Average complexity of the Best Response Algorithm in Potential Games

Stéphane Durand, Bruno Gaujal

Average complexity of the Best Response Algorithm in Potential Games [*]

Stéphane Durand[1,2,3]   Bruno Gaujal[3,1]

[1] Univ. Grenoble Alpes, LIG, F-38000 Grenoble, France
[2] CNRS, LIG and Gipsa-Lab F-38000 Grenoble, France
[3] Inria

stephane.durand@inria.fr, bruno.gaujal@inria.fr

Mots-clés : *Potential Games, Best Response, Average Complexity*

1   Introduction

Potential games have been introduced in [6] and have proven very useful, especially in the context of routing games, first mentioned in [1] and exhaustively studied ever since, in the transportation as well as computer science litterature, see for example [3,5,8] and for distributed optimization (see for example [7]).

The Best Response Algorithm (BRA) is probably the most popular algorithm that converges to a pure Nash equilibrium (NE) in potential games [4]. However, its efficiency (speed of convergence) has attracted surprisingly little attention.

In this paper, we analyze the performance of BRA used over a potential game with N players, each with A possible actions, and we provide answers to the following questions.

1. How many steps does the Best Response Algorithm take before it stops in a Nash Equilibrium ?

2. A related question is how many comparisons are done during the execution of BRA ?

In a nutshell, the answers to these questions are respectively,

1. less than $\log(N) + e^\gamma$ on average (γ is the Euler constant) ;

2. $e^\gamma A N$ on average.

These numbers say that BRA is a very efficient distributed algorithm to compute NE. Our analysis is based on two ingredients, one is the construction of an approximation of the behavior of BRA, where each state is examined at most once and the second is the use of a continuous-state space-discrete-time Markov chain to analyze the average complexity.

2   Best Response Algorithm and Potential games

We consider a game with a finite number N of players and a finite strategy space for each player, each of size A, and the corresponding utility functions.

$\mathfrak{G} \stackrel{\text{def}}{=} \mathfrak{G}(\mathcal{N}, \mathcal{A}, u)$ will be a tuple consisting of
– a finite set of *players* $\mathcal{N} = \{1, \dots, N\}$;
– a finite set $\mathcal{A}_k$ of *actions* (or *pure strategies*) for each player $k \in \mathcal{N}$; The set of *actions profiles* or *states* of the game is $\mathcal{A} \stackrel{\text{def}}{=} \prod_k \mathcal{A}_k$;
– the players' *payoff functions* $u_k : \mathcal{A} \to \mathbb{R}$.

We define the classical *best response correspondence* $\mathbf{br}_k(x)$ as the set of all actions that maximizes the payoff for player k under profile x :

$$\mathbf{br}_k(x) \stackrel{\text{def}}{=} \left\{ \underset{\alpha \in \mathcal{A}_k}{\text{argmax}}\, u_k(\alpha; x_{-k}) \right\}. \tag{1}$$

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

Iteratively playing a best response may not converge in general. We consider here the specific class of potential games for which convergence is ensured.

Définition 1 (Potential games) *A game is a* potential game *[4] if it admits a function (called the potential)* $\Phi : \mathcal{A} \to \mathbb{R}$ *such that for any player k and any* unilateral *deviation of k from action profile x to x', $u_k(x) - u_k(x') = \Phi(x) - \Phi(x')$.*

We consider a general version of *Best Response Algorithm* (BRA) where the next player is selected according to a *revision sequence* $(R_t)_{t\in\mathbb{N}}$. We assume that this sequence of players is *weakly fair* : each players appears infinitely often in the sequence. The revision sequence can be deterministic (for example, Round-Robbin) or random (for example, the next player is chosen according to a probability distribution ρ (the *revision law*) : $\forall k \in \mathcal{N}, \mathbb{P}(\text{selected player} = k) = \rho_k$). In that case, we assume that the probability of choosing any player k is strictly positive ($\forall k \in \mathcal{N}, \rho_k > 0$). This insures that the revision sequence is weakly fair almost surely.

