

Approximate Nash Equilibria for Multi-player Games [★]

Sébastien Hémon^{1,2}, Michel de Rougemont^{1,3}, and Miklos Santha^{1,4}

¹ CNRS-LRI, Univ. Paris-Sud, F-91405 Orsay

² LRDE-EPITA F-94276 Le Kremlin-Bicetre

³ Univ. Paris II F-75005 Paris

⁴ CQT, Nat. Univ. of Singapore, Singapore 117543

Abstract. We consider games of complete information with $r \geq 2$ players, and study approximate Nash equilibria in the additive and multiplicative sense, where the number of pure strategies of the players is n . We establish a lower bound of $r^{-1} \sqrt{\frac{\ln n - 2 \ln \ln n - \ln r}{\ln r}}$ on the size of the support of strategy profiles which achieve an ε -approximate equilibrium, for $\varepsilon < \frac{r-1}{r}$ in the additive case, and $\varepsilon < r-1$ in the multiplicative case. We exhibit polynomial time algorithms for additive approximation which respectively compute an $\frac{r-1}{r}$ -approximate equilibrium with support sizes at most 2, and which extend the algorithms for 2 players with better than $\frac{1}{2}$ -approximations to compute ε -equilibria with $\varepsilon < \frac{r-1}{r}$. Finally, we investigate the sampling based technique for computing approximate equilibria of Lipton et al. [12] with a new analysis, that instead of Hoeffding's bound uses the more general McDiarmid's inequality. In the additive case we show that for $0 < \varepsilon < 1$, an ε -approximate Nash equilibrium with support size $\frac{2r \ln(nr+r)}{\varepsilon^2}$ can be obtained, improving by a factor of r the support size of [12]. We derive an analogous result in the multiplicative case where the support size depends also quadratically on g^{-1} , for a lower bound g on the payoffs of the players at some given Nash equilibrium.

1 Introduction

Classical games of complete information with r players model situations where r decision makers interact and pursue well-defined objectives. A Nash equilibrium describes strategies for each player such that no player has any incentive to change her strategy. The algorithmic study of Nash equilibria started with the work of Lemke and Howson [11] in the 1960's, for the case of two players. This classical algorithm is exponential in the number of strategies (see [15]). Computing a Nash equilibrium is indeed not an easy task. It was proven recently that this computation is complete for the class PPA, first for $r \geq 4$ in [7], then for $r \geq 3$ in [6] and [2], and finally for $r \geq 2$ in [4]. Therefore it is unlikely to be feasible in polynomial time.

Approximate Nash equilibria have been studied both in the additive and the multiplicative models of approximation. An ε -approximate Nash equilibrium describes strategies for each player such that by changing her strategy unilaterally, no player can improve her gain by more than ε . Lipton et al. [12] studied additive approximate Nash equilibria for r -player games by considering small-support strategies, and obtained an approximation schema which computes an ε -approximate equilibrium in the additive sense, in time $n^{O(\frac{\ln n}{\varepsilon^2})}$, where n is the maximum number of pure strategies. It is known that there is no Fully Polynomial Time Approximation Schema for this problem [3], but it is open to decide if there is a PTAS. Daskalakis et al. [8] gave a simple algorithm for computing an additive $\frac{1}{2}$ -approximate equilibrium in 2-player games, using strategies with support at most 2. Feder et al. [10] showed that the factor $\frac{1}{2}$ was optimal when the size of the support could not

[★] Research supported by the European Commission IST Integrated Project Qubit Applications (QAP) 015848, and by the ANR Blanc AlgoQP grant of the French Research Ministry.

exceed $\log n - 2 \log \log n$. Breaking the $\frac{1}{2}$ barrier required approximation strategies with larger support size. In [9] Papadimitriou et al. have exhibited an additive $\frac{3-\sqrt{5}}{2}$ -approximate polynomial time algorithm, using linear programming. Further improvements for the approximation of the equilibrium in 2-player game were obtained by Bosse et al. [1] and Tsaknakis et al. [16], but the case of polynomial time approximation in games with more than 2 players was not investigated. The case of the multiplicative approximation has been studied by Chien and Sinclair [5] for dynamic strategies.

Here we study approximate Nash equilibria for r -player games, where the number of pure strategies of the players is n . First we extend the lower bounds on the factors of approximations for strategies with small support size. In Theorem 1 we prove that no ε -approximate equilibrium can be achieved with strategy profiles of support size less than $r^{-1} \sqrt{\frac{\ln n - 2 \ln \ln n - \ln r}{\ln r}}$ if $\varepsilon < \frac{r-1}{r}$ in the additive case, and $\varepsilon < r - 1$ in the multiplicative case.

