

Efficiency of Best Response Dynamics with High Playing Rates in Potential Games

Stéphane Durand, Federica Garin, Bruno Gaujal

► **To cite this version:**

Stéphane Durand, Federica Garin, Bruno Gaujal. Efficiency of Best Response Dynamics with High Playing Rates in Potential Games. IFIP WG 7.3 Performance 2018 - 36th International Symposium on Computer Performance, Modeling, Measurements and Evaluation, Dec 2018, Toulouse, France. pp.1-2. hal-01940150

HAL Id: hal-01940150

<https://hal.inria.fr/hal-01940150>

Submitted on 30 Nov 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Efficiency of Best Response Dynamics with High Playing Rates in Potential Games*

Stéphane Durand
Univ. Grenoble Alpes, Inria,
CNRS, Grenoble INP
LIG, GIPSA-Lab
Grenoble, France
stephane.durand@inria.fr

Federica Garin
Univ. Grenoble Alpes, Inria,
CNRS, Grenoble INP
GIPSA-Lab
Grenoble, France
federica.garin@inria.fr

Bruno Gaujal
Univ. Grenoble Alpes, Inria,
CNRS, Grenoble INP
LIG
Grenoble, France
bruno.gaujal@inria.fr

ABSTRACT

In this paper we design and analyze distributed best response dynamics to compute Nash equilibria in potential games. This algorithm uses local Poisson clocks for each player and does not rely on the usual but unrealistic assumption that players take no time to compute their best response. If this time (denoted δ) is taken into account, distributed best response dynamics (BRD) may suffer from *overlaps*: one player starts to play while another player has not changed its strategy yet. An overlap may lead to a decrease of the potential but we can show that they do not jeopardize eventual convergence to a Nash equilibrium. Our main result is to use a Markovian approach to show that the average execution time of the algorithm $\mathbb{E}[T_{BRD}]$ can be bounded: $2\delta n \log n / \log \log n + O(n) \leq \mathbb{E}[T_{BRD}] \leq 4e^\gamma \delta n \log n / \log \log n + O(n)$, where γ is the Euler constant, n is the number of players and δ is the time taken by one player to compute its best response. These bounds are obtained by using an asymptotically optimal playing rate $\hat{\lambda}$. Our analytic bound shows that $2\delta\hat{\lambda} = \log \log n - \log C$, where C is a constant. This induces a large probability of overlap ($\hat{p} = 1 - C / \log^{1/2} n$). In practice, numerical simulations also show that using high playing rates is efficient, with an optimal probability of overlap $p_{opt} \approx 0.78$, for n up to 250. This shows that best response dynamics are unexpectedly efficient to compute Nash equilibria, even in a distributed setting.

Categories and Subject Descriptors

F.7 [Theory of Computation]: Theory and algorithms for application domains—*Algorithmic game theory and mechanism design*

Keywords

Potential Games, Best Response Dynamics, Distributed Algorithms, Markov Chains

1. POTENTIAL GAMES

A game $\mathfrak{G} \stackrel{\text{def}}{=} \mathfrak{G}(\mathcal{N}, \mathcal{A}, u)$ is a triplet consisting of: A finite set of *players* $\mathcal{N} = \{1, \dots, n\}$; A finite set \mathcal{A} of *actions* (or

*This work has been partially supported by the LabEx PERSYVAL-Lab

pure strategies) ; The set of *actions profiles* or *states* of the game is \mathcal{A}^n and the players' *payoff functions* $u_k : \mathcal{A}^n \rightarrow \mathbb{R}$.

The *best response correspondence* $\mathcal{BR}_k(x)$ is the set of actions maximizing the payoff for player k under state \mathbf{x} :

$$\mathcal{BR}_k(\mathbf{x}) \stackrel{\text{def}}{=} \left\{ \arg \max_{\alpha_k \in \mathcal{A}} u_k(\alpha_k; \mathbf{x}_{-k}) \right\}.$$

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

DEFINITION 1 (Potential games). *A game is a (best response) potential game [2] if there exists a function $\Phi : \mathcal{A}^n \rightarrow \mathbb{R}$ such that for any player k and action profile \mathbf{x}*

$$\mathcal{BR}_k(\mathbf{x}) = \left\{ \arg \max_{\alpha_k \in \mathcal{A}} \Phi(\alpha_k, \mathbf{x}_{-k}) \right\}.$$

The main property of best response in potential games is that each time a player plays its best response, the potential increases. This implies that repeating this for all players will eventually lead to a configuration where the potential cannot be improved, hence a Nash Equilibrium.

2. EXECUTION TIME OF DISTRIBUTED BEST RESPONSE DYNAMICS

We consider an algorithmic version of Best Response called *Best Response Dynamics* (BRD) in the following.

A revision sequence is an infinite sequence of players chosen according to some rule. The round robin sequence, i.e., the cyclic sequence $1, \dots, n, 1, \dots, n, 1, \dots$ has been studied in [1]. Here, we consider a distributed version of the best response dynamics, given in Algorithm 1, where players act according to individual Poisson clocks with rate λ/n .

Let δ be the time taken by any player to compute its best response. If two players become active during a time interval smaller than δ , their plays *overlap*: one player starts to play before the previous one has yet updated its state. This implies that the best response of the second player may be corrupted. This may lead to a potential decrease. Convergence to a Nash equilibrium still holds with probability one: Using the Borel Cantelli Lemma, a sequence of plays of arbitrary length and with no overlap will eventually occur. Of course it may take a very long time before such an event occurs. We show next that this is not the case. On the contrary, the average execution time is surprisingly small: convergence occurs before all players have played 4 times.

