Playing Large Games Using Simple Strategies

Richard J. Lipton[*] *Evangelos Markakis*[†] *Aranyak Mehta*[†]

Abstract

We prove the existence of ϵ-Nash equilibrium strategies with support logarithmic in the number of pure strategies. Furthermore, the payoffs to all players in any (exact) Nash equilibrium can be ϵ-approximated by the payoffs to the players in some such logarithmic support ϵ-Nash equilibrium. These strategies are also uniform on a multiset of logarithmic size and therefore this leads to a quasi-polynomial algorithm for computing an ϵ-Nash equilibrium. To our knowledge this is the first subexponential algorithm for finding an ϵ-Nash equilibrium. Our results hold for any multiple-player game as long as the number of players is a constant (i.e., it is independent of the number of pure strategies). A similar argument also proves that for a fixed number of players m, the payoffs to all players in any m-tuple of mixed strategies can be ϵ-approximated by the payoffs in some m-tuple of constant support strategies.

We also investigate the question: when does a game have a "small" support (exact) equilibrium? We derive a sufficient condition for two person games: if the payoff matrices of a two person game have low rank then the game has a Nash equilibrium with small support. This also implies that if the payoff matrices can be well approximated by low rank matrices then the game has an ϵ-equilibrium with small support.

1 Introduction

Non-cooperative game theory has been extensively used to analyze situations of strategic interactions. Recently, it has been pointed out [20, 11, 24] that many internet related problems can be studied within the framework of this theory. The most important solution concept in non-cooperative games is the notion of Nash equilibrium.

In this paper we consider the following two issues concerning Nash equilibria:

First, it is currently not known if Nash equilibria can be computed efficiently. For two player games the known algorithms [8, 9, 10, 12, 13, 14, 16] either have exponential worst-case running time (in the number of available pure strategies) or it is unknown whether they run in polynomial time. For three player games, the problem seems to be even more difficult. Although for two player games it can be formalized as a Linear Complementarity

[*]Address: College of Computing, Georgia Institute of Technology, Atlanta, GA 30332–0280 and Telcordia Research, Morristown, NJ 07960. E-mail: rjl@cc.gatech.edu. Research supported by NSF grant CCR-0002299.

[†]Address: College of Computing, Georgia Institute of Technology, Atlanta, GA 30332–0280, E-mail: {vangelis,aranyak}@cc.gatech.edu

Problem (and hence some of the algorithms above) the problem for three player games is a Non-linear Complementarity Problem. Furthermore there exist examples of small three player games with rational payoff matrices in which all Nash equilibria are irrational. Algorithms for approximating equilibria in multiple player games (among others, [23, 30]) are also believed to be exponential. The problem of computing Nash equilibria has been of considerable interest in the computer science community and has been called one of the central open problems in computational complexity (Papadimitriou [20]). In fact it is known that the problem for two-person games lies in some class between P and NP [21]. It is also known that determining the existence of a Nash equlibrium with some additional natural properties (e.g. maximizing payoff sum, maximizing support) is NP-hard [5, 2]. For surveys on computational issues of Nash equilibria see [28, 17].

A second and related issue is the need to play simple strategies. Even if Nash strategies can be computed efficiently, they may be too complicated to implement. This has been pointed out, among others, by Simon [27] and later by Rubinstein [25] in the context of bounded rationality. Players tend to prefer strategies as simple as possible. They might prefer to play a sub-optimal strategy (with respect to rationality) instead of following a complex plan of action which might be difficult to learn or to implement. In this paper we consider normal form games and our notion of simple strategies is strategies which are uniform on a small support set. The importance of small support strategies becomes clear if we consider the pure strategies to be resources. In this case an equilibrium is almost impractical if a player has to use a mixed strategy which randomizes over a large set of pure strategies. The problem with the requirement of small strategies, of course, is that there exist games whose Nash equilibria are completely mixed (i.e., a player has to randomize over all his available pure strategies).

We address both these problems (namely, the need for efficient algorithms and the need for simple strategies), by using the weaker concept of ϵ-equilibrium (strategies from which each player has only an ϵ incentive to defect). More precisely:

Our main result (Section 3) is that for any two-person game there exists an ϵ-equilibrium with only logarithmic support (in the number of available pure strategies). Moreover the strategy of each player in such an equilibrium can be expressed in polylogarithmically many bits. In our opinion, this is an interesting observation on the structure of competitive behavior in various scenarios - namely, extremely simple approximate solutions exist. This result directly yields a quasi-polynomial ($n^{O(\ln n)}$, where n is the number of available pure strategies) algorithm for computing such an approximate equilibrium. To our knowledge this is the first subexponential algorithm for ϵ-equilibria. In addition to being small, our approximate equilibria provide both players with a good payoff too: the payoff that each player gets using these strategies is almost the same as that in some exact Nash equilibrium. Finally, our result holds not only for two person games but also for games in which the number of players is independent of the number of pure strategies. It is interesting to note that although the problem of finding exact equilibria seems to become more difficult in the "transition" from two player games to three and more, this is not the case for approximate equilibria. Computing ϵ-equilibria is important since they behave almost as well as exact Nash equilibria in several scenarios. In Section 3.2 we provide an interesting example based on the recent work by Vetta [29].