Then we exhibit polynomial time algorithms for the additive approximation. Our results are based on the algorithm of Theorem 2 which extends approximations for r -player games to approximations for $(r + 1)$ -player games. As a consequence, we design in Corollary 3 a polynomial time algorithm which computes an $\frac{r-1}{r}$ -approximate equilibrium with support size at most 2, and in Corollary 4 extend the algorithms breaking the $\frac{1}{2}$ -approximation threshold in 2-player games into algorithms breaking the $\frac{r-1}{r}$ approximation threshold in r -player games.

Finally, we investigate the sampling based technique for computing approximate additive equilibria of Lipton et al. [12]. We propose a new analysis of this technique that instead of the Hoeffding's bound uses the more general McDiarmid's inequality [13] which enables us to bound the deviation of a function of independent random variables from its expectation. In Theorem 4 we show that for $0 < \varepsilon < 1$, an ε -approximate Nash equilibrium with support size $\frac{2r \ln(nr+r)}{\varepsilon^2}$ can be obtained, improving by a factor r the support size of [12]. We also establish a result analogous to the additive case in Theorem 5, where we show that for $0 < \varepsilon < 1$, a multiplicative ε -approximate Nash equilibrium with support size $\frac{9r \ln(nr+r)}{2g^2\varepsilon^2}$ can be achieved where g is a lower bound on the payoffs of the players at some given Nash equilibrium. In Remark 2 we argue that some dependence on g is necessary if we want the support of the approximate equilibrium to be included in the support of the given Nash equilibrium.

2 Preliminaries

For a natural number n , we denote by $[n]$ the set $\{1, \dots, n\}$. For an integer $r \geq 2$, an r -player game in normal form is specified by a set of pure strategies S_p , and a utility or payoff function $u_p : S \rightarrow \mathbb{R}$, for each player $p \in [r]$, where $S = S_1 \times \dots \times S_r$ is the set of pure strategy profiles. For $s \in S$, the value $u_p(s)$ is the payoff of player p for pure strategy profile s . Let $S_{-p} = S_1 \times \dots \times S_{p-1} \times S_{p+1} \times \dots \times S_r$, the set of all pure strategy profiles of players other than p . For $s \in S$, we set the partial pure strategy profile s_{-p} to be $(s_1, \dots, s_{p-1}, s_{p+1}, \dots, s_r)$, and for s' in S_{-p} , and $t_p \in S_p$, we denote by (s'_{-p}, t_p) the combined pure strategy profile $(s'_1, \dots, s'_{p-1}, t_p, s'_{p+1}, \dots, s'_r) \in S$. Let $B = \{e_1, \dots, e_n\}$ be the canonical basis of the vector space \mathbb{R}^n . We will suppose that each player has n pure strategies and that $S_p = B$, for all $p \in [r]$, and therefore $S = B^r$.

A *mixed strategy* for player p is a probability distribution over S_p , that is a vector $x_p = (x_p^1, \dots, x_p^n)$ such that $x_p^i \geq 0$, for all $i \in [n]$, and $\sum_{i \in [n]} x_p^i = 1$. We define $\text{supp}(x_p)$, the *support* of the mixed strategy x_p , as the set of indices i for which $x_p^i > 0$. Following [12], a mixed strategy x_p is called *k-uniform*, for some $k \in [n]$, if for every $i \in [n]$, there is an integer $0 \leq l \leq k$ such that $x_p^i = \frac{l}{k}$. Obviously, the size of the support of a k -uniform strategy is at most k . We denote by Δ_p the set of mixed strategies for p , and we call $\Delta = \Delta_1 \times \dots \times \Delta_r$ the set of *mixed strategy profiles*. For a mixed strategy profile $x = (x_1, \dots, x_r)$ we set $\text{supp}(x) = \text{supp}(x_1) \times \dots \times \text{supp}(x_r)$, and $\text{size}(x) = \max\{|\text{supp}(x_p)| : p \in [r]\}$. For a mixed strategy profile $x = (x_1, \dots, x_r)$ and pure strategy profile $s \in S$, the product $x_s = x_1^{s_1} x_2^{s_2} \dots x_r^{s_r}$ denotes the probability of s in x . We will consider the multilinear extension of the payoff functions from S to Δ defined by $u_p(x) = \sum_{s \in S} x_s u_p(s)$. The set Δ_{-p} , the partial mixed strategy profile x_{-p} for $x \in \Delta$, and the combined mixed strategy profile (x', x_p) for $x' \in \Delta_{-p}$ and $x_p \in \Delta_p$ are defined analogously to the pure case. The pure strategy s_p is a *best response* for player p against the partial mixed strategy profile x_{-p} if it maximizes $u_p(x_{-p}, \cdot)$. We will denote by $\text{br}(x_{-p})$ the set of best responses against x_{-p} .