Algorithm 1: Best Response Algorithm (BRA)

Input :
Game utilities $(u_i(\cdot))$,
Initial state $(x(0))$,
Infinite revision sequence $R = (R_t)_{t\in\mathbb{N}}$ of players.

foreach *player $k \in K$* **do**
 $\quad$ $stop_k := false$
repeat
 $\quad$ Pick next player $k := R_{t+1}$
 $\quad$ Select new action $\alpha_k := \mathbf{br}_k(x(t))$
 $\quad$ $stop_k := \mathbf{1}_{\{\alpha_k = x_k(t)\}}$;
 $\quad$ $x_k(t+1) := \alpha_k$;
until $stop_1 \wedge stop_2 \wedge \cdots \wedge stop_N$;

A famous result first proved in [4] states that for any potential game $\mathfrak{G}$, Algorithm 1 converges in finite time (a.s.) to a Nash Equilibrium of $\mathfrak{G}$.

3 Complexity

In this section, we analyse the time complexity of BRA. More precisely, we consider three measures (related to each other). The most important quantity is T_{BRA}, the number of iterations (or the number of times that the function **br** was called) before BRA reaches a Nash equilibrium. A related measure is the total number of comparisons used (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. This is the same (up to one) as the number of moves (denoted M_{BRA}) taken by BRA before convergence to a NE. Of course, $M_{BRA} \leqslant T_{BRA}$. The proofs of the propositions 1 and 2 are not provided due to lack of space. They are available in a research report [2].

In the worst case, for random weakly fair revision sequences R, T_{BRA} is unbounded because the revision sequence R can be arbitrarily bad : one player might appear too few times to guarantee convergence in any bounded time. Indeed, the weak fairness assumption is too weak

to prevent unbounded gaps between the turns of one player). Even when this is not the case, for example when R is a round-robin sequence, the time for convergence can still be very large, as shown in the following proposition.

Proposition 1 *In the worst case, under round robin revisions, $T_{BRA} = NA^{N-1}$.*

The worst case complexity does not say much about the typical behavior of BRA. In the following we will analyze its average complexity when the revision sequence is round-robin. Other revision sequences behave similarly (up to a multiplicative factor).

3.1 Randomization

In the following we will randomize over the potential game over which BRA is used. Since the behavior of BRA only depends on the potential function, we randomize directly over the potential Φ. The natural randomization is to consider all possible total orderings of the set $\{\Phi(x), x \in \mathcal{A}\}$ (there are $(A^N)!$ of them) and pick one uniformly. This is equivalent to pick iid potentials in all states, uniformly distributed in $[0, 1]$.

3.2 Markovian Analysis

We will be analysing the intersection-free approximation of the behavior of BRA (where no state is visited twice) whose behavior is asymptotically the same as BRA.

Let y be the potential of the current state x : ($y \overset{\text{def}}{=} \Phi(x)$). If $k - 1$ players have already played best response without changing the profile, then the evolution at the next step of BRA is as follows. The k-th player computes its best response. This player has $a \overset{\text{def}}{=} A - 1$ new actions whose potential must be compared with the current potential (y). With probability y^a none of the new actions beat the current choice. The state remains at y 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 actions is the best response. The current state moves to a new state with a larger potential and the number of players for which the new state is a best response is set back to 1. More precisely, with probability $1 - z^a$ the potential moves from y to a value larger than z.

This says that the couple $(Y(t), K(t))$ is a Markov chain, where $Y(t)$ is the potential at step t, in $[0, 1]$ and $K(t)$ is the current number of players whose best response did not change the current state (in $\{1, 2, \dots, N\}$). Its transitions are :

$$\mathbb{P}\Big((Y(t+1), K(t+1)) = (y, k+1) \Big| (Y(t), K(t)) = (y, k) \Big) = y^a,$$

and, if $z > y$,

$$\mathbb{P}\Big((Y(t+1), K(t+1)) \in ([z, 1], 1) \Big| (Y(t), K(t)) = (y, k) \Big) = 1 - z^a.$$

Now, let $M(y, k)$ be the average number of moves of BRA before convergence when the current potential equals y and k players have played without changing their action, at potential y.

The quantity $M(y, k)$ satisfies the forward heat equation of the Markov chain :

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

Let us also 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 action. The forward 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 $C(y, N) = 0$.