A second result (Section 3) is that if the players are allowed to communicate and "sign treaties" then there are *constant* support strategies which approximate the payoffs that each player gets in an equilibrium (in fact there are constant support strategies that approximate the payoffs of *any* pair of strategies). In real life, such treaties are not unknown (though often tacit) - this result can be considered as an explanation of why certain small strategies behave well and are used in real games, as opposed to a large and complicated Nash equilibrium.

A third question we investigate is: "when does a game have small support exact Nash equilibria?" In Section 4 we give a sufficient condition for two person games: if the payoff matrices of the players have low rank then there exists a Nash equilibrium with small support. This is a generalization of a result due to Raghavan [22]. We also prove that if the matrices can be well approximated by low rank matrices, then there exists an approximate equilibrium with small support.

The problem of looking for small support equilibria has been studied earlier. Koller and Megiddo [8] prove that for two person games in *extensive form* there exist equilibrium strategies whose support is at most the number of leaves of the game tree. However, not all games can be represented in the extensive form with a small number of leaves (where by small we mean logarithmic in the number of pure strategies). Our result guarantees the existence of equilibria with logarithmic support for any two person normal form game (and also for multiple players as stated above) but the equilibria are only approximate.

It should be noted that since Nash equilibria are fixed points of a certain map [19], ϵ-equilibria can be found using Scarf's algorithm [26], a general algorithm for finding approximate fixed points of continuous mappings. However, no sub-exponential upper bounds are known for approximating equilibria using this algorithm. In fact, Scarf's algorithm is known to take exponential time in the worst case for a general fixed point approximation ([6]).

For the class of two-person zero-sum games, results for approximate minmax strategies have been proved independently by Lipton and Young [15] and Althöfer [1]. In fact the proofs of Section 3 use the same technique (sampling). While [1] gives no details, the author claims that a similar result holds for non-zero sum two person games. The implication from approximate minmax strategies to ϵ-Nash equilibria which also approximate the payoffs in some exact Nash equilibrium does not seem to be direct. Furthermore our result holds for multiple player games too and not only for bimatrix games, which is interesting because multiple player games seem to be more difficult.

The remaining of the paper is structured as follows: In Section 2 we give the relevant definitions. In Section 3 we prove our main result. In Section 4 we prove that low rank payoff matrices imply the existence of equilibria with small support.

2 Notation and Definitions

Consider a two person game G, where for simplicity the number of available (pure) strategies for each player is n. We will refer to the two players as the row and the column player and we will denote their payoff matrices by R, C respectively. The results of Section 3 are also generalized for multiple person games in which the players do not have the same number of pure strategies.

A *mixed strategy* (or a randomized strategy) for a player is a probability distribution over the set of his pure strategies and will be represented by a vector $x = (x_1, x_2, \cdots, x_n)$, where $x_i \geq 0$ and $\sum x_i = 1$. Here x_i is the probability that the player will choose his ith pure strategy. If $x_i > 0$ we say that the mixed strategy x *uses* the ith pure strategy. The *support* of x ($Supp(x)$) is the set of pure strategies that it uses. A mixed strategy is called k*-uniform* if it is the uniform distribution on a multiset S of pure strategies, with $|S| = k$. For a mixed strategy pair x, y, the payoff to the row player is the expected value of a random variable which is equal to R_{ij} with probability $x_i y_j$. Therefore the payoff to the row player is (x, Ry), where $(.\,,.)$ denotes the inner product of two n-dimensional vectors. Similarly the payoff to the column player is (x, Cy).

The notion of a Nash equilibrium [19] is formulated as follows:

Definition 1 *A pair of strategies x^*, y^* is a Nash equilibrium point if:*

(i) For every (mixed) strategy $\bar{x}$ of the row player, $(\bar{x}, Ry^*) \leq (x^*, Ry^*)$, and

(ii) For every (mixed) strategy $\bar{y}$ of the column player, $(x^*, C\bar{y}) \leq (x^*, Cy^*)$

Similarly we can define ϵ-equilibria (this definition is well known in the literature):

Definition 2 *For any $\epsilon > 0$ a pair of mixed strategies x', y' is called an ϵ-Nash equilibrium point if:*

(i) For every (mixed) strategy $\bar{x}$ of the row player, $(\bar{x}, Ry') \leq (x', Ry') + \epsilon$ and