Algorithm 1: Distributed BRD (for player k)

```
1 Input: Game utilities ( $u_k(\cdot)$ ); Initial state ( $\mathbf{x} := \mathbf{x}(0)$ );  
2 Initialize  $t := 0$ ; List of satisfied players  $L := \emptyset$ ;  
3 Local clock, ticking w.r.t. a Poisson process with rate  
    $\lambda/n$ ;  
4 while  $size(L) \neq n$  do  
5   On each tick of the local clock  
6     if  $k \notin L$  then  
7       if  $x_k = \mathcal{BR}_k(\mathbf{x})$  then  
8          $L := L \cup \{k\}$   
9       else  
10         $x_k := \mathcal{BR}_k(\mathbf{x})$ ;  $L := \{k\}$ 
```

Randomization. In the following we will randomize over the potential games over which Best Response Dynamics (BRD) is used. Since the behavior of BRD only depends on the potential function, we randomize directly over the potential Φ .

THEOREM 1. *Under rate λ , the overlap probability is $p = 1 - e^{-\delta\lambda}$ and the expected execution time satisfies*

$$\mathbb{E}[T_{BRD}] \leq \frac{e}{\lambda p} \left(n \log(n) + \frac{C}{1-p} \left(n + n^p + \frac{n^{2p-1}}{p(1-p)} \right) \right),$$

where C is a constant, and $\mathbb{E}[T_{BRD}] \geq n \log(n)/\lambda$. Choosing an adequate global playing rate $2\delta\lambda = \log(\log(n)) - \log(C)$, inducing an overlap probability $\hat{p} = 1 - \frac{C}{\log^{1/2}(n)}$, the average execution time of BRD satisfies

$$2 \frac{\delta n \log(n)}{\log(\log(n))} (1+o(1)) \leq \mathbb{E}[T_{BRD}] \leq 4e^\gamma \frac{\delta n \log(n)}{\log(\log(n))} (1+o(1))$$

where γ is the Euler constant.

3. MARKOVIAN APPROXIMATION

The direct analysis of the behavior of BRD is difficult because, over time, more and more states have been visited by the algorithm. Thus, its behavior is non-homogeneous in time. We introduce two successive simplifications whose behavior can be compared (using two different coupling schemes) with the best response dynamics over a potential game with n players.

Under the Intersection-Free approximation (IFA), every time a new player who is not satisfied (in other words, a player not in L) has to compute its best response in a state (say \mathbf{x}), it compares $\Phi(\mathbf{x})$ with the potential of its other strategies, as for the original BRD. In IFA, however, we will assume that none of these states have yet been visited during the previous steps of the algorithm. Under BRD, there is a chance that some of these states have already been visited. Under BRD, the potentials of all states is generated first, and the game is played given these values of the potentials. Under IFA, each time a state is (re)visited, its iid potential is (re)generated. Using an appropriate coupling between BRD and IFA, the expected convergence time of brd ($\mathbb{E}T_{brd}$) and of ifa ($\mathbb{E}T_{ifa}$) verify: $\mathbb{E}T_{brd} \leq c_1 \mathbb{E}T_{ifa}$, for some constant c_1 .

We next construct the restart approximation (RST) by considering that whenever an overlap occurs, the potential is reset to a uniform potential in $[0, 1]$, disregarding the previously obtained potentials. In other words, we restart from

step 0. Again a coupling technique can be used to show that the time to reach a Nash equilibrium under RST is stochastically larger than under IFA.

The construction of the restart approximation ensures that the future behavior of the system only depends on its current state, namely the current potential and the list of satisfied players. Therefore, the convergence time of (RST) can be evaluated by modeling its behavior by a Markov chain over a hybrid state space (with a continuous and a discrete component). The expected hitting time of a Nash equilibrium for this Markov chain satisfies a Poisson equation. This Poisson equation can be transformed into an ordinary differential equation with an implicit initial condition whose solution can be computed in integral form. The derivation of the asymptotic behavior of this solution in n (the number of players) leads to Theorem 1.

4. SENSITIVITY TO λ

We have simulated the behavior of BRD over a large number of potential games whose potentials are generated uniformly over all possible potentials. These simulations are carried to assess tightness of our bound and to study the sensitivity of BRD to the value of its unique parameter λ .




Figure 1: Mean execution time of distributed BRD with 100 players (with 95% confidence intervals), as a function of the expected time between two plays (inverse of the playing rate), in log-log scale. Comparison with the analytic asymptotic bound

Our experiments show that λ_{opt} does not depend on n : for any n , $1/\delta\lambda_{opt} \approx 0.63$ implying $p_{opt} \approx 0.77$. However it was very hard to get a good precision on λ_{opt} . Even if these experiments took several days of computations over a modern desktop computer, with 1000 runs for each couple (n, λ) , our accuracy on the optimal value of $\lambda_{opt}(n)$ is between 5% and 10%, depending on the value of n .

5. REFERENCES

- [1] Stéphane Durand and Bruno Gaujal. Complexity and optimality of the best response algorithm in random potential games. In *Symposium on Algorithmic Game Theory (AGT) 2016*, pages 40–51, Liverpool, United Kingdom, September 2016.
- [2] Mark Voorneveld. Best-response potential games. *Economics letters*, 66(3):289–295, 2000.