Solving these equations leads to the following proposition.

Proposition 2 *The average number of moves in BRA before convergence can be bounded :* $\mathbb{E}M_{BRA} \leqslant \log(N) + e^{\gamma} + O(1/N)$, *and this bound is tight, up to an additive constant.*

The average number of comparisons made during an execution of BRA is asymptotycally $\mathbb{E}C_{BRA} = e^{\gamma}(A-1)(N-1) + o(A)$.

The average number of steps is $\mathbb{E}T_{BRA} = e^{\gamma}(N-1) + o(1)$.

3.3 Numerical experiments

The following figures correspond to simulations on games, all with $A = 30$. The number of players N ranges from 0 to 250. These games have potentials chosen uniformly in $[0, 1]$. For each value of N, algorithm BRA is run 5000 times. the errorbars correspond to confidence interval at 95%.
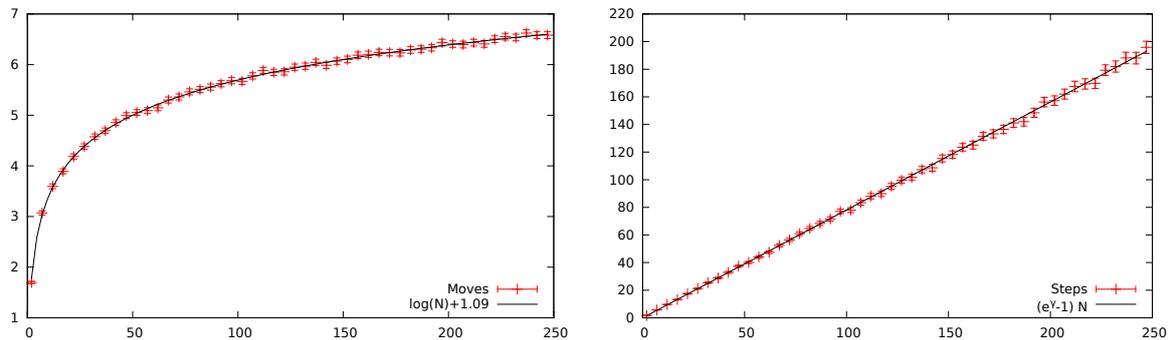
(a) Average number of moves.
The best fit function is $\log(N) + 1.09$

(b) Average number of steps before reaching the NE.
The predicted value is $(e^{\gamma} - 1)N$.

FIG. 1 – Simulation results for Round-Robin BRA

As one can see in Figure 1a the average number of moves for BRA with Round-Robin approaches $\log N + 1.09$, while Theorem 2 predicts that it should be of the form $\log N + Cte$, with $Cte \leqslant e^{\gamma} \approx 1.7$. As predicted by Theorem 2, Figure 1b shows that the number of comparisons is $e^{\gamma}(A-1)N$, with a very tight confidence interval. This provides numerical evidence that the variance is small (and the distribution is not spread).

Références

[1] M. Beckman, C. B. McGuire, and C. B.Winsten. *Studies in the Economics of Transportation.* Yale University Press, 1956.

[2] Stéphane Durand and Bruno Gaujal. Efficiency of the best response algorithm in potential games. Technical report, Inria, 2015.

[3] R. G. Gallager. A minimum delay routing algorithm using distributed computation. *IEEE Transactions on Communications*, 25(1) :73–85, 1977.

[4] Dov Monderer and Lloyd Shapley. Potential games. *Games and economic behavior, Elsevier*, 14(1) :124–143, 1996.

[5] A. Orda, R. Rom, and N. Shimkin. Competitive routing in multeuser communication networks. *IEEE/ACM Trans. on Networking*, 1(5) :510–521, 1993.

[6] Robert W. Rosenthal. A class of games possessing pure-strategy nash equilibria. *Int. J. of Game Theory, Springer*, 2(1) :65–67, 1973.

[7] T. Roughgarden. *Selfish Routing and the Price of Anarchy.* MIT Press, 2005.

[8] J.G. Wardrop. Some theoretical aspects of road traffic research. part ii. *Proc. of the Institute of Civil Engineers*, 1 :325–378, 1954.