(ii) For every (mixed) strategy $\bar{y}$ of the column player, $(x', C\bar{y}) \leq (x', Cy') + \epsilon$

3 The Main Result

For the present we assume that all entries of R and C are between 0 and 1. Our main result is:

Theorem 1 *For any Nash equilibrium x^*, y^* and for any real number ϵ between 0 and 1, there exists, for every $k \geq \frac{12 \ln n}{\epsilon^2}$, a pair of k-uniform strategies x', y', such that:*

1. *x', y' is an ϵ-equilibrium,*
2. *$|(x', Ry') - (x^*, Ry^*)| < \epsilon$, (row player gets almost the same payoff as in the Nash equilibrium)*
3. *$|(x', Cy') - (x^*, Cy^*)| < \epsilon$, (column player gets almost the same payoff as in the Nash equilibrium)*

Proof :

The proof is based on the probabilistic method. For the given $\epsilon > 0$, fix $k \geq 12 \ln n/\epsilon^2$. Form a multiset A by sampling k times from the set of pure strategies of the row player, independently at random according to the distribution x^*. Similarly form a multiset B by sampling k times from the pure strategies of the column player, independently at random according to the distribution y^*.

Let x' be the mixed strategy for the row player which assigns probability $1/k$ to each member of A and 0 to other pure strategies. Let y' be the mixed strategy for the column player which assigns probability $1/k$ to each member of B and 0 to other pure strategies. Clearly, if a pure strategy occurs α times in the multiset, then it is assigned probability α/k.

Denote by x^i the ith pure strategy of the row player, and by y^j the jth pure strategy of the column player. In order to analyze the probability that x', y' is an ϵ-Nash equilibrium it suffices to consider only deviations to pure strategies.

We define the following events:

$$\begin{aligned}
\phi_1 &= \{\mid (x', Ry') - (x^*, Ry^*) \mid < \epsilon/2\} \\
\pi_{1,i} &= \{(x^i, Ry') < (x', Ry') + \epsilon\}, \qquad (i = 1, \cdots, n) \\
\phi_2 &= \{\mid (x', Cy') - (x^*, Cy^*) \mid < \epsilon/2\} \\
\pi_{2,j} &= \{(x', Cy^j) < (x', Cy') + \epsilon\}, \qquad (j = 1, \cdots, n) \\
GOOD &= \phi_1 \cap \phi_2 \bigcap_{i=1}^{n} \pi_{1,i} \bigcap_{j=1}^{n} \pi_{2,j}
\end{aligned}$$

We wish to show that $Pr[GOOD] > 0$. This would mean that there exists a choice of A and B such that the corresponding strategies x' and y' satisfy all three conditions in the statement of the theorem.

In order to bound the probabilities of the events ϕ_1^c and ϕ_2^c we introduce the following events:

$$\begin{aligned}
\phi_{1a} &= \{\mid (x', Ry^*) - (x^*, Ry^*) \mid < \epsilon/4\} \\
\phi_{1b} &= \{\mid (x', Ry') - (x', Ry^*) \mid < \epsilon/4\} \\
\phi_{2a} &= \{\mid (x^*, Cy') - (x^*, Cy^*) \mid < \epsilon/4\} \\
\phi_{2b} &= \{\mid (x', Cy') - (x^*, Cy') \mid < \epsilon/4\}
\end{aligned}$$

Note that $\phi_{1a} \cap \phi_{1b} \subseteq \phi_1$ hence $\phi_1^c \subseteq \phi_{1a}^c \cup \phi_{1b}^c$. The expression (x', Ry^*) is essentially a sum of k independent random variables each of expected value (x^*, Ry^*). Each such random variable takes value between 0 and 1. Therefore we can apply a standard tail inequality [7] and get:

$$Pr[\phi_{1a}^c] \leq 2e^{-k\epsilon^2/8}$$

Using a similar argument we have:

$$Pr[\phi_{1b}^c] \leq 2e^{-k\epsilon^2/8}$$

Therefore $Pr[\phi_1^c] \leq 4e^{-k\epsilon^2/8}$ and the same holds for the event ϕ_2^c.

In order to bound the probabilities of the events $\pi_{1,i}$'s and $\pi_{2,j}$'s we define the following auxilliary events:

$$\psi_{1,i} = \{(x^i, Ry') < (x^i, Ry^*) + \epsilon/2\}, \qquad (i = 1, \cdots, n)$$

$$\psi_{2,j} = \{(x', Ry^j) < (x^*, Ry^j) + \epsilon/2\}, \qquad (j = 1, \cdots, n)$$

