

Complexity and Optimality of the Best Response Algorithm in Random Potential Games

Stéphane Durand¹ and Bruno Gaujal²(✉)

¹ Univ. Grenoble Alpes, 38000 Grenoble, France
stephane.durand@inria.fr

² Inria, Grenoble, France
bruno.gaujal@inria.fr

Abstract. In this paper we compute the worst-case and average execution time of the Best Response Algorithm (BRA) to compute a pure Nash equilibrium in finite potential games. Our approach is based on a Markov chain model of BRA and a coupling technique that transform the average execution time of this discrete algorithm into the solution of an ordinary differential equation. In a potential game with N players and A strategies per player, we show that the worst case complexity of BRA (number of moves) is exactly NA^{N-1} , while its average complexity over random potential games is equal to $e^\gamma N + O(N)$, where γ is the Euler constant. We also show that the effective number of states visited by BRA is equal to $\log N + c + O(1/N)$ (with $c \leq e^\gamma$), on average. Finally, we show that BRA computes a pure Nash Equilibrium faster (in the strong stochastic order sense) than any local search algorithm over random potential games.

1 Introduction

The question of computing Nash Equilibria (NE) in games is a central question in algorithmic game theory and has been investigated of many papers. The most classical result is in [1], showing that the problem of computing NE in finite games is PPAD complete.

Potential games have been introduced in [2] and have proven very useful, especially in the context of routing problems in networks, first mentioned in [3] and exhaustively studied ever since, in the transportation as well as computer science literature, see for example [4–6]. They have also been heavily investigated in the context of distributed optimization (see for example [7]). In [8, 9] the authors show that the computation of NE for general potential games is PLS complete (Polynomial Local Search complete). As for PPAD, this complexity class is believed to be different from P .

The best response dynamics is one of the most basic tool in game theory. The original proof of the existence of a Nash Equilibrium by Nash [10] can be seen as the proof of existence of a fixed point of the best response correspondence (best response is called *countering* in [10]). It has been well-known for a long time that

the Best Response Algorithm converges in finite time to a pure NE in potential games [11]. So BRA is a natural candidate for computing Nash equilibria.

In this paper, we analyze the performance of BRA over potential games with N players, each with A possible strategies. It is well known that the convergence time of BRA over potential games can be exponential in the number of players (see for example [8]). Here, we confirm this by showing that the worst case complexity of BRA (number of plays) is exactly NA^{N-1} . Special cases, such as graphical potential games have been analyzed in [12] by showing an equivalence between the potential of such games and Markov fields. In other special cases such as scheduling congestion games with identical tasks, it has been shown that BRA takes at most N steps before finding a NE [13]. Extensions with positive and negative externalities also have a linear complexity [14].

However the *average* complexity of BRA over all potential games has attracted surprisingly little attention. Random (non potential) games with two players have been studied in [15]: With two IID utility matrices of size $A \times A$, the computation of a NE is $O(A^3 \log \log A)$ with high probability using a rather sophisticated algorithm.

Our main contribution is to show that for potential games with N players, $\mathbb{E}[M_{BRA}]$, the average number of strategy profiles visited by BRA before convergence, is $\mathbb{E}[M_{BRA}] = \log(N) + C + O(1/N)$ (where $C \leq e^\gamma$, γ being the Euler constant). We also show that the average number of comparisons performed by the algorithm is equal to $e^\gamma(A-1)(N-1) + o(AN)$. This could be intuitively explained by the fact that random potential games have a lot of pure NE [16]. In our framework, potentials are IID random variables so that, on average, one action profile out of $(A-1)N + 1$ profiles is a NE while in the worst case, a potential game may have a single pure NE. This is only a partial explanation, however. This does not explain the fact that the complexity does not depend on the number of actions, nor the value of the constant factor, $e^\gamma \approx 1.78$.

We further show that the Best Response Algorithm computes a pure Nash Equilibrium faster than any algorithm based on player's local information, not only in average but also in the strong stochastic order sense.

Missing proofs and additional details (numerical simulations, analysis of alternative algorithms) are given in a long version of this paper, available in HAL Archive [17].

