Convergence of Best-Response Dynamics in Games with Conflicting Congestion Effects

Michal Feldman[1] and Tami Tamir[2]

[1] Hebrew University of Jerusalem and Harvard University.
E-mail: mfeldman@huji.ac.il.

[2] School of Computer science, The Interdisciplinary Center, Herzliya, Israel.
E-mail: tami@idc.ac.il

Abstract. We study the model of resource allocation games with conflicting congestion effects introduced by Feldman and Tamir (2012). In this model, an agent's cost consists of its resource's load (which increases with congestion) and its share in the resource's activation cost (which decreases with congestion). The current work studies the convergence rate of best-response dynamics (BRD) in the case of homogeneous agents. Even within this simple setting, interesting phenomena arise. We show that, in contrast to standard congestion games with identical jobs and resources, the convergence rate of BRD under conflicting congestion effects might be super-linear in the number of jobs. Nevertheless, a specific form of BRD is proposed, which is guaranteed to converge in linear time.

1 Introduction

Resource allocation is considered to be a fundamental problem in algorithmic game theory, and has naturally been the subject of intensive research within this field. Most of the game-theoretic literature on resource allocation settings emphasizes either the negative or the positive congestion effects on the individual cost of an agent. The former approach assumes that the cost of a resource is some non-decreasing function of its load. This literature includes job scheduling and routing models [11, 20]. In these cases an individual user will attempt to avoid sharing its resource with others as much as possible. The second approach, in stark contrast, assumes that a resource's cost is a decreasing function of its load. This is the case, for example, in network design and cost sharing connection games, in which each resource has some *activation cost*, which should be covered by its users [2, 6]. In these cases, an individual user wishes to share its resource with as many other users as possible in attempt to decrease its share in the cost.

In reality, most applications have cost functions that exhibit both negative and positive congestion effects. Accordingly, more practical models that integrate the two congestion effects into a unified cost function have been considered [1, 9, 15]. The present paper studies the resource allocation setting that is introduced by Feldman and Tamir [9], in which the individual cost of an agent is composed of two components, one that exhibits positive externalities, and the other that exhibits negative externalities. More specifically, every resource has some activation cost, that is shared among all the agents using it. The individual cost of an agent is the sum of its chosen resource's load (reflecting the negative externalities) and its share in the resource's activation cost (reflecting the positive externalities). This model is applicable to a large set of applications, including job scheduling, network routing, and network design settings.

The induced game, unlike its two "parent games," is not a potential game[3] when played by heterogeneous agents. Indeed, it has been shown in [9] that best-response dynamics (BRD) do not necessarily converge in this setting. Yet, in the special case where agents are identical, the induced game is a potential game; consequently, any BRD is guaranteed to converge to a Nash equilibrium [9]. The rate of the convergence, however, has been overlooked thus far. It is argued that the convergence rate is crucial for the Nash equilibrium hypothesis to hold; that is, it is more plausible that a Nash equilibrium will be reached if natural dynamics lead to such an outcome within a small number of moves.

In this paper, we study the convergence rate of BRD in a job scheduling game with conflicting congestion effects and identical agents.

1.1 Our results

It is fairly easy to see that for unit-size jobs, convergence to a Nash equilibrium is linear in the number of jobs in both of the "parent" models; namely, if the the cost function equals the resource's load or if it equals the job's share in the resource's activation cost. We find that if the cost function takes both components into consideration, the convergence rate might be super-linear. We then introduce a specific form of BRD, referred as *max-cost*, where the job that incurs the highest cost is the one to perform its best move. The motivation behind this BRD is clear: the job that incurs the highest cost has the strongest incentive to improve its state. For *max-cost*-BRD, linear convergence rate is guaranteed. Due to space constraints, we defer some proofs to the full version [10].