A *Nash equilibrium* is a mixed strategy profile x^* such that for all $p \in [r]$, and for all $x_p \in \Delta_p$,

$$u_p(x_{-p}^*, x_p) \leq u_p(x^*).$$

An equivalent condition is $u_p(x_{-p}^*, s_p) \leq u_p(x^*)$ for every $s_p \in \text{br}(x_{-p}^*)$. Nash has shown [14] that for games with a finite number of players there exists always an equilibrium. It is immediate that the set of Nash equilibria is invariant by translation and positive scaling of the utility functions. Therefore we will suppose that they take values in the interval $[0, 1]$.

Several relaxations of the notion of equilibrium have been considered in the form of additive and multiplicative approximations. Let $\varepsilon > 0$. An *additive ε -approximate equilibrium* is a mixed strategy profile x^* such that for all $p \in [r]$, and for all $x_p \in \Delta_p$,

$$u_p(x_{-p}^*, x_p) \leq u_p(x^*) + \varepsilon.$$

A *multiplicative ε -approximate equilibrium* is a mixed strategy profile x^* such that for all $p \in [r]$, and for all $x_p \in \Delta_p$,

$$u_p(x_{-p}^*, x_p) \leq (1 + \varepsilon)u_p(x^*).$$

Since by our convention $0 \leq u_p(x^*) \leq 1$, a multiplicative ε -approximate equilibrium is always an additive ε -approximate equilibrium, but the converse is not necessarily true.

The input of an r -player game is given by the description of rn^r rational numbers. Here we will consider the computational model where arithmetic operations and comparisons have unit cost.

3 Inapproximability results for small support size

In [10] Feder, Nazerzadeh and Saberi have shown that there are 2-player games where for $\varepsilon < 1$, no multiplicative ε -approximation can be achieved with support size less than $\ln n - 2 \ln \ln n$. We generalize this result for r -player games in both models of approximation.

Theorem 1. *For $r \in o(n)$ there exists an r -player game such that no mixed strategy profile x with $\text{size}(x) < r^{-1} \sqrt{\frac{\ln n - 2 \ln \ln n - \ln r}{\ln r}}$ can be an additive ε -approximate equilibrium for $\varepsilon < \frac{r-1}{r}$, or a multiplicative ε -approximate equilibrium for $\varepsilon < r - 1$.*

Proof. We use the probabilistic method and will show that a random game from an appropriately chosen probabilistic space satisfies the claimed properties with positive probability. The space is defined as follows: for every pure strategy profile $s = (s_1, \dots, s_r) \in S$, choose a uniformly random $p \in [r]$ and set $u_p(s) = 1$ and $u_q(s) = 0$ for all $q \neq p$. This defines a random r -player 0/1 game with constant sum 1.

Fix $k < r^{-1} \sqrt{\frac{\ln n - 2 \ln \ln n - \ln r}{\ln r}}$, and set $S^{\leq k} = \{K_1 \times \dots \times K_r \subseteq S : |K_p| \leq k \text{ for } p \in [r]\}$.

Clearly $\text{size}(x) \leq k$ exactly when $\text{supp}(x) \in S^{\leq k}$. We define $S_{-p}^{\leq k}$ analogously. The event E_p is defined as follows: For all $K \in S_{-p}^{\leq k}$, there exists a pure strategy $t_p \in B$ such that for all $s' \in K$, we have $u_p(s', t_p) = 1$. Let $E = \bigwedge_{p \in [r]} E_p$. When E is realized, then for every x with $\text{size}(x) \leq k$, each player can increase her payoff to 1 by changing her strategy. Since the total payoff of the players is 1, at least one player has payoff at most $1/r$, and therefore x is not an additive ε -approximate equilibrium for $\varepsilon < \frac{r-1}{r}$, nor a multiplicative ε -approximate equilibrium for $\varepsilon < r - 1$.