We can easily see that

$$\psi_{1,i} \cap \phi_1 \subseteq \pi_{1,i}, \qquad (i = 1, \cdots n)$$

$$\psi_{2,j} \cap \phi_2 \subseteq \pi_{2,j}, \qquad (j = 1, \cdots n)$$

Hence for $i = 1, \cdots, n, j = 1, \cdots, n$:

$$Pr[\pi_{1,i}^c] \leq Pr[\psi_{1,i}^c] + Pr[\phi_1^c] \tag{1}$$

$$Pr[\pi_{2,j}^c] \leq Pr[\psi_{2,j}^c] + Pr[\phi_2^c] \tag{2}$$

Using the Hoeffding bound again we get:

$$Pr[\psi_{1,i}^c] \leq e^{-k\epsilon^2/2}$$

$$Pr[\psi_{2,j}^c] \leq e^{-k\epsilon^2/2}$$

Now by equations 1 and 2 we see that:

$$Pr[GOOD^c] \leq Pr[\phi_1^c] + Pr[\phi_2^c] + \sum_{i=1}^{n} Pr[\pi_{1,i}^c] + \sum_{j=1}^{n} Pr[\pi_{2,j}^c]$$

$$\leq 8e^{-k\epsilon^2/8} + 2n[e^{-k\epsilon^2/2} + 4e^{-k\epsilon^2/8}] < 1$$

Thus $Pr[GOOD] > 0$. □

Note that the strategies x', y' do not only form an ϵ-equilibrium, but they also provide both players a payoff ϵ-close to the payoffs they would get in some Nash equilibrium. In fact, the payoffs of every Nash equilibrium can be thus approximated by a small strategy ϵ-equilibrium. This provides another incentive for the players to remain in the ϵ-Nash equilibrium. Furthermore x', y' are k-uniform, which implies the following corollary:

Corollary 2 *For a 2-person game, there exists a quasi-polynomial algorithm for computing all k-uniform ϵ-equilibria (by Theorem 1 at least one such equilibrium exists).*

Proof : Given an $\epsilon > 0$, fix $k = \frac{12 \ln n}{\epsilon^2}$. By an exhaustive search, we can compute all k-uniform ϵ-equilibria (by Theorem 1 at least one such equilibrium exists; verifying ϵ-equilibrium condition is easy as we need to check only for deviations to pure strategies).

The running time of the algorithm is quasi-polynomial since there are $\binom{n+k-1}{k}^2$ possible pairs of multisets to look at. $\square$

To our knowledge this is the first subexponential algorithm for finding an approximate equilibrium. Furthermore, given the payoffs of any Nash equilibrium the algorithm can find an ϵ-Nash equilibrium in which both players receive payoffs ϵ-close to the given values.

When the entries of R and C are not between 0 and 1 the ϵ incentive to defect and the ϵ change in payoff both get magnified by $R_{max} - R_{min}$ for the row player and by $C_{max} - C_{min}$ for the column player. Here R_{max} and R_{min} denote the maximum and minimum entry of R, and similarly for C. Additionally if the players do not have the same number of pure strategies (say n_1, n_2) then the same result holds with $k \geq \frac{12 \ln max\{n_1,n_2\}}{\epsilon^2}$.

Our results can also be generalized to games with more than two players. In particular for an m-person game:

Theorem 3 *Let $s_1^*, \cdots, s_m^*$ be a Nash equilibrium in an m-person game. Let $p_1^*, \cdots, p_m^*$ be the payoffs to the players in the Nash equilibrium. Then for any real number ϵ between 0 and 1, there exists, for every $k \geq \frac{3m^2 \ln m^2 n}{\epsilon^2}$, a set of k-uniform strategies $s_1', s_2', \cdots, s_m'$, such that:*

1. *$s_1', s_2', \cdots, s_m'$ is an ϵ-equilibrium,*
2. *$|p_i' - p_i^*| < \epsilon$ for $i = 1, \cdots, m$*

where $p_1', \cdots, p_m'$ are the payoffs to the players if they play strategies s_i'.

As we see from Theorem 3 we can guarantee an ϵ-equilibrium with logarithmic support only when m is independent of n. It seems to us that the technique of sampling cannot help us prove a more general theorem than that. It is an interesting question to see whether this can be done using a different technique. However, it is still interesting that we can prove the existence of simple approximate equilibria even for three player games. This is so because the problem of finding exact equilibria for three player games seems to be much more difficult than for two player games due to irrational equilibria and non-linearity of the Complementarity Problem.

Corollary 2 also generalizes to games with a constant number of players since in this case the number of combinations of multisets that the algorithm has to look at is still quasi-polynomial . Again it would be interesting if a more general result could be proved.

3.1 Approximating Payoffs of Nash equilibria with Constant Support

In terms of the size of the support we can do much better, if we have weaker requirements. There may be applications in which we would not even insist on an approximate equilibrium. All we would care for is to approximate the payoffs in an actual Nash equilibrium. The next result is in that direction:

Theorem 4 *For any Nash equilibrium x^*, y^* and any ϵ between 0 and 1, there exists, for every $k \geq 5/\epsilon^2$, a pair of k-uniform strategies (x, y), such that*

1. $|(x, Ry) - (x^*, Ry^*)| < \epsilon$ *(row player gets almost the same payoff), and*
2. $|(x, Cy) - (x^*, Cy^*)| < \epsilon$ *(column player gets almost the same payoff),*

Again this result can be generalized to multiple player games. For an m-person game the support of the k-uniform strategies will be $O(m^2 \ln m)$.