1.2 Related work

A lot of research has been conducted in the analysis of job-scheduling applications using a game-theoretic approach, where the jobs are owned by players who choose the machine to run on. The questions that are commonly analyzed under this approach are Nash equilibrium existence, the convergence of best-response dynamics to a Nash equilibrium, and the inefficiency of Nash equilibria (quantified mainly by the price of anarchy [16, 18] and price of stability [2] measures).

It is well known that every congestion game is a potential game [19, 17], and therefore admits a pure Nash equilibrium, and every best-response dynamics converges to a pure Nash equilibrium. However, the convergence time may, in general, be exponentially long [1, 8, 21]. This observation has led to a large amount of work that identified special classes of congestion games, where best-response dynamics converge to a Nash equilibrium in polynomial time or even linear time. This agenda has been the focus of [7, 12] in a setting with negative congestion effects, and was also studied in a setting of positive congestion effects [2]. In particular, it has been shown that it takes at most n steps (where n is the

[3] Potential games have been introduced by [17].

number of users) to converge to a Nash equilibrium if the network is composed of parallel links [7], and this result has been later extended to extension-parallel networks [12]. For resource selection games (i.e., where feasible strategies are composed of singletons), it has been shown in [14] that better-response dynamics converge within at most mn^2 steps for general cost functions (where m and n are the number of resources and users, respectively). In addition to standard better- and best-response dynamics, a few variants have been explored. One example is the study of the convergence rate of α-Nash dynamics to an approximate Nash equilibrium [5] and to an approximate optimal solutions [3]. Also, the robustness of best-response convergence to altruistic agents has been studied in [13], where it has been shown that BRD may cycle as a result of altruism.

In this paper we study the congestion models with conflicting congestion effects introduced in [9] and studied also in [4]. This model can also be seen as a special case of the model introduced in [2], where the network is composed of parallel links and the setup cost is determined through the cost-sharing rule.

2 Model and Preliminaries

We consider a *job-scheduling* setting with identical machines and identical (unit-size) jobs. There is a set of machines $M = \{M_1, M_2, \ldots\}$ of unlimited size,[4] each associated with an *activation cost*, B. An instance of our problem is given as a tuple (n, B), where n denotes the number of jobs. An assignment method produces an assignment $s = (s_1, \ldots, s_n)$ of jobs into machines, where $s_j \in M$ denotes the machine to which job j is assigned. We use the terms assignment, schedule, and profile interchangeably. The load of a machine M_i in a schedule s, denoted $L_i(s)$, is the number of jobs assigned to M_i in s.

Given a job-scheduling setting and an activation cost B, a *job-scheduling game* is induced where the set of players is the set of jobs, and the action space of each player is the set of machines. The cost function of job j in a given schedule is the sum of two components: the load on j's machine and j's share in the machine's activation cost. It is assumed that the activation cost B is shared equally between all the jobs that use a particular machine. That is, given a profile s in which $s_j = M_i$, the cost of job j is $c_j(s) = L_i(s) + \frac{B}{L_i(s)}$. We denote the cost of a job that is assigned to a machine with load x by $c(x)$, where $c(x) = x + \frac{B}{x}$. It can be easily verified that the cost function exhibits the following structure.

Observation 1 *The function $c(x) = x + B/x$ for $x > 0$ attains its minimum at $x = \sqrt{B}$, is decreasing for $x \in (0, \sqrt{B})$, and increasing for $x > \sqrt{B}$.*

Practically, the input to the cost function is an integral value. If B is a perfect square, then the integral load achieving the minimal cost is exactly $\sqrt{B}$. For example, if $B = 100$, then being assigned to a machine with load 10 is optimal. In general, however, the optimal integral load (i.e., the load that minimizes the cost function) may be either $\left\lfloor \sqrt{B} \right\rfloor$ or $\left\lceil \sqrt{B} \right\rceil$, and for some values of B it may be both. For example, if $B = 12$ then both 3 and 4 are optimal loads, as

[4] In any instance, though, the number of machines will clearly be less than n.