1.1 Coupling and Markovian Analysis

The main idea of our approach is to see the evolution of BRA in a random environment as a dynamical system, whose behavior can be computed using differential equations. This will allow us to compute the exact asymptotics of the average complexity in N and A , not only $O(\cdot)$ bounds. Second moments of T_{BRA} and of M_{BRA} can also be computed by the same approach (see [17]).

The first step (Sect. 4.2) is to construct an approximation of the behavior of BRA over a potential game. This approximation is called IFA in the paper, for Intersection-Free Approximation because it discards strategies already explored

by BRA. We show that the execution time of BRA is smaller than the execution time of its IFA approximation for the strong stochastic order. This is done by constructing a non-trivial *coupling* between both executions. This powerful technique is exploited to our great benefit here.

The second and most important step (Sect. 4.4) is to consider one run of the IFA approximation of BRA as a trajectory of a Markov chain over the continuous space of potentials. Doing so, the average complexity is transformed into the average hitting time of an absorbing state of the Markov chain. The theory of Markov chains implies that this average hitting time satisfies a Poisson differential equation. Thus, the average complexity of BRA is given by the solution of a system of ordinary differential equations. This system happens to have a solution in closed form whose asymptotics in N and A can be computed by taking integrals over initial states.

As for the proof of optimality of BRA among all local search algorithms (Sect. 5), our approach is based once again on a coupling argument. While using coupling techniques is more classical in this context (comparison of algorithms), this particular case retains some originality because the coupling used here is not built off-line but is being constructed on the fly while the algorithm runs.

2 Best Response Algorithm and Potential Games

We consider a game with a finite number N of players, each with A strategies.

Definition 1 (N -player game). A game is a tuple $\mathfrak{G} \stackrel{\text{def}}{=} \mathfrak{G}(N, \mathcal{A}, u)$ with

- a finite set of **players** $\mathcal{N} = \{1, \dots, N\}$;
- a finite set \mathcal{A}_k of pure **strategies** for each player $k \in \mathcal{N}$.

The set of **strategy profiles** or **states** of the game is $\mathcal{A} \stackrel{\text{def}}{=} \mathcal{A}_1 \times \mathcal{A}_2 \times \dots \times \mathcal{A}_N$.

- The players' **payoff** functions $u_k : \mathcal{A} \rightarrow \mathbb{R}$.

We define the *best response correspondence* $\mathbf{br}_k(x)$ as the set of all strategies that maximizes the payoff for player k under profile $x = (x_1, \dots, x_N)$: $\mathbf{br}_k(x) \stackrel{\text{def}}{=} \left\{ \operatorname{argmax}_{\alpha \in \mathcal{A}_k} u_k(\alpha; x_{-k}) \right\}$.

A *Nash equilibrium* (NE) is a fixed point of this correspondence, i.e. a profile x^* such that $x_k^* \in \mathbf{br}_k(x^*)$ for every player k .

Definition 2 (Potential games and its generalizations). A game is an (exact) potential game [11] if it admits a function (called the potential) $\Phi : \mathcal{A} \rightarrow \mathbb{R}$ such that for any player k and any unilateral deviation of k from strategy profile x to x' : $u_k(\alpha, x_{-k}) - u_k(\alpha', x_{-k}) = \Phi(\alpha, x_{-k}) - \Phi(\alpha', x_{-k})$.

A game is a generalized ordinal potential game [11] (or G -potential game for short) if there exists a potential function $\Phi : \mathcal{A} \rightarrow \mathbb{R}$ such that, for any player k and any state x , $u_k(\alpha, x_{-k}) > u_k(\alpha', x_{-k}) \Rightarrow \Phi(\alpha, x_{-k}) > \Phi(\alpha', x_{-k})$.

A game is a best-response potential game [18] (or BR -potential game for short) if there is $\Phi : \mathcal{A} \rightarrow \mathbb{R}$ such that for any player k and strategy profile x ,

$$\mathbf{br}_k(x) = \left\{ \operatorname{argmax}_{\alpha \in \mathcal{A}_k} \Phi(\alpha, x_{-k}) \right\}.$$

As shown in [18], exact potential games are BR-potential games, but there exist G-potential games that are not BR-potential games. In the following, we will consider the most general case (*i.e.* all games that are either BR-potential or G-potential games) and call them potential games for simplicity.