Theorem 4 establishes the existence of *constant* support strategies which approximate the payoffs that both players get in a Nash equilibrium. The techniques used to prove this are the same as those used to prove Theorem 1, and the proof is omitted. Again, we assume that the entries of R and C are between 0 and 1 (in the general case we get a magnification by $R_{max} - R_{min}$ and $C_{max} - C_{min}$ as before). Note that Theorem 4 is true for any pair of strategies x^*, y^*, not necessarily for Nash equilibria.

A situation in which this result could be applicable is the following: Consider a game between two players both having a very large number of pure strategies at their disposal. Let v_1, v_2 be the payoffs in a Nash equilibrium to the row and column player respectively. If the support of the equilibrium strategies is very big, then it would be preferable for both players to sign a "bilateral treaty" (multilateral treaty in case of multiple player games) and use only a small number of strategies, as provided by the result. In that case, both players would still receive a payoff close to v_1 and v_2 respectively, while using a small number of strategies. Furthermore, each player will be able to check, during the game, if the other player has violated the treaty, in which case he can switch to any other strategy.

3.2 An Interesting Application

The fact that we can compute ϵ-equilibria in quasi-polynomial time is very important, as they are "almost as good as" exact equilibria in several scenarios. Thus the difficulty in computing exact equilibria can be conveniently sidestepped. As an example we note the recent result of Vetta [29] on the social performance of Nash equilibria. This line of research was initiated by Koutsoupias and Papadimitriou [11]. The setting is a multi-player game together with a social utility function. In this scenario we would like to know how suboptimal a Nash equilibrium can be in terms of maximizing the social utility function. Results of this flavor for traffic routing problems were given in [4, 11, 24]. Vetta [29] proves that in any valid utility system with a non-decreasing and submodular social utility function, any Nash equilibrium gives at least half of the social optimum.

In this context the fact that we can compute ϵ-equilibria in quasi-polynomial time can be useful. In particular a simple generalization of the proof of Theorem 3.4 in [29] to ϵ-Nash equilibria together with our Corollary 2 gives the following result (for details and definitions see [29]):

Corollary 5 *Consider a valid utility system with a non-decreasing submodular utility function. Let OPT be the maximum value of the utility function. Then, we can compute in time* *quasi-polynomial in the number of pure strategies, a set of mixed strategies which form an ϵ-Nash equilibrium and which provide a social payoff P, s.t $P \geq 1/2OPT - m\epsilon$, where m is the number of players.*

4 Low Rank Implies Small Support Exact Equilibria

In this section we investigate the question: when does a two person game have small support exact Nash equilibria? We show that if the payoff matrices have low rank then the game has a small support Nash equilibrium. Furthermore we show that if the payoff matrices can be approximated by low rank matrices then the game has a small support approximate equilibrium (where the approximation factor depends on how well the matrices can be approximated).

Theorem 6 can be seen as a generalization of [22] which deals with "completely mixed equilibria", i.e. equilibria which use all the pure strategies. The generalization is based on a careful Gaussian elimination type step.

Denote again by R, C the payoff matrices for the row and column player respectively. Suppose that R and C are $n \times n$ matrices. The results of this section can be easily generalized for $m \times n, (m \neq n)$ matrices.

Theorem 6 *Let x^*, y^* be a Nash equilibrium. If $rank(C) \leq n-k$, then there exists a mixed strategy x for the row player with $|Supp(x)| \leq n-k+1$ such that x, y^* is an equilibrium point. Similarly, if $rank(R) \leq n-k$, then there exists a mixed strategy y for the column player with $|Supp(y)| \leq n-k+1$ such that x^*, y is an equilibrium point. Furthermore the payoff that both players receive in the equilibria x, y^* and x^*, y is the same as in the initial equilibrium x^*, y^*.*

Proof : Let $v_1 = (x^*, Ry^*), v_2 = (x^*, Cy^*)$ be the payoffs to the players. Since $rank(C) \leq n-k$ there exist at least k linearly independent vectors $\pi_1, \pi_2, \cdots \pi_k$ which satisfy $C^T \pi_i = 0$, $i = 1, ..., k$. At least $k-1$ of them are also linearly independent of x^*. Let $\pi_1, \cdots \pi_{k-1}$ be these vectors. Let:

$$\begin{aligned} x^* &= (x_1^*, x_2^*, \cdots, x_n^*) \\ \pi_1 &= (\pi_{11}, \pi_{12}, \cdots, \pi_{1n}) \\ \pi_2 &= (\pi_{21}, \pi_{22}, \cdots, \pi_{2n}) \\ &\vdots \\ \pi_{k-1} &= (\pi_{k-1,1}, \pi_{k-1,2}, \cdots, \pi_{k-1,n}) \end{aligned}$$

Our goal is to obtain a vector $x = f(x^*, \pi_1, \cdots, \pi_{k-1})$ such that the following hold:

1. x has at least $k-1$ coordinates equal to 0.