We will prove that $\Pr[\overline{E}_p] < 1/r$ for all $p \in [r]$, and therefore $\Pr[E] > 0$. For fixed $K \in S_{-p}^{\leq k}$ and $t_p \in B$, the probability that there exists $s' \in K$ with $u_p(s', t_p) = 0$ is

$$1 - \Pr[\forall s' \in K \ u_p(s', t_p) = 1] \leq 1 - \frac{1}{r^{k^{r-1}}}.$$

Since the payoff functions are set independently, using the union bound we get

$$\Pr[\overline{E}_p] \leq \binom{n}{k}^{r-1} \left(1 - \frac{1}{r^{k^{r-1}}}\right)^n.$$

To prove the bound on $\Pr[\overline{E}_p]$ as claimed we bound the logarithm of the right hand side of the above inequality. This is at most

$$k(r-1) \ln n - \frac{n}{2r^{k^{r-1}}},$$

which can easily be seen to be no more than $-\ln r$ for the chosen value of k by rearranging, and taking logarithms again.

Corollary 1. *For $r \in O(1)$ there exists an r -player game such that for some constant $c > 0$, no mixed strategy profile x with $\text{size}(x) < c r^{-1} \sqrt{\ln n}$ can be an additive ε -approximate equilibrium for $\varepsilon < \frac{r-1}{r}$, or a multiplicative ε -approximate equilibrium for $\varepsilon < r - 1$.*

How essential are the restrictions on r and ε in Theorem 1? As we will show in the next section, for r fixed, the bound on ε is optimal in the case of additive approximation. The optimality of the bound for the multiplicative error remains open, and we don't know either if the restriction $r \in o(n)$ is necessary. Observe, however, that the case $r \geq n$ is anyhow of limited interest, since the uniform distribution on the pure strategies, for each players, is clearly an additive $\frac{n-1}{n}$ -approximation, and a multiplicative $(n-1)$ -approximation.

4 Polynomial time additive approximations

We know from the previous section that no strategy profile of constant support size can achieve a better than $\frac{r-1}{r}$ -approximate additive Nash equilibrium. We will prove here on the other hand that additive $\frac{r-1}{r}$ -approximate Nash equilibrium of constant support size does exist, and it can be computed in polynomial time. It is also shown that there are polynomial time computable additive η -approximate equilibria for some $\eta < \frac{r-1}{r}$. These results are based on an algorithm which extends any additive approximation for r -player games to an approximation for $(r + 1)$ -player games.

Theorem 2. *Given an algorithm \mathcal{A} that computes in time $q(r, n)$ an additive ε -approximate equilibrium for r -player games, there exists an algorithm \mathcal{A}' that computes in time $q(r, n) + O(n^{r+1})$ an additive $\frac{1}{2-\varepsilon}$ -approximate equilibrium for $(r + 1)$ -player games. Moreover, the support of the last player is of size at most 2, and the size of the support of the first r players is the same as in algorithm \mathcal{A} .*

Proof. Let s_{r+1} an arbitrary pure strategy of player $r + 1$. This induces an r -player game for the other players, assuming that player $p + 1$ is restricted to s_{r+1} . Algorithm \mathcal{A} finds for the induced game an additive ε -approximate equilibrium, say $x = (x_1, \dots, x_r)$. Compute now in time $O(n^{r+1})$ a pure strategy t_{r+1} for the last player which is in $\text{br}(x_1, \dots, x_r)$. Let us define the mixed strategy $x_{r+1} = \frac{1}{2-\varepsilon}s_{r+1} + \frac{1-\varepsilon}{2-\varepsilon}t_{r+1}$. We claim that $x^* = (x, x_{r+1})$ is an $\frac{1}{2-\varepsilon}$ -approximate equilibrium.

Consider any of the first r players. She can earn an additional payoff at most ε when player $r + 1$ plays s_{r+1} , and an additional payoff at most 1 when the chosen strategy is t_{r+1} . Therefore the overall gain by changing strategy is at most $\frac{\varepsilon}{2-\varepsilon} + \frac{1-\varepsilon}{2-\varepsilon} = \frac{1}{2-\varepsilon}$.

The last player has no way to increase her payoff when she plays her best response strategy t_{r+1} . Therefore her overall gain by changing strategy is at most $\frac{1}{2-\varepsilon}$.

Corollary 2. *Given an algorithm \mathcal{A} that computes in time $q(n)$ an additive ε -approximate equilibrium for 2-player games, there exists an algorithm that computes for any $r \geq 3$, in time $q(n) + O(n^r)$ an additive $\frac{(r-2)-(r-3)\varepsilon}{(r-1)-(r-2)\varepsilon}$ -approximate equilibrium for r -player games. Moreover, the supports of all but the first two players are of size at most 2, and the support size of the first two players is the same as in algorithm \mathcal{A} .*

Proof. We apply Theorem 2 inductively. Let ε_l be the approximation obtained for l -player games. Then $\varepsilon_2 = \varepsilon$ and $\varepsilon_{l+1} = \frac{1}{2-\varepsilon_l}$. Solving the recursion gives the result.