We consider a general version of the *Best Response Algorithm* (BRA) with uniform choice over all possible best responses when ties occur and where the next player is selected according to a *revision function* $R(\cdot)$, that may depend of the whole past of the algorithm. We assume that this function is *weakly fair*: each player appears infinitely often in the sequence of plays induced by R , almost surely. This revision function can be deterministic (for example, round-robin: $R(t) = t \bmod N$) or random (for example, Bernoulli where the next player is chosen according to an probability distribution ρ (the *revision law*): $\forall k \in \mathcal{N}, \mathbb{P}(R(t) = k) = \rho_k$). In that case, weak fairness implies that the probability of choosing any player k is strictly positive ($\forall k \in \mathcal{N}, \rho_k > 0$).

Algorithm 1. Best Response Algorithm (BRA)

Input: Game utilities ($u_k(\cdot)$); Initial state ($x := x(0)$);
 Weakly fair revision function R ;
 List of satisfied customers, initially empty: $L := \emptyset$;

repeat

Pick next player $k := R(t)$; $t := t + 1$;
if $x_k \notin \mathbf{br}_k(x)$ **then**
 Update strategy for player k to any $x_k \in \mathbf{br}_k(x)$;
 $L := \emptyset$;
 $L := L \cup \{k\}$;

until $size(L) = N$;

It is well known (see [11]) that for any potential game \mathfrak{G} , Algorithm 1 converges in finite time, almost surely, to a Nash Equilibrium of \mathfrak{G} .

3 Worst Case Complexity

In this section, we analyze the time complexity of BRA. More precisely, we consider three measures (related to each other). The first one is T_{BRA} , the number of iterations (or the number of times that the function \mathbf{br} is called) before BRA reaches a Nash equilibrium. A related measure is the total number of comparisons used by BRA (denoted C_{BRA}). One should expect that $C_{BRA} \approx (A - 1)T_{BRA}$. Finally, another interesting quantity is the number of different states visited by BRA (denoted M_{BRA}). This is called the number of *moves* done by BRA before convergence to a Nash equilibrium (NE). Of course, $M_{BRA} \leq T_{BRA}$.

These quantities depend on the game over which BRA is run, on the initial state x^0 and on the infinite sequence of revision players R . It should also be clear that they are functions of the game only through the potential Φ , so we denote by $T_{BRA}(\Phi, x^0, R)$ the number of steps before convergence of Algorithm BRA

for a game with potential Φ , starting in state x^0 , under the condition that the sequence of players is given by R .

In the worst case, for some weakly fair revision functions R , $T_{BRA}(\Phi, x^0, R)$ can be unbounded because the revision sequence induced by R can be arbitrarily bad: one player might appear too few times to guarantee convergence in any bounded time. When R is the round-robin function, the time for convergence is finite but can still be very large, as shown in the following theorem.

Theorem 1. *In the worst case, $T_{BRA}(\Phi, x^0, \text{round-robin}) = NA^{N-1}$.*

It is well known that the worst case complexity of BRA is exponential in the number of players (see for example [8]). The version of this result given here (for round robin revision and generalized potential games) is only given for the record (the proof is given in [17]).

4 Average Complexity of BRA

4.1 Randomization

In the following we will randomize over the potential games over which *BRA* is used. Since the behavior of BRA only depends on the potential function, we randomize directly over the potential Φ .

We consider a randomization over all games, uniformly over all possible orders for the potentials. On one hand, this is the classical average complexity approach when no additional information is known about the games (the same approach is used in [15] for 2 player games for example). This yields IID potential for all profiles, as explained below. On the other hand, some may argue that uniformly random games are not generic in some sense and a good performance of BRA on average does not necessarily translate in good performances for “real word” games. In any case, this is a first step that must be taken in absence of additional information about specific games that one may want to study.