2. All the coordinates of x are nonnegative.

3. $\sum_{i=1}^{n} x_i = 1$.

4. x, y^* is a Nash equilibrium.

We find x as a suitable linear combination of $x^*, \pi_1, \pi_2, ..., \pi_{k-1}$. The procedure to find x is shown in Figure 1.

In the procedure we essentially use one vector $\pi \in P$ in each iteration of the while loop to make at least one more coordinate of x equal to 0, while keeping x a probability distribution. We then discard that π from P. The "If $|P| \geq k - i$" block is used to preprocess P so that at the end of the if-block x has i 0's, $|P| = (k-1) - i$, and every $\pi \in P$ has a 0 in every coordinate in which x has a 0.

In each iteration, after the preprocessing step, we choose a vector π from P and compute our new vector x as a linear combination of the old x and π. We consider two cases depending on whether the entries of π sum up to 0 or not. The choices of λ in both cases ensure that x remains a probability distribution and that at least one more coordinate in x becomes 0. No zero coordinates of x become non-zero because all $\pi \in P$ have 0 in such a coordinate (by the preprocessing step). It can be verified that these λ exist in either case.

Note also that in every step the elements of P are linearly independent since we start with linearly independent vectors and add to π_i only multiples of π_{k_0} which is the vector then removed from P. Linear independence is essential to ensure that none of the vectors in P will ever become the zero vector and also that we can always find a λ as defined in the procedure. It can be verified that at the end of the procedure, conditions (1), (2) and (3) above are satisfied.

It remains to prove that x, y^* is a Nash equilibrium (condition (4)). To do this we need to prove that (a) $(x^i, Ry^*) \leq (x, Ry^*)$ and (b) $(x, Cy^j) \leq (x, Cy^*)$ for all $i, j = 1, \cdots, n$. Since x^*, y^* is a Nash equilibrium, $(x^i, Ry^*) \leq v_1$. Consider the vector Ry^*. For every pure strategy of the row player that is used in x^* the corresponding coordinate of Ry^* is equal to v_1 (otherwise the row player would have a better response to y^*). Furthermore by the procedure of Table 1, we can see that if a pure strategy is not used by x^* then it is also not used by x. Therefore $(x, Ry^*) = v_1$, which implies (a). For (b) notice that at the end of the procedure we have: $x = x^* + \mu_1\pi_1 + \cdots \mu_{k-1}\pi_{k-1}$ for some real numbers $\mu_1, \cdots \mu_{k-1}$. That is, $x = x^* + \eta$ where η belongs to the null space of C. Thus $(x, Cy^j) = (x^*, Cy^j) + (\eta, Cy^j) = (x^*, Cy^j)$. Also $(x, Cy^*) = (x^*, Cy^*) + (\eta, Cy^*) = (x^*, Cy^*)$. But $(x^*, Cy^j) \leq (x^*, Cy^*)$ implying (b). Hence x, y^* is an equilibrium and the players receive payoff v_1, v_2 respectively as in the initial Nash equilibrium x^*, y^*. □

It also follows from Theorem 6 that there exists a small support Nash equilibrium if there is a game with low rank payoff matrices strategically equivalent to R, C [18].

Definition 3 *For $n \times n$ matrices C, D, D is an ϵ-approximation of C if $C = D + E$, where $|E_{ij}| \leq \epsilon$ for $i, j = 1, ..., n$.*

$P := \{\pi_1, ..., \pi_{k-1}\}, x := x^*$
$i :=$ number of 0's in x
While $i < k-1$:
begin
  If $|P| \geq k-i$ (begin preprocessing step)
    For each j s.t. $x_j = 0$
      $Q := \{\pi_k : \pi_{k,j} \neq 0, \pi_k \in P\}$
      If $|Q| = 0$
        $P := P \backslash \pi_{k_0}$, for some $\pi_{k_0} \in P$
      If $|Q| = 1$
        $P := P \backslash \pi_{k_0}$, for $\pi_{k_0} \in Q$
      If $|Q| > 1$
        Choose any $\pi_{k_0} \in Q$
        $\pi_k := \pi_k - \pi_{k_0}(\pi_{k,j}/\pi_{k_0,j}) \quad [\forall \pi_k \in Q]$
        $P := P \backslash \pi_{k_0}$
  (end of preprocessing step)