$c(3) = c(4) = 12$. We denote an optimal load by $\ell^* = \ell^*(B)$. Assuming a unique integral optimal load, it is easy to verify that the cost function is decreasing for $x \in [1, \ell^*]$ and increasing for $x \geq \ell^*$. For two optimal integral loads, $\ell^* - 1$ and ℓ^*, the cost function is decreasing for $x \in [1, \ell^* - 1]$ and increasing for $x \geq \ell^*$.

An assignment $s \in S$ is a *pure Nash equilibrium* (NE) if no job $j \in N$ can benefit from unilaterally deviating from its machine to another machine (possibly a new machine). In our game, this implies that for every job j assigned to M_i and every $i' \neq i$, it holds that $c(L_i(s)) \leq \min(c(1), c(L_{i'}(s) + 1))$.

3 Convergence of Best-Response Dynamics

Best-Response Dynamics (BRD) is a local search method where in each step some player plays its best-response, given the strategies of the others. In systems where the agents always reach a Nash equilibrium after repeatedly performing improvement steps, the notion of a pure Nash equilibrium is well justified. This section explores the convergence rate of best-response dynamics into a pure NE.

In the general case, in which jobs have arbitrary lengths and the activation cost of a machine is shared by the jobs proportionally to their length, BRD is not guaranteed to converge to a NE [9]. In contrast, if the jobs are identical, then the induced game is equivalent to a congestion game with n resources [19]. One can easily verify that the function $\Phi(s) = \sum_i \left(B \cdot H_{\ell_i} + \frac{1}{2}\ell_i^2\right)$, where ℓ_i denotes the number of jobs on machine i, $H_0 = 0$, and $H_k = 1 + 1/2 + \ldots + 1/k$, is a potential function for the game. Convergence to a NE is guaranteed in potential games, but the convergence time might be exponential.

Here, we study the convergence time of BRD of unit-length jobs. We show that the convergence in general might take $\Omega(n \log \frac{n}{B})$ moves, and propose a specific BRD that ensures convergence within a linear number of moves. Specifically,

> **Max-cost BRD:** At every time step, a job that incurs the highest cost among those who can benefit from migration, is chosen to perform its best-response move (where ties are broken arbitrarily).

The analysis of the convergence rate of BRD and max-cost BRD (MC-BRD hereafter) is quite complicated and requires several preparations and terminology. Recall that all jobs assigned to a machine with load x incur the same cost $c(x) = x + B/x$. We denote by ℓ^* a load achieving minimal cost. By Observation 1, ℓ^* may be either $\lfloor \sqrt{B} \rfloor$ or $\lceil \sqrt{B} \rceil$, and for some values of B it may be both. For simplicity, throughout this section we assume a unique optimal load. All the results hold also for the case of two optimal loads, where minor straightforward modifications are required in the proofs.

We denote by ℓ_i^t the load of machine M_i at time t, i.e., *before* the migration of iteration t takes place. A machine that has load at least (respectively, smaller than) ℓ^* is said to be a *high* (*low*) machine.

We observe that if at some iteration a job migrates to a low machine, then in subsequent iterations that machine will attract more jobs up to load at least ℓ^*. Indeed, since $c(\ell + 1) < c(\ell)$ for $\ell < \ell^*$, a low best-response machine continues to be a best response until it is filled up to load at least ℓ^*. Formally,

Observation 2 *If at some iteration t there is a migration to a low machine M_i such that $\ell_i^t = \ell^* - x$ for some $x > 0$, then the following $x - 1$ iterations will involve migrations to M_i.*[5]

Properties of MC-BRD: By the design of the MC-BRD process and as a direct corollary of Observation 1, every migration in the MC-BRD process is from either the lowest or the highest machine into either the lowest-high or the highest-low machine (see Figure 3).