There are several equivalent ways to do this randomization. The first one is based on the fact that the complexity of the algorithm does not depend on the actual values of the potential of the states but only on the comparisons between them. When two potentials are equal, a strict order between them is chosen uniformly. Therefore, the natural randomization is to consider the linear extensions (total orders) of all possible partial orders over the set \mathcal{A} and pick one uniformly. The number of total orders on \mathcal{A} is the number of permutations on \mathcal{A} , namely $(A^N)!$.

The second (equivalent) randomization is the following: The potentials of all states x are chosen independent, identically distributed according to an arbitrary distribution F admitting a density w.r.t. the Lebesgue measure.

Both randomizations are equivalent. Indeed, take any k states x_1, \dots, x_k in \mathcal{A} . In both cases, $\mathbb{P}(\Phi(x_1) > \Phi(x_2) > \dots > \Phi(x_k)) = 1/k!$. Now, since F is increasing, F^{-1} is well-defined and we get $\mathbb{P}(\Phi(x) > \Phi(x')) = \mathbb{P}(F^{-1}(\Phi(x)) > F^{-1}(\Phi(x')))$. Note that $F^{-1}(\Phi(x))$ is uniformly distributed on $[0, 1]$. Therefore, with no loss of generality, one can assume that the potential of all the states are i.i.d., uniformly distributed on $[0, 1]$. This randomization is used in the following.

4.2 Intersection-Free Approximation

The direct analysis of the behavior of BRA over a random potential is difficult because, over time, more and more states have been visited by the algorithm. Thus, its behavior is non-homogeneous in time. To avoid this difficulty, we consider a new model, called the *Intersection-Free Approximation* (IFA) in the following. Under the Intersection-Free Approximation, every time a new player (say k) has to compute its best response in a state (say x), it compares $\Phi(x)$ with the potential of its $A - 1$ other possible strategies, as for the real BRA. Here however, we assume that those $A - 1$ states have not yet been visited during the previous steps of the algorithm. Note that under the real behavior of BRA, it could happen that some of these states have already been compared at a previous step of the algorithm, by another player (this will be called an *intersection* in the following). Under the Intersection-Free Approximation, the states visited by the algorithm are always “new” states, never compared before with any other states.

More formally, the algorithm BRA under IFA can be written as follows.

Algorithm 2. BRA algorithm under IFA

Input: Initial state ($x(0)$); Revision function R ;
Set of satisfied players, initially empty $L := \emptyset$.

repeat

 Pick next player $k := R(t)$; $t := t + 1$;

if $k \notin L$ **then**

Generate IID potentials $\Phi(\alpha, x_{-k}), \alpha \in \mathcal{A}_k \setminus \{x_k\}$ unif. on $[0, 1]$;

 Compute best response: $\alpha_k := \operatorname{argmax}_{\beta \in \mathcal{A}_k} \Phi(\beta; x_{-k})$;

if $\alpha_k = x_k$ **then**

$L := L \cup \{k\}$

else

$L := \{k\}$; $x_k := \alpha_k$;

until $\text{size}(L) = N$;

Let us recall that C_{BRA} (resp. T_{BRA}, M_{BRA}) is the number of comparisons (resp. number of steps, number of moves) taken by BRA before convergence and let us define C_{IFA} (resp. T_{IFA}, M_{IFA}) to be the number of comparisons (steps, moves) of BRA under the intersection-free approximation. By definition, the worst case complexity of IFA under a round-robin revision sequence is infinite. However, its average complexity is the same as for BRA, as shown by the following lemma.

Lemma 1 (BRA and IFA are asymptotically equivalent). *Under the foregoing notations and using a round-robin revision function, the following comparisons hold, where \leq_{st} is the strong stochastic order:*

1. $C_{BRA} \leq_{st} C_{IFA}$ (equivalently, [19] $\forall t \in \mathbb{R}, \mathbb{P}(C_{BRA} > t) \leq \mathbb{P}(C_{IFA} > t)$).
2. If I is the total number of intersections in BRA, then $T_{BRA} \leq_{st} T_{IFA} + \frac{I}{A-1}$.
3. $\mathbb{E}[T_{BRA}] = \mathbb{E}[T_{IFA}] + o(1)$ and $\mathbb{E}[C_{BRA}] = \mathbb{E}[C_{IFA}] + o(1)$,
4. $\mathbb{E}[M_{BRA}] = \mathbb{E}[M_{IFA}] + o(1)$.