Corollary 2 never returns a better than $\frac{r-2}{r-1}$ -approximate Nash equilibrium. And, the procedure yields for r -players an additive ε -approximation with $\varepsilon \geq \frac{r-2}{r-1}$, only if the original two-player algorithm \mathcal{A} computes an additive η -approximation with $\eta \leq \frac{(r-2)-(r-1)\varepsilon}{(r-2)\varepsilon-(r-3)}$.

Corollary 3. *There exists an algorithm which computes an additive $\frac{r-1}{r}$ -approximate equilibrium for r -player games in time $O(n^r)$. Moreover the support of all players is of size at most 2.*

Proof. We apply Corollary 2 to the algorithm of [8] which computes in time $O(n^2)$ an additive $\frac{1}{2}$ -approximation for 2-player games with support size at most 2.

Let us stress here that though the complexity of the algorithm of Corollary 3 is exponential in r , it is sublinear in the input size.

Corollary 4. *There exist algorithms which in polynomial time compute an additive ε -approximate equilibrium for r -player games for some constant $\varepsilon < \frac{r-1}{r}$.*

Proof. Apply Corollary 2 to any of the polynomial time algorithms for 2-player games, such as [9], [1] or [16], which obtain an additive η -approximate equilibrium for some $\eta < \frac{1}{2}$. □

5 Subexponential time additive and multiplicative approximation

In one of the most interesting works on approximate equilibria, Lipton, Markakis and Mehta [12] have shown that for r -player games, for every $\varepsilon > 0$, there exists a k -uniform additive ε -approximation whenever $k > \frac{3r^2 \ln(r^2 n)}{\varepsilon^2}$. The result is proven by averaging, for all players, independent samples of pure strategies according to any Nash equilibrium.

Here we improve their bound by a factor r by showing that for $0 < \varepsilon < 1$, an additive ε -approximation exists already when $k > \frac{2r \ln(rn+r)}{\varepsilon^2}$. We also establish an analogous result for multiplicative ε -approximation when $k > \frac{9r \ln(rn+r)}{2g^2 \varepsilon^2}$, where g is a lower bound on the payoffs of the players at the equilibrium.

The proof is based on the probabilistic method and is analogous to the one given in [12]. The main difference is that instead of the Hoeffding's bound, we use the more general Mc Diarmid's inequality [13] which bounds the deviation of a function of several independent random variables from its expectation. It specializes to the Hoeffding's bound when the function is the sum of the variables. It is stated as follows:

Theorem 3 (McDiarmid). *Let Y_1, \dots, Y_m be independent random variables on a finite set A , and let $f : A^m \rightarrow \mathbb{R}$ be a function with the property that there exist real numbers c_1, \dots, c_m such that for all $(a_1, \dots, a_m, b) \in A^{m+1}$ and $1 \leq l \leq m$:*

$$|f(a_1, \dots, a_l, \dots, a_m) - f(a_1, \dots, b, \dots, a_m)| \leq c_l.$$

Then, for every $\varepsilon > 0$,

$$\Pr[f(Y_1, \dots, Y_m) - \mathbb{E}[f(Y_1, \dots, Y_m)] > \varepsilon] \leq e^{-\frac{2\varepsilon^2}{\sum_l c_l^2}}.$$

Theorem 4. *For all $0 < \varepsilon < 1$, and for all $k > \frac{2r \ln(rn+r)}{\varepsilon^2}$, there exists a k -uniform additive ε -approximate equilibrium.*

Proof. For every $p \in [r]$, let X_p^1, \dots, X_p^k be k copies of the random variable that takes the pure strategy $e_i \in B$ with probability x_p^i . We define $\mathcal{X}_p = \frac{1}{k} \sum_{j=1}^k X_p^j$, and let $\mathcal{X} = (\mathcal{X}_1, \dots, \mathcal{X}_r)$. Observe that $\mathbb{E}[u_p(\mathcal{X})] = u_p(x)$. For $p \in [r]$ and $i \in [n]$, we consider the events

$$E_p : |u_p(\mathcal{X}) - u_p(x)| < \frac{\varepsilon}{2},$$

$$F_p^i : |u_p(\mathcal{X}_{-p}, e_i) - u_p(x_{-p}, e_i)| < \frac{\varepsilon}{2},$$

and we define E as the conjunction of all them.