  Choose any $\pi_{k_0} \in P$
  Normalize π_{k_0} so that $\sum_{j=1}^{j=n} \pi_{k_0,j}$ is either 0 or 1.
  If $\sum_{j=1}^{j=n} \pi_{k_0,j} = 0$
    $\frac{1}{\lambda} := max_{x_j \neq 0} \frac{\pi_{k_0,j}}{x_j}$
    $x := x - \lambda \pi_{k_0}$
  If $\sum_{j=1}^{j=n} \pi_{k_0,j} = 1$
    Let λ be such that $\frac{1+\lambda}{\lambda} = max_{x_j \neq 0} \frac{\pi_{k_0,j}}{x_j}$
    $x := (1+\lambda)x - \lambda \pi_{k_0}$
  For every $\pi_j \in P$, $\pi_j \neq \pi_{k_0}$: $\pi_j := \pi_j - \frac{\pi_{j,l}}{\pi_{k_0,l}} \pi_{k_0}$, where l is the index which achieved the max in either case.
  $P := P \backslash \pi_{k_0}$
  $i :=$ number of 0's in x

end while

Figure 1: Obtaining an equilibrium with (rank+1)-size support.

Lemma 7 *Let D be an ϵ-approximation of C. Let x^*, y^* be a Nash equilibrium for the game with payoff matrices R, D. Then x^*, y^* is a 2ϵ-Nash equilibrium for the game with payoff matrices R, C.*

Proof : Clearly $(x^*, Ry^*) \geq (\bar{x}, Ry^*)$, $\forall\, \bar{x}$. For any strategy $\bar{y}$:

$$(x^*, Cy^*) = (x^*, Dy^*) + (x^*, Ey^*) \geq (x^*, D\bar{y}) + (x^*, Ey^*)$$

Since $|E_{ij}| \leq \epsilon$, $\forall\, i, j$,

$$(x^*, E\bar{y}) - (x^*, Ey^*) \leq 2\epsilon$$

Hence,

$$(x^*, Cy^*) \geq (x^*, D\bar{y}) + (x^*, E\bar{y}) - 2\epsilon = (x^*, C\bar{y}) - 2\epsilon$$

□

Corollary 8 *For any game R, C, and for any $k < n$, if C can be ϵ-approximated by a rank k matrix then there exists a 2ϵ-equilibrium x, y with $|Supp(x)| \leq k + 1$. Similarly for R.*

Proof : This follows immediately from Theorem 6 and Lemma 7. □

In particular, we can use the Singular Value Decomposition to approximate the payoff matrices R, C by rank k matrices for any k. The approximation factor ϵ of Corollary 8 is then a function of the singular values of the matrices.

A useful corollary arises from the observation that for 2-person games, if we know the support of a Nash equilibrium, then we can compute the exact equilibrium strategies in polynomial time. This is because an equilibrium strategy y for the column player equalizes the payoff that the row player gets for every pure strategy in his support and vice versa. Hence we can write a linear program and compute the Nash equilibrium with the given support. The following is a direct consequence of this observation and Theorem 6.

Corollary 9 *If the payoff matrices R, C have constant rank, then we can compute an exact Nash equilibrium in polynomial time. In particular if one of the players has a constant number of pure strategies, we can compute a Nash equilibrium in polynomial time.*

5 Discussion

Another attempt to prove the results of Section 3 would be to approximate the vectors of a Nash equilibrium by vectors of small support. It is not difficult to see that we can approximate any probability distribution vector by a vector of logarithmic support in the l_∞ norm with error at most $1/\log n$. However, approximating an equilibrium x^*, y^* in this manner does not imply that the approximating vectors will form an ϵ-equilibrium, for any given fixed ϵ. On the other hand it can be shown that an ϵ-approximation in the l_1 norm does yield an ϵ-equilibrium, but such an approximation is not always possible (e.g. if the Nash strategies are the uniform distributions).

An interesting open question is whether we can generalize the results of Section 3 to games where the number of players is an increasing function of n. Another question would be to generalize the result so that the incentive to defect won't depend on the range of the payoff matrices (which can be much higher than the expected payoff in any equilibrium). Finally it would be interesting to extend the low-rank result of Section 4 to multiple player games.

Acknowledgements. We would like to thank George Karakostas, Dana Randall and Vijay Vazirani for valuable discussions, and Milena Mihail and Craig Tovey for useful pointers to the literature.

References

[1] Althöfer, I. On Sparse Approximations to Randomized Strategies and Convex Combinations. *Linear Algebra and Applications* 199, pp. 339-355, 1994.

[2] Conitzer, V., Sandholm, T. Complexity Results about Nash Equilibria. *Technical report CMU-CS-02-135*, 2002.

[3] Cottle, R. W., Pang, J., Stone, R. E. *The Linear Complementarity Problem* Academic Press, 1992.