The proof of the lemma is available in [17]. It is based on the construction of a *coupling* between the executions of BRA with and without IFA. The assumption that the revision function is round-robin for BRA and for IFA does not play a big role in the proof, and it could be removed. However, the following section, asserting the optimality of round-robin implies that extending the proof to more general revision functions has a limited interest.

4.3 Round-Robin and Other Revision Sequences

As for the worst case analysis, the revision sequence influences the average time complexity of the algorithm. We show that on average round-robin is asymptotically the best one.

Lemma 2 (Asymptotic optimality of round-robin). *For any revision function R , $\mathbb{E}[T_{BRA}(\Phi, x^0, \text{round-robin})] \leq \mathbb{E}[T_{BRA}(\Phi, x^0, R)] + \epsilon(N)$, where the expectation is taken over all potentials Φ and all initial states x^0 and $\epsilon(N)$ goes to zero when N goes to infinity.*

The proof is again available in [17]. It uses the comparison with IFA. In the rest, we focus on round-robin revision functions and omit it in the notations, unless specified otherwise.

4.4 Complexity Analysis

We will be analyzing the intersection-free approximation of the behavior of BRA, under a round-robin revision sequence, with no further reference to this.

Let us consider the intersection-free approximation and let y be the potential of the current state x : ($y \stackrel{\text{def}}{=} \Phi(x)$). Let k be the number of players that have already played best response without changing the profile. This number of “satisfied” players can replace the explicit set L used in Algorithm 2 when the revision sequence is round-robin. The evolution at the next step of BRA under IFA is as follows. The k th player computes its best response. The player has $A - 1$ new strategies whose potential must be compared with the current potential (y). As mentioned before, we can assume that the potentials of those $a \stackrel{\text{def}}{=} A - 1$ strategies are IID, uniformly distributed in $[0, 1]$.

With probability y^a none of the new strategies beat the current choice. The state remains at y , one more player is satisfied and it is the turn of the $k + 1$ -st player to try its best response.

With probability $1 - y^a$, one of the new strategies is the best response. The current state moves to a new state where the number of satisfied players is set

back to 1 and the potential increases to a value larger than $u > y$ with probability $1 - u^a$.

Let Y_t be the potential at step t ($Y_t \in [0, 1]$) and K_t be the current number of consecutive players whose best response did not change the current potential ($K_t \in \{1, 2, \dots, N\}$) (number of satisfied players). The previous discussion says that the couple (Y_t, K_t) is a discrete-time, continuous-space Markov chain whose kernel is:

$$\mathbb{P}\left((Y_{t+1}, K_{t+1}) = (y, k + 1) \mid (Y_t, K_t) = (y, k)\right) = y^a,$$

and, for any $u > y$,

$$\mathbb{P}\left((Y_{t+1}, K_{t+1}) \in ([u, 1], 1) \mid (Y_t, K_t) = (y, k)\right) = 1 - u^a.$$

All the other transitions have a null probability.

Let $m(y, k)$ be the number of moves of IFA before convergence when the current state of the Markov chain is equal to (y, k) .

With probability y^a , the next player does not change its choice so that $m(y, k) = m(y, k + 1)$.

With probability density au^{a-1} the next player finds a new best response with potential u so that one move is taken and $m(y, k) = 1 + m(u, 1)$.

Let $M(y, k) = \mathbb{E}[m(y, k)]$. The previous one step analysis of $m(y, k)$ makes $M(y, k)$ satisfy a forward Poisson equation:

$$M(y, k) = y^a M(y, k + 1) + \int_y^1 au^{a-1}(M(u, 1) + 1)du.$$

By definition, the boundary conditions are: $\forall y, M(y, N) = 0$ (the current state is NE when all players agree on this) and $\forall k, M(1, k) = 0$ (the potentials are all bounded by 1, so a state with potential 1 is guaranteed to be a NE).