For every $p \in [r]$ and $i \in [n]$, the event F_p^i , the fact that x is a Nash equilibrium, and the event E_p imply that

$$|u_p(\mathcal{X}_{-p}, e_i) - u_p(\mathcal{X})| < \varepsilon.$$

Therefore, when E is realized, \mathcal{X} is an additive ε -approximate Nash equilibrium.

We prove that event E occurs with strictly positive probability. We start by bounding the probability of \overline{E}_p . We use McDiarmid's inequality with $m = rk$, when A is the canonical basis B , and the function f is defined as

$$f(a_1^1, \dots, a_1^k, \dots, a_r^1, \dots, a_r^k) = u_p \left(\frac{1}{k} \sum_{j=1}^k a_1^j, \dots, \frac{1}{k} \sum_{j=1}^k a_r^j \right).$$

Observe that $f(X_1^1, \dots, X_1^k, \dots, X_r^1, \dots, X_r^k) = u_p(\mathcal{X}_1, \dots, \mathcal{X}_r)$ and therefore

$$\mathbb{E}[f(X_1^1, \dots, X_1^k, \dots, X_r^1, \dots, X_r^k)] = u_p(x).$$

We claim that the values c_p^j can be chosen as $1/k$. Let a_p^j for $j \in [k]$ and $p \in [r]$ be some pure strategies. Fix $j \in [k]$, $p \in [r]$ and let b_p^j be a pure strategy. For $q \neq p$, we define the mixed strategies $\alpha_q = \frac{1}{k} \sum_{j=1}^k a_q^j$. Then, using the multilinearity of u_p , we have

$$f(a_1^1, \dots, a_p^j, \dots, a_r^k) - f(a_1^1, \dots, b_p^j, \dots, a_r^k) = \frac{1}{k} (u_p(\alpha_1, \dots, a_p^j, \dots, \alpha_r) - u_p(\alpha_1, \dots, b_p^j, \dots, \alpha_r)).$$

This implies the claim, because u_p takes values in $[0, 1]$. Since $\sum_{j,p} (c_p^j)^2 = kr \frac{1}{k^2} = \frac{r}{k}$, by McDiarmid's inequality we have

$$\Pr[\overline{E}_p] \leq e^{-\frac{\varepsilon^2 k}{2r}}.$$

For bounding from above the probability of \overline{F}_p^i , just observe that McDiarmid's inequality can be applied analogously for a function defined with $(r-1)k$ variables, and we get

$$\Pr[\overline{F}_p^i] \leq e^{-\frac{\varepsilon^2 k}{2(r-1)}}.$$

It follows from the union bound that

$$\Pr[\overline{E}] \leq r(n+1)e^{-\frac{\varepsilon^2 k}{2r}}.$$

The right side of this inequality is smaller than 1 when $k > \frac{2r \ln(rn+r)}{\varepsilon^2}$.

□

Theorem 5. *Let x be a Nash equilibrium for an r -player game and let $g > 0$ be a lower bound on the payoff of each player at the equilibrium. Then, for all $0 < \varepsilon < 1$, and for all $k > \frac{9r \ln(rn+r)}{2g^2 \varepsilon^2}$, there exists a k -uniform multiplicative ε -approximate equilibrium.*

Proof. The proof is a slight modification of the previous one. The random variable \mathcal{X} is defined identically. We set $\eta = 1 - \frac{1}{\sqrt{1+\varepsilon}}$ and $\zeta = \sqrt{1+\varepsilon} - 1$. The events E_p and F_p^i are defined as

$$\begin{aligned} E_p &: |u_p(\mathcal{X}) - u_p(x)| < \eta u_p(x), \\ F_p^i &: |u_p(\mathcal{X}_{-p}, e_i) - u_p(x_{-p}, e_i)| < \zeta u_p(x), \end{aligned}$$

and E as the conjunction of all them.

Recursively applying E_p , we get for every integer $m > 0$,

$$u_p(x) < \eta^m u_p(x) + u_p(\mathcal{X}) \sum_{l < m} \eta^l.$$

Therefore, using that $\frac{1}{1-\eta} = 1 + \zeta$, we have

$$u_p(x) \leq (1 + \zeta) u_p(\mathcal{X}).$$

The event F_p^i and the fact that x is a Nash equilibrium imply that

$$u_p(\mathcal{X}_{-p}, e_i) < (1 + \zeta) u_p(x).$$

Since $(1 + \zeta)^2 = 1 + \varepsilon$, it follows from the last two inequalities that \mathcal{X} is a multiplicative ε -approximate Nash equilibrium when E is realized.