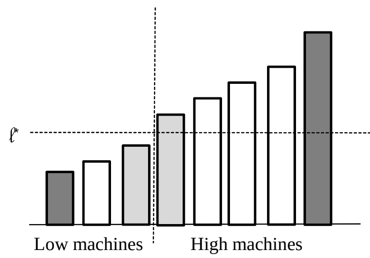

Fig. 1. MC-BRD process. Every migration is from one of the extreme machines into one of the middle grey machines.

Since all jobs on a particular machine share the same cost, the MC-BRD process can be described as if it acts on machines rather than on jobs. Specifically, in every iteration t, one job migrates from machine M_i to machine $M_k, k \neq i$, where (i) $c(\ell_k^t + 1)$ is minimal, (ii) $c(\ell_k^t + 1) < c(\ell_i^t)$, and (iii) $c(\ell_i^t)$ is maximal among all the machines from which a beneficial migration exists. While the MC-BRD process does not specify which job is migrating from M_i, for simplicity we will assume a LIFO (last in first out) job selection rule. Specifically, the job that entered M_i last is the one to migrate. If all jobs on M_i were assigned to it in the initial configuration, then an arbitrary job is selected. Since the BRD-process can be characterized by the load-vector of the machines in every time step, the number of iterations is independent of the job-selection rule. Consequently, our analysis of the convergence rate of MC-BRD applied with a LIFO job-selection rule is valid for any MC-BRD process.

Note that M_i, the machine from which a job is selected to migrate in iteration t, is not necessarily the machine for which $c(\ell_i^t)$ is maximal. For example, suppose that $B = 100$ and there are two active machines, a low one with load 3, and a high one with load 33. It is easy to verify that $c(4) < c(33) < c(3) < c(34)$. In this case, $c(3)$ is the maximal cost, but jobs on the low machine have no beneficial move (since $c(34) > c(3)$). On the other hand, jobs on the high machine wish to migrate to the low one (since $c(4) < c(33)$). Thus, the high machine is the one selected by MC-BRD to perform a migration, although the low machine is the one incurring max-cost. Clearly, such a case can only occur if the machine that incurs the max-cost is itself the best-response machines, as summarized in the following observation.

[5] It is possible that the system reaches a NE and the BRD process terminates before $x - 1$ iterations are performed.

Observation 3 *If at time t the machine M_i that incurs max-cost is not the one from which a job is selected to migrate in MC-BRD, then $c(\ell_i^t + 1)$ is the best-response, in particular, this implies that M_i is low.*

We next observe that in MC-BRD, if at some iteration a job leaves some low machine, then in the following iterations all the jobs assigned to that machine leave it one by one until the machine empties out. To see this, note that $c(\ell-1) > c(\ell)$ for $\ell < \ell^*$; thus, if a low machine incurs the highest cost, it continues to incur the highest cost after its load decreases. It remains to show that if a beneficial move out of M_i exists when it has load $\ell < \ell^*$, then it is also beneficial to leave M_i when it has load $\ell - 1$. This is ensured by Observation 3. Specifically, if it is not beneficial, then $c(\ell)$ is the cost of the best-response machine. But this is impossible since $c(\ell)$ was the max-cost in the previous iteration.

Observation 4 *If at some iteration t there is a migration from a low machine M_i such that $\ell_i^t = \ell^* - x$ for some $x > 0$, then the following $\ell^* - x - 1$ iterations will involve migrations from M_i.*

We are now ready to state the bound on the convergence rate of MC-BRD. As shown in the full version [10], the following bound is almost tight.

Theorem 1. *For every job scheduling game with identical jobs, every MC-BRD process converges to a NE within at most $\max\{\frac{3n}{2} - 3, n - 1\}$ steps.*

In contrast to MC-BRD, the convergence time of arbitrary BRD, might not be linear in n.

Theorem 2. *There exists a job scheduling game with identical jobs and a BRD process such that the convergence time to a NE is $\Omega(n \log \frac{n}{B})$.*

While the convergence rate of general BRD is super-linear, the following theorem establishes an upper bound of n^2. Closing the gap remains open.