By setting $B(y) \stackrel{\text{def}}{=} \int_y^1 au^{a-1}(M(u, 1) + 1)du$, we get the following system of integral equations

$$\begin{cases} M(y, 1) & = y^a M(y, 2) + B(y), \\ M(y, 2) & = y^a M(y, 3) + B(y), \\ \vdots & = \vdots \\ M(y, N - 2) & = y^a M(y, N - 1) + B(y), \\ M(y, N - 1) & = B(y). \end{cases} \quad (1)$$

Successive substitution of $M(y, 2), \dots, M(y, N - 1)$ in the first equality yields $M(y, 1) = B(y)H(y)$ where $H(y) \stackrel{\text{def}}{=} 1 + y^a + \dots + y^{a(N-2)}$. Differentiating w.r.t. y , one gets an ordinary differential equation in $M(y, 1)$:

$$\frac{dM(y, 1)}{dy} + (ay^{a-1}H - \frac{1}{H} \frac{dH}{dy})M(y, 1) = -ay^{a-1}H.$$

The equation is of the form $\dot{f} + gf = h$. Using the boundary condition $M(1, 1) = 0$, its generic solution is

$$M(y, 1) = e^{-Q(y)} \int_y^1 au^{a-1} H(u) e^{Q(u)} du. \tag{2}$$

where

$$Q(y) \stackrel{\text{def}}{=} \int_0^y \left(au^{a-1} H(u) - \frac{1}{H(u)} \frac{dH(u)}{du} \right) du = -\log(H(y)) + \int_0^y au^{a-1} H(u) du. \tag{3}$$

The average number of profile changes in the execution of the algorithm starting from an arbitrary profile is $\mathbb{E}[M_{IFA}] = \int_0^1 M(y, 1) dy$. Since $M(y, 1)$ is decreasing in y , $\mathbb{E}[M_{IFA}]$ is upper-bounded by $M(0, 1)$. Using $Q(0) = 0$, $H(0) = 1$ and replacing Q and H by their values,

$$\begin{aligned} M(0, 1) &= \int_0^1 \exp\left(\sum_{i=0}^{N-2} \frac{u^{a(i+1)}}{i+1}\right) au^{a-1} du = \int_0^1 \exp\left(\sum_{i=0}^{N-2} \frac{v^{i+1}}{i+1}\right) dv \quad (\text{with } v = u^a) \\ &= \int_0^{1-\frac{1}{N}} \exp\left(\sum_{i=1}^{N-1} \frac{v^i}{i}\right) dv + \int_{1-\frac{1}{N}}^1 \exp\left(\sum_{i=1}^{N-1} \frac{v^i}{i}\right) dv \\ &\leq \int_0^{1-\frac{1}{N}} \exp\left(\sum_{i=1}^{\infty} \frac{v^i}{i}\right) dv + \frac{1}{N} \exp\left(\sum_{i=1}^{N-1} \frac{1}{i}\right) \end{aligned} \tag{4}$$

$$= \int_0^{1-\frac{1}{N}} \frac{dv}{1-v} + e^\gamma + O(1/N) = \log(N) + e^\gamma + O(1/N). \tag{5}$$

Furthermore, this bound is tight, up to an additive constant (see [17]).

Let us now consider the average number of comparisons made by BRA under the intersection-free assumption. Let $C(y, k)$ be the average number of comparisons starting in a state with potential y and k players have played without changing their strategy. The Poisson equation for $C(y, k)$ is :

$$C(y, k) = y^a(C(y, k + 1) + a) + \int_y^1 au^{a-1}(C(u, 1) + a)du,$$

with the boundary conditions $C(1, 1) = a(N - 1)$ and $C(y, N) = 0$.

The solution of this differential system can be obtained in closed form, using a similar approach as for $M(y, 1)$.

$$C(y, 1) = a \left(\sum_{i=0}^{N-2} y^{ai} \right) \exp\left(-\sum_{i=1}^{N-1} \frac{y^{ai} - 1}{i}\right).$$

The average number of comparisons is $\mathbb{E}[C_{IFA}] = \int_0^1 C(y, 1) dy$.

