k}{2}((1+\epsilon)\log(1+\epsilon)+(1-\epsilon)\log(1-\epsilon))} \\ &\ge \frac{1}{\sqrt{k}} e^{-c_2 k\epsilon^2} \end{aligned}$$

Now we can choose a constant c such that $k = \frac{\log n}{c\epsilon^2}$, and $P[s_j \ge s] \ge n^{-1/2}$. So the probability that none of the columns has at least $(1+\epsilon)\frac{k}{2}$ 1's in these k rows is less than $(1-n^{-1/2})^n$. The probability that this happens in some choice of k columns is less than $(1-n^{-1/2})^n n^k << 1$. This proves the result for uniform strategies over the support.

For nonuniform strategies, we limit the strategy space to the strategies that the probability of playing any rows or columns is equal to $\frac{j}{k^2}$, for some integer $0 \le j \le k^2$. This gives an additive error of $\frac{1}{k}$ in approximation of s_j's, which is negligible for large n. The number of such strategies can be closely approximated by $(k^2)^k$. Without loss of generality, assume that rows in S are the rows 1 to k. Let $r = (r_1, \ldots, r_k)$ be the distribution of the mixed strategy. For $1 \le i \le k$, the i'th *cyclic permutation* of r is the mixed strategy with distribution $(r_{1+i}, \ldots, r_{k+i})$ where addition is in module k. Therefore, the payoff of column player by playing column j, when the row player strategy is the i'th cyclic permutation, is $\sum_{l=1}^{k} r_{l+i}c_{lj}$. Because the average payoffs of all of the k cyclic permutations is equal to s_j, at least one of these permutations guarantees payoff s_j for the column player. The entries of C are chosen independently at randomly, so by (1), $P[s_j \ge s] \ge \frac{1}{k\sqrt{k}} e^{-c_2 k\epsilon^2}$. So we can choose constant c' such that $k = \frac{\log n}{c'\epsilon^2}$, and $P[s_j \ge s] \ge n^{-1/2}$. The probability that the total payoff of the best responses of the players, for all pair of the strategies be less than $1+\epsilon$, is at most $2(1-P[s_j \ge s])^n (k^2)^k n^k = o(1)$ which completes the proof. $\square$

THEOREM 3. *Let R and C be arbitrary matrices of size at most n. There exists a function $f(n) = (2+o(1))^n$ such that for any $0 < \epsilon < \frac{1}{4nf(n)}$ there is a pure row strategy and a mixed column strategy that gives an α-approximate Nash equilibrium with $\alpha = \frac{1}{2}(1+\epsilon)$. This approximate equilibrium can be computed in polynomial time.*

PROOF. Without loss of generality, we assume $r_{11} = 1$ is the biggest element in R and the rest of the elements are non-negative. Let multi-set $S_1 = \{1\}$. Assume that row player chooses row 1 and column player plays the column in S_1 with probability $\frac{1}{2}(1-\epsilon)$ and plays column y_1 with probability $\frac{1}{2}(1+\epsilon)$, where c_{1y_1} is the highest entry in the first row. This guarantees α-approximation for column player with respect to the best response. If this pair of strategies also gives a similar guarantee to row player then we are done. Otherwise, there is another row, let us say row

2, which gives a payoff greater than $\frac{1}{2}(1+\epsilon)$ times the current payoff. Then row player chooses this row. Now column player chooses uniformly among the columns in multi-set $S_2 = S_1 \cup y_1$ with probability $\frac{1}{2}(1-\epsilon)$ and plays column y_2 with probability $\frac{1}{2}(1+\epsilon)$ where c_{2y_2} is the highest entry in row 2. While there exists a row that by playing it row player is better off by at least a factor α, she chooses this row and column player changes her strategy accordingly.

Without loss of generality assume that row player chooses the rows in the increasing order from 1 to n. Let μ_i be the average of the entries of row i in S_i, i.e. $\mu_i = \frac{1}{i}\sum_{j\in S_i} r_{ij}$. By induction, we show $\mu_i \geq 1 - f(i)\epsilon$. The basis clearly holds. Because row player preferred row $i+1$ to row i we have $\frac{(1-\epsilon)\mu_i}{(1-\epsilon)\mu_{i+1}+(1+\epsilon)y_{i+1}} < \frac{1}{2}(1+\epsilon)$. Also, y_{i+1} is at most 1 which implies $\mu_{i+1} > \frac{2}{1+\epsilon}\mu_i - \frac{1+\epsilon}{1-\epsilon} > 1-(2+o(1))f(i)\epsilon$.

Moreover, row player never chooses a row that she has chosen before. Note that if row player gives up playing row i, then μ_i is less than $\frac{1}{2}(1+\epsilon)$. If she is currently playing row k then

$$\begin{aligned}\frac{(1-\epsilon)\mu_k}{(1-\epsilon)\frac{i\mu_i+k-i}{k}+(1+\epsilon)} &\geq \frac{(1-f(k)\epsilon)}{1-\frac{i(1-\epsilon)}{2k}+\frac{(1+\epsilon)}{(1-\epsilon)}} \\ &\geq \frac{1-\frac{1}{4n}}{2-\frac{1-\epsilon}{n}+\frac{\epsilon^2}{1-\epsilon}} \\ &\geq \frac{1}{2}(1+\epsilon)\end{aligned}$$

With this observation row player will finally stay at one row and the players reach an α-approximate equilibrium. □

Acknowledgment. We would like to thank Mohsen Bayati, Costis Daskalakis, Aranyak Mehta, and Christos Papadimitriou for very useful discussions.

The last two authors are supported by NSF award 0546889 and a gift from Google.

3. REFERENCES

[1] I. Althöfer. On sparse approximations to randomized strategies and convex combinations. *Linear Algebra and its Applications, 199*, 1994.

[2] X. Chen, X. Deng. Settling the complexity of two-player nash equilibrium. In *Proceeding of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2006.

[3] X. Chen, X. Deng, and S. Teng. Computing Nash equilibria:approximation and smoothed complexity. In *Proceeding of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2006.

[4] C. Daskalakis, A. Mehta, C. Papadimitriou. A note on approximate Nash equilibria. In *Proceedings of the Second International Workshop on Internet and Network Economics (WINE)*, 2006.

[5] C. Daskalakis, A. Mehta, C. Papadimitriou. Progress in Approximate Nash equilibrium. In *Proceedings of 8th ACM Conference on Electronic Commerce (EC)*, 2007.

[6] C. Daskalakis, P. Goldberg, and C. Papadimitriou. The complexity of computing a Nash equilibrium. In *Proceedings of the 38th ACM Annual Symposium on Theory of Computing (STOC)*, 2006.

[7] R. Lipton, V. Markakis, A. Mehta. Playing large games using simple strategies. In *Proceedings of the 4th Annual ACM Conference on Electronic Commerce (EC)*, 2003.

[8] J. Nash. Non-Cooperative Games. *The Annals of Mathematics*, 54(2):286-295, 1950.