Using that g is a lower bound for $u_p(x)$, by McDiarmid's inequality we get

$$\Pr[\overline{E}_p] \leq e^{-\frac{2g^2\eta^2k}{r}},$$

and

$$\Pr[\overline{F}_p^i] \leq e^{-\frac{2g^2\zeta^2k}{r-1}}.$$

As $\eta = \frac{\zeta}{1+\zeta}$ we have that $\eta < \zeta$. Also, it is not hard to see that $0 < \varepsilon < 1$ implies $\frac{\varepsilon}{3} < \eta$. Therefore

$$\Pr[\overline{E}] \leq r(n+1)e^{-\frac{2g^2\varepsilon^2k}{9r}},$$

and $\Pr[E] > 0$ when $k \geq \frac{9r \ln(rn+r)}{2g^2\varepsilon^2}$.

□

Remark 1. The condition $\varepsilon < 1$ in Theorem 5 is not a real restriction, since when $\varepsilon \geq 1$ then $\eta > \frac{1}{4}$, and therefore there exists a k -uniform multiplicative ε -approximate equilibrium for $k > \frac{8r \ln(rn+r)}{g^2}$.

Remark 2. If we require in Theorem 5 that the support of the approximate equilibrium is a subset of the support of the Nash equilibrium, the dependence on g is indeed necessary. Consider the following two players game given in the standard bimatrix representation where the number of the pure strategies of the players is $2n$:

$$M_1 = \begin{pmatrix} O_n & \frac{1}{2}I_n \\ \frac{1}{n}I_n & A_{n \times n} \end{pmatrix} \quad M_2 = \begin{pmatrix} I_n & O_n \\ O_n & I_n \end{pmatrix}.$$

Here, O_n denotes the $n \times n$ matrix with everywhere 0's, I_n is the $n \times n$ identity matrix, and $A_{n \times n}$ is the $n \times n$ matrix with everywhere $1/n$ except on its diagonal where all entries are 0. The game has a Nash equilibrium $x = (x_1, x_2)$ where $x_1 = x_2 = \frac{1}{n} \sum_{i=1}^n e_{n+i}$. The payoffs of the first and second player are respectively $u_1(x, y) = \frac{1}{n} - \frac{1}{n^2}$ and $u_2(x_1, x_2) = \frac{1}{n}$ and therefore, the minimum of the payoffs is $g = \Theta(1/n)$. Let $0 < \varepsilon < 1$, and let $y = (y_1, y_2)$ be a multiplicative ε -approximate Nash equilibrium. Let k denote the size of $\text{supp}(y_2)$, we claim that $k \geq \frac{n}{2(1+\varepsilon)}$. For this, observe first that $u_1(y_1, y_2) \leq 1/n$. Since $\text{supp}(y_2) \subseteq \{n+1, \dots, 2n\}$, there exists $i \in [n]$ such that $y_2^{n+i} \geq 1/k$, and therefore $u_1(e_i, y_2) \geq \frac{1}{2k}$. Since y is a multiplicative ε -approximate equilibrium, we have that $\frac{1}{2k} \leq \frac{1+\varepsilon}{n}$ and the statement follows. Observe on the other hand that there exists, for any $\varepsilon > \frac{2}{n-2}$, multiplicative ε -approximate equilibria with support size only 2 if we let the support be outside this of the Nash equilibrium.

In [12] it was already observed that when the number of players is constant, the sampling method yields an additive ε -approximation, for all constant $\varepsilon > 0$, in time $n^{O(\ln n)}$. When $g = \Omega(1)$, Theorem 5 implies a similar result for the multiplicative approximation. This condition is satisfied for example if all the utility functions are bounded from below by a constant.

Corollary 5. *If in an r -player game, where r is constant, there exists a Nash-equilibrium at which all the players payoffs are bounded from below by a constant then for all constant $\varepsilon > 0$, a multiplicative ε -approximation can be found in time $n^{O(\ln n)}$.*

It could be interesting to compare the complexities of the two additive approximation algorithms based on the Lipton, Markakis and Mehta sampling technique. Let $\mathcal{A}(r)$ be the additive ε -approximation r -player algorithm based on Theorem 4 which searches exhaustively through all the $\frac{2r^2 \ln(nr+r)}{\varepsilon^2}$ -uniform strategies. Let $\mathcal{B}(r)$ be the r -player algorithm we obtain by applying the iterative construction technique of Corollary 2 to $\mathcal{A}(2)$. Since by this technique we will never obtain a better than $\frac{r-2}{r-1}$ -approximation, let us fix some $\varepsilon > \frac{r-2}{r-1}$. The overall running time of $\mathcal{A}(r)$ is $O(n^{\frac{2r^2 \ln(nr+r)}{\varepsilon^2}})$ since the search is applied to the r players independently. The complexity of algorithm $\mathcal{B}(r)$ is $O(n^{\frac{r \ln(nr+r)}{\zeta^2}} + n^r)$ where $\zeta = \frac{(r-2)-(r-1)\varepsilon}{(r-2)\varepsilon-(r-3)}$. A simple computation shows that for all $r \geq 3$, algorithm $\mathcal{B}(r)$ has a smaller complexity.