For all $y < 1$,

$$C(y, 1) = a \left(\sum_{i=0}^{\infty} y^{ai} \right) \exp \left(\sum_{i=1}^{N-1} 1/i \right) \exp \left(- \sum_{i=1}^{\infty} y^{ia}/i \right) + o(aN) \quad (6)$$

$$= a \frac{1}{1 - y^a} (N - 1) e^{\gamma} (1 - y^a) + o(aN) + O(1) \quad (7)$$

$$= a(N - 1) e^{\gamma} + o(aN), \quad (8)$$

where γ is the Euler constant ($\gamma \approx 0.5772\dots$). Therefore, the same equality holds for the integral, equal to $\mathbb{E}[C_{IFA}]$.

The results of this section, together with Lemma 1, lead to the following theorem, the main result of the section.

Theorem 2 (Average complexity of BRA). *Under the round-robin revision sequence, the average complexity of BRA over a potential game satisfies:*

- (i) *Average number of moves: $\mathbb{E}[M_{BRA}] = \log(N) + c + O(1/N)$, where $c \leq e^{\gamma}$*
- (ii) *Average number of comparisons: $\mathbb{E}[C_{BRA}] = e^{\gamma} AN + o(AN)$.*
- (iii) *Average number of steps: $\mathbb{E}[T_{BRA}] = e^{\gamma} N + o(N)$.*

The average complexity $\mathbb{E}[T_{BRA}]$ can be split into two parts: The number of plays before reaching a NE and the number of plays needed to check if a state is indeed a NE. This last part takes exactly $N - 1$ steps in Algorithm 1: The players have to play one by one to fill up set L . This means that a NE equilibrium is reached on average as soon as $e^{\gamma} - 1 \approx 78\%$ of the players have played once. The second moments of the number of steps and the number of moves under IFA can be computed similarly (see [17]). In both cases, the standard deviations are of the same order as the means.

5 Optimality of BRA

In this section, we prove that BRA finds a Nash equilibrium faster than any local search algorithm (defined in Sect. 5), in the strong stochastic order sense.

By definition a *Local Search Algorithm* can only access the payoff matrix, one player at a time. This access is often called a *query* in the literature. Once the payoff of a strategy profile has been obtained, it is stored in memory and can re-used later by the algorithm without an additional query.

In addition to queries, a local search algorithm can use any arithmetic operation, draw random variables and choose a strategy for all players.

Any local search algorithm can be written in the following form, based on the history of the execution, \mathcal{H}_t , that corresponds to the amount of information gathered by the algorithm up to step t .

Algorithm 3. A general local search algorithm

Initial storage reduced to the initial profile: $\mathcal{H}_0 := \{(x(0))\}$.
repeat
 Select next player: $k := R(\mathcal{H}_t)$;
 Query payoff vector of k under current state: $u_k(\cdot, x_{-k}(t))$;
 Store the new visited states and their payoffs in memory:
 $\mathcal{H}_{t+1} := \mathcal{H}_t \cup \{((\alpha, x_{-k}(t)), u_k(\alpha, x_{-k}(t)))_{\alpha \in \mathcal{A}_k}\}$;
 jump to next state $x(t+1) := J(\mathcal{H}_{t+1})$;
 Set $stop := 1$ if the current state is a NE;
 $t := t + 1$;
until $stop$;

The functions J and R used in the inner loop are arbitrary functions that choose the next state as well as the next player to play, according to the whole history of the process. These functions can be deterministic or random. Testing if $x(t+1)$ is NE is not detailed. Notice, however, that it can only be done when all the payoff vectors for all the players in state $x(t+1)$ have been stored in memory.

The complexity of a local search algorithm A is defined as the total number of its payoff vector queries (denoted T_A).

Theorem 3 (Optimality of BRA). *Let A be any local search algorithm that computes a Nash Equilibria in potential games. Under the foregoing randomization, and choosing the starting point x^0 uniformly among all states, $\forall t \geq 0$, $\mathbb{P}(T_{BRA} \geq t | R = R_A) \leq \mathbb{P}(T_A \geq t)$, where R_A is the revision sequence constructed in A .*

The proof is reported in [17]. Combining this theorem with Lemma 2 establishes the optimality of BRA with round-robin.