Theorem 3. *For every job scheduling game with identical jobs, every BRD process converges to a NE within at most n^2 steps.*

It is interesting to compare our results to those established for the standard model that considers only the negative congestion effects (i.e., where a job's cost is simply the load of its chosen machine). It has been shown by [7] that if the order of the jobs performing their best-response moves is determined according to their lengths (i.e., longer job first), then best-response dynamics reaches a pure Nash equilibrium within at most n improvement steps. In contrast, if the jobs move in an arbitrary order, then convergence to a Nash equilibrium might take an exponential number of steps. These results imply that for the special case of equal-length jobs, convergence occurs within at most n steps. Our results provide evidence that when there are conflicting congestion effects, it might take longer to reach a Nash equilibrium. Nevertheless, for the special case of max-cost BRD, the consideration of positive congestion effects (through activation costs) does not lead to a longer convergence time.

Acknowledgement This work was partially supported by the Israel Science Foundation (# 1219/09), by the Leon Recanati Fund of the Jerusalem School of Business

Administration, the Google Inter-university center for Electronic Markets and Auctions, and the People Programme (Marie Curie Actions) of the European Unions 7th Framework Programme (FP7/2007-2013) under REA grant agreement number 274919.

References

1. H. Ackermann, H. Röglin, and B. Vöcking. On the impact of combinatorial structure on congestion games. In FOCS'06, pp. 613-622 (2006).
2. E. Anshelevich, A. Dasgupta, J. M. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In FOCS'04, pp. 295–304 (2004).
3. B.Awerbuch, Y.Azar, A.Epstein, V.S.Mirrokni, and A.Skopalik. Fast convergence to nearly optimal solutions in potential games. In ACMEC'08, pp. 264–273 (2008).
4. B. Chen and S. Gürel. Efficiency analysis of load balancing games with and without activation costs. *Journal of Scheduling*, 15(2) (2011).
5. S. Chien and A. Sinclair. Convergence to approximate Nash equilibria in congestion games. In SODA'07, pp. 169–178 (2007).
6. A. Epstein, M. Feldman, and Y. Mansour. Strong equilibrium in cost sharing connection games. *Games and Economic Behavior*, 67(1) (2009).
7. E. Even-Dar, A. Kesselman, and Y. Mansour. Convergence time to nash equilibria. In ICALP'03, pp. 502–513 (2003).
8. A. Fabrikant, C. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In STOC'04, pp. 604-612 (2004).
9. M. Feldman and T. Tamir. Conflicting congestion effects in resource allocation games. *Operations Research*, 60(3): 529-540, (2012).
10. M. Feldman and T. Tamir. Conflicting congestion effects in resource allocation games. *Full version. http://www.faculty.idc.ac.il/tami/Papers/BRDcoco.pdf.*
11. D. Fotakis, S Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. Spirakis. The structure and complexity of nash equilibria for a selfish routing game. In ICALP'02, pp. 510–519 (2002).
12. D. Fotakis. Congestion games with linearly independent paths: convergence time and price of anarchy. In *Theory Comput. Syst.*, 47(1):113–136 (2010).
13. M. Hoefer and A. Skopalik. Stability and convergence in selfish scheduling with altruistic agents. In WINE'09, pp. 616–622 (2009).
14. S. Ieong, R.Mcgrew, E. Nudelman, Y. Shoham and Q. Sun, Fast and Compact: A simple class of congestion games. In AAAI'05, pp. 489–494 (2005).
15. R. Johari and S. Kumar. Congestible services and network effects. In EC'10, pp. 93–94 (2010).
16. E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. *Computer Science Review*,3(2): 65–69 (2009).
17. D. Monderer and L.S. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143 (1996).
18. C. Papadimitriou. Algorithms, games, and the Internet. In STOC'01, pp.749–753 (2001).
19. R. W. Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2:65–67 (1973).
20. T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236 – 259 (2002).
21. V. Syrgkanis. The complexity of equilibria in cost sharing games. In WINE'10, pages 366–377 (2010).