[4] Czumaj, A., Voecking, B. Tight Bounds for Worst-Case Equilibria. *13th Annual Symposium on Discrete Algorithms*, pp. 413-420, 2002.

[5] Gilboa, I., Zemel, E. Nash and Correlated Equilibria: Some Complexity Considerations . *Games and Economic Behavior*, 1989.

[6] Hirsch, M. D., Papadimitriou, C. H., Vavasis, S.A. Exponential Lower Bounds for finding Brouwer Fixed Points. *Journal of Complexity* 5, pp. 379-416, 1989.

[7] Hoeffding, W. Probability Inequalities for Sums of Bounded Random Variables. *American Statistical Journal*, pp. 13-30, March 1963.

[8] Koller, D., Megiddo, N. Finding Mixed Strategies with Small Support in Extensive Form Games. *International Journal of Game Theory* 25, pp. 73-92, 1996.

[9] Koller, D., Megiddo, N., von Stengel, B. Fast Algorithms for Finding Randomized Strategies in Game Trees. *26th ACM Symposium on the Theory of Computing*, pp. 750-759, 1994.

[10] D. Koller, N. Megiddo, and B. von Stengel. Efficient computation of equilibria for extensive two-person games. *Games and Economic Behavior*, 14:2, pp. 247-259, 1996.

[11] Koutsoupias, E. Papadimitriou, C.,H. Worst case Equilibria. *16th Annual Symposium on Theoretical Aspects of Computer Science* pp. 404-413, 1999.

[12] Kuhn, H. W. An Algorithm for Equilibrium Points in Bimatrix Games. *Proceedings of the National Academy of Sciences* 47, pp. 1657-1662, 1961.

[13] Lemke, C. E. Bimatrix Equilibrium Points and Mathematical Programming. *Management Science* 11, pp. 681-689, 1965.

[14] Lemke, C. E., Howson, J. T. Equilibrium Points of Bimatrix Games. *Journal of the Society of Industrial and Applied Mathematics* 12, pp. 413-423, 1964.

[15] Lipton, R. J., Young, N. Simple Strategies for Zero-Sum Games with Applications to Complexity Theory. *26th ACM Symposium on the Theory of Computing* 26, pp. 734-740, 1994.

[16] Mangasarian, O. L. Equilibrium Points of Bimatrix Games. *Journal of the Society for Industrial and Applied Mathematics*, 12:4, pp. 778-780, 1964.

[17] McKelvey, R., McLennan, A. Computation of Equilibria in Finite Games. In Amman, H., Kendrick, D., Rust, J. eds *Handbook of Computational Economics* 1, 1996.

[18] Moulin, H., Vial, L. Strategically Zero-Sum Games: The Class of Games Whose Completely Mixed Equilibria Cannot be Improved Upon. *International Journal of Game Theory* 7, 3/4, pp. 201-221, 1978

[19] Nash, J. F. Non-Cooperative games. *Annals of Mathematics* 54, pp. 286-295, 1951.

[20] Papadimitriou, C., H. Algorithms, Games, and the Internet. *Proceedings of the 33th Annual ACM Symposium on the Theory of Computing* **33**, pp. 749-753, 2001.

[21] Papadimitriou, C., H. On the Complexity of the Parity Argument and Other inefficient Proofs of Existence. *Journal of Computer and System Sciences* **48**(3), 1994.

[22] Raghavan, T.E.S. Completely Mixed Strategies in Bimatrix Games. *Journal of London Math Society* 2(2), 709-712, 1970.

[23] Rosenmuller, J. On a Generalization of the Lemke-Howson Algorithm to Non-Cooperative Games. *SIAM Journal of Applied Mathematics* 21, pp. 73-79, 1971.

[24] Roughgarden, T., Tardos, E. How bad is Selfish Routing. *41st Annual Symposium on Foundations of Computer Science* , pp. 93-102, 2000.

[25] Rubinstein, A. *Modeling Bounded Rationality*, Cambridge, Mass. MIT Press, 1998.

[26] Scarf, H. The Approximation of Fixed Points of a Continuous Mapping. *SIAM Journal of Applied Mathematics* 15, pp. 1328-1343, 1967.

[27] Simon, H. *Models of Bounded Rationality*, Volume 2, Cambridge, Mass., MIT Press, 1982.

[28] von Stengel, B. Computing Equilibria for two-person Games. In Aumann, R. and Hart, S. eds, *Handbook of Game Theory* 3, 2002.

[29] Vetta, A. Nash Equilibria in Competitive Societies, with Applications to Facility Location, Traffic Routing and Auctions. *43rd Annual Symposium on Foundations of Computer Science* , 2002.

[30] Wilson, I. Computing Equilibria of N-Person Games. *SIAM Journal of Applied Mathematics* 21, pp. 80-87, 1971.