6 Conclusion and Perspectives

The best response algorithm is one on the most basic object in game theory. In this paper, we prove it has a linear complexity on average over uniformly randomized potential games. Furthermore, BRA is optimal in the class of local search algorithms when one has no information about the structure of the potential game.

Does all this make BRA the perfect algorithm to compute NE in general? We believe that the answer is no because BRA suffers from several drawbacks. First, it does not tolerate simultaneous plays. Second, it requires to know the entire payoff vector of a player before choosing its strategy. Other drawbacks include

high sensitivity on the order of play and on noisy perturbations on the payoffs. Designing algorithms that do not suffer from these drawbacks is the object of our future investigations.

Acknowledgement. This work was partially supported by LabEx Persyval-Lab.

References

1. Daskalakis, P.G.C., Papadimitriou, C.: The complexity of computing a nash equilibrium. *SIAM J. Comput.* **39**(3), 195–259 (2009)
2. Rosenthal, R.W.: A class of games possessing pure-strategy nash equilibria. *Int. J. Game Theory* **2**(1), 65–67 (1973)
3. Beckman, M., McGuire, C.B., Winsten, C.B.: *Studies in the Economics of Transportation*. Yale University Press, New Haven (1956)
4. Orda, A., Rom, R., Shimkin, N.: Competitive routing in multi user communication networks. *IEEE/ACM Trans. Netw.* **1**(5), 510–521 (1993)
5. Gallager, R.G.: A minimum delay routing algorithm using distributed computation. *IEEE Trans. Commun.* **25**(1), 73–85 (1977)
6. Wardrop, J.: Some theoretical aspects of road traffic research, part ii. In: *Proceedings of the Institute of Civil Engineers*, vol. 1, pp. 325–378 (1954)
7. Roughgarden, T.: *Selfish Routing and the Price of Anarchy*. MIT Press, Cambridge (2005)
8. Fabrikant, A., Papadimitriou, C., Talwar, K.: The complexity of pure nash equilibria. In: *Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing*, ser. (STOC 2004), pp. 604–612. ACM (2004)
9. Ackermann, H., Röglin, H., Vöcking, B.: On the impact of combinatorial structure on congestion games. In: *FOCS 2006*, pp. 613–622 (2006)
10. Nash, J.: Equilibrium points in n-person game. *Proc. Nat. Acad. Sci.* **38**, 48–49 (1950)
11. Monderer, D., Shapley, L.: Potential games. *Games Econ. Behav.* **14**(1), 124–143 (1996). Elsevier
12. Babichenko, Y., Tamuz, O.: Graphical potential games. arXiv.org, Tech. Rep.: 1405.1481v2
13. Even-Dar, E., Kesselman, A., Mansour, Y.: Convergence time to nash equilibria. In: Baeten, J.C.M., Lenstra, J.K., Parrow, J., Woeginger, G.J. (eds.) *ICALP 2003*. LNCS, vol. 2719. Springer, Heidelberg (2003)
14. Feldman, M., Tamir, T.: Convergence of best-response dynamics in games with conflicting congestion effects. In: Goldberg, P.W. (ed.) *WINE 2012*. LNCS, vol. 7695, pp. 496–503. Springer, Heidelberg (2012)
15. Imre, B., Vempala, S., Vetta, A.: Nash equilibria in random games. *Random Structures Algorithms* **31**(4), 391–405 (2007)
16. Rinott, Y., Scarsini, M.: On the number of pure strategy nash equilibria in random games. *Games Econ. Behav.* **33**(2), 274–293 (2000)
17. Durand, S., Gaujal, B.: Complexity, optimality of the best response algorithm in random potential games. Inria, Research Report RR-8925 (2016). <https://hal.inria.fr/hal-01330805>
18. Voorneveld, M.: Best-response potential games. *Econ. Lett.* **66**(3), 289–295 (2000)
19. Müller, A., Stoyan, D.: *Comparison Methods for Stochastic Models and Risks*. Series in Probability and Statistics. Wiley, Chichester (2002)