Obviously, $\mathcal{A}(r)$ and $\mathcal{B}(r)$ are not the fastest algorithms when Corollary 4 yields a polynomial time procedure for computing an additive approximate Nash equilibrium. A simple analysis shows that for each two-player polynomial time η -approximation, when $\eta > 0$, Corollary 4 gives a polynomial time algorithm computing an additive $(\frac{r-2}{r-1} + \eta')$ -approximate equilibrium in an r -player game, for some $\eta' > 0$. This implies that currently, when ε is in some right neighborhood of $\frac{r-2}{r-1}$, algorithm $\mathcal{B}(r)$ is the most efficient for computing an additive ε -approximate Nash equilibrium.

6 Conclusion and open problems

In this paper, we have started the study of approximate Nash equilibria for r -player games when $r > 2$. The main open problem, just as in the two-player case, is the existence of a PTAS. We enumerate a few other, possibly much simpler problems, left also open:

1. Does the lower bound of Theorem 1 on the size of the strategy profiles hold also in the case when $r = cn$, for a constant $c \leq 1$?
2. Can we reduce the gap on the support size between the lower bound of Theorem 1 and the upper bound of Theorems 4 and 5 ? For example, when $r = \Theta(1)$, the lower bound is $\Omega(\sqrt[r-1]{\ln n})$ and the upper bound is $O(\ln n)$. When $r = 2$, these bounds are tight.
3. Is there a polynomial time algorithm which computes a multiplicative $(r - 1)$ -approximate Nash equilibrium ?

References

1. H. Bosse, J. Byrka, and E. Markakis. New algorithms for approximate Nash equilibria in bimatrix games. *International Workshop on Internet and Network Economics*, 2002.
2. X. Chen and X. Deng. 3-NASH is PPAD-complete. *Electronic Colloquium in Computational Complexity*, 2005.
3. X. Chen and X. Deng. Computing Nash equilibria: approximation and smoothed complexity. *IEEE Symposium on Foundations of Computer Science*, 2006.
4. X. Chen and X. Deng. Settling the complexity of two-player Nash equilibrium. *IEEE Symposium on Foundations of Computer Science*, 2006.
5. S. Chien and A. Sinclair. Convergence to approximate Nash equilibria in congestion games. *Proceedings of the eighteenth annual ACM-SIAM Symposium On Discrete Algorithms*, 2007.
6. C. Daskalakis, P. Goldberg, and C. Papadimitriou. Three players games are hard. *Electronic Colloquium on Computational Complexity*, 2005.
7. C. Daskalakis, P. Goldberg, and C. Papadimitriou. The complexity of computing a Nash equilibrium. *ACM Symposium on Theory of Computing*, 2006.
8. C. Daskalakis, A. Mehta, and C. Papadimitriou. A note on approximate Nash equilibria. *International Workshop on Internet and Network Economics*, 2006.
9. C. Daskalakis, A. Mehta, and C. Papadimitriou. Progress in approximate Nash equilibria. *ACM Conference on Electronic Commerce*, 2007.
10. T. Feder, H. Nazerzadeh, and A. Saberi. Approximating Nash equilibria using small-support strategies. *ACM Conference on Electronic Commerce*, 2007.
11. C.E. Lemke and J.T. Jr. Howson. Equilibrium points of bimatrix games. *Journal of the Society of Industrial and Applied Mathematics*, 1964.
12. R. Lipton, E. Markakis and A. Mehta. Playing large games using simple strategies. *ACM Conference on Electronic Commerce*, 2003.
13. C. McDiarmid. On the method of bounded differences. *Surveys in combinatorics, London Math. Soc. Lectures Notes 141, Cambridge Univ. Press*, 1989.
14. J.F. Nash. Non-cooperative games. *Annals of mathematics*, 54, 1951.
15. R. Savani and B. von Stengel. Exponentially many steps for finding a Nash equilibrium in a bimatrix game. *IEEE Symposium on Foundations of Computer Science*, 2004.
16. H. Tsaknakis and P. Spirakis. An optimization approach for approximate Nash equilibria. *International Workshop on Internet and Network Economics*, To appear.