Information Processing Letters 130 (2018) 36–39

Contents lists available at ScienceDirect

Information Processing Letters

www.elsevier.com/locate/ipl

Generalized mirror descents with non-convex potential functions in atomic congestion games: Continuous time and discrete time

Po-An Chen

Institute of Information Management, National Chiao Tung University, Taiwan

ARTICLE INFO

Article history:
Received 27 August 2016
Accepted 16 October 2017
Available online 20 October 2017
Communicated by Tsan-sheng Hsu

Keywords:
Analysis of algorithms
Mirror descents
Congestion games
Non-convex
Convergence

ABSTRACT

When playing certain specific classes of no-regret algorithms such as multiplicative updates and replicator dynamics in atomic congestion games, some previous convergence analyses were done with the standard Rosenthal potential function in terms of mixed strategy profiles (i.e., probability distributions on atomic flows), which could be non-convex. In several other works, the convergence, when playing the mirror-descent algorithm (a more general family of no-regret algorithms including multiplicative updates, gradient descents, etc.), was guaranteed with a convex potential function in terms of nonatomic flows as an approximation of the Rosenthal one. The convexity of the potential function provides convenience for analysis. One may wonder if the convergence of mirror descents can still be guaranteed directly with the non-convex Rosenthal potential function. In this paper, we answer the question affirmatively for *discrete-time* generalized mirror descents with the *smoothness* property (similarly adopted in many previous works for congestion games and markets) and for *continuous-time* generalized mirror descents with the *separability* of regularization functions.

1. Introduction

Playing learning algorithms in repeated games has been extensively studied within this decade, especially with generic no-regret algorithms [3,8] and various specific no-regret algorithms [9,10,4–6,12]. Multiplicative updates are played in atomic congestion games to reach pure Nash equilibria with high probability with full information in [9], and in load-balancing games to converge to certain mixed Nash equilibria with bulletin-board posting in [10]. The family of mirror descents [1], including multiplicative updates, gradient descents, and many more classes of algorithms by choosing the corresponding regularization functions, are generalized in the bulletin-board model, and even with only bandit feedbacks in (respectively, nonatomic and atomic) congestion games to guarantee convergence to approximate equilibria [4–6].

For the analyses for multiplicative updates and replicator dynamics (which can be seen as a continuous variant of multiplicative updates) in [9,12] were accomplished with a standard Rosenthal potential function, in terms of mixed strategy profiles (probability distributions on atomic flows), which may be *non-convex*. On the other hand, the convergence analyses for multiplicative updates in [10] and mirror descents (even more general than multiplicative updates) in [4–6] were done with a *convex* potential function in terms of nonatomic flows as an approximation of the Rosenthal one.[1] It can be seen that though with different

[1] There is a tradeoff of an error from the nonlinearity of cost functions if the implication of reaching approximate equilibria, which is not our focus here, is needed besides the convergence guarantee ([10,5,6]).

E-mail address: poanchen@nctu.edu.tw.

https://doi.org/10.1016/j.ipl.2017.10.003

techniques, the properties from convexity and others help there, especially for bounding the convergence time.

One may wonder if the convergence of mirror descents can still be guaranteed directly with the non-convex Rosenthal potential function. Nevertheless, as far as we are concerned, there has not been an explicit convergence analysis. Recently in [13,11], it is shown that gradient descent converges to a local minimizer almost surely with random initialization if it ever converges. Their focus in there is to ensure that when gradient descent does converge, it is to minimisers, never to saddle points; still, there is no convergence guarantee for mirror descents in atomic congestion games. In this paper, we affirmatively answer the question whether generalized mirror descents converge with the non-convex Rosenthal potential function in atomic congestion games: convergences guaranteed for *discrete-time* generalized mirror descents with the help of the *smoothness* property (similarly adopted in [2,7,4–6]) and for *continuous-time* generalized mirror descents with the *separability* of regularization functions (as in [7]). Note that although we do not show what our generalized mirror descent converges to here, the result of [13,11] indicates that it converges to minimizers at least for gradient descents.

2. Preliminaries

We need to formally define the game and potential function before we proceed. We consider the following atomic congestion game, described by $(N, E, (\mathcal{S}_i)_{i\in N}, (c_e)_{e\in E})$, where N is the set of players, E is the set of m edges (resources), $\mathcal{S}_i \subseteq 2^E$ is the collection of allowed paths (subsets of resources) for player i, and c_e is the cost function of edge e, which is a nondecreasing function of the amount of flow on it. Let us assume that there are n players, each path has length at most m, and each player has a flow of amount $1/n$ to route.

The mixed strategy of each player i is to send her entire flow on a single path, chosen randomly according to some distribution over her allowed paths, which can be represented by a $|\mathcal{S}_i|$-dimensional vector $p_i = (p_{i\gamma})_{\gamma\in\mathcal{S}_i}$, where $p_{i\gamma} \in [0,1]$ is the probability of choosing path γ. Let $\mathcal{K}_i$ denote the feasible set of all such $p_i \in [0,1]^{|\mathcal{S}_i|}$ for player i, and let $\mathcal{K} = \mathcal{K}_1 \times \ldots \times \mathcal{K}_n \subseteq \mathbb{R}^d$ be the feasible set of all such joint strategy profiles $p = (p_1, \ldots, p_n)$ of the n players.

Γ_i is a random path, and Γ_{-i} is a vector $\Gamma = (\Gamma_i)_{i\in N}$ except Γ_i. We consider the following Rosenthal potential function in terms of mixed strategy profile p ([9,12]), where c_i is player i's cost function and c_e is the cost function of edge e.

$$\Psi(p) = \mathbf{E}_{\Gamma_{-i}}[\Phi_{-i}(\Gamma_{-i})] + \mathbf{E}_{\Gamma}[c_i(\Gamma)] \tag{1}$$

$$= \mathbf{E}_{\Gamma_{-i}}[\sum_{e\in E}\sum_{j=1}^{K_i(e)} c_e(j/n)] + \mathbf{E}_{\Gamma}[c_i(\Gamma)], \tag{2}$$

where $K_i(e)$ is a random variable defined as $|\{j : j \neq i, e \in \Gamma_j\}|$ and $\Phi_{-i}(\Gamma_{-i})$ is defined as $\sum_{e\in E}\sum_{j=1}^{K_i(e)} c_e(j/n)$. We assume that any edge cost function e satisfies the conditions that for any $y \in [0,1]$, $c_e(0) = 0$, $c_e(1) = 1$, and $c_e(y) \le by$ for a nonnegative constant b. Let Γ_i have value γ, which is a strategy for i, with probability $p_{i\gamma}$, and player i's expected individual cost $\mathbf{E}_\Gamma[c_i(\Gamma)]$ is defined as follows.

$$\mathbf{E}_\Gamma[c_i(\Gamma)] = \sum_{\gamma\in\mathcal{S}_i} p_{i\gamma} c_{i\gamma}, \tag{3}$$

where player i's expected individual cost for choosing γ over the randomness from the other players is $c_{i\gamma} = \mathbf{E}_{\Gamma_{-i}}[c_i(\gamma, \Gamma_{-i})] = \sum_{e\in\gamma} \mathbf{E}_{\Gamma_{-i}}[c_e(\frac{1+K_i(e)}{n})]$.

Given the essential definitions above, we have the following properties to be used in Section 3.

$$\frac{\partial \Psi(p)}{\partial p_{i\gamma}} = c_{i\gamma}, \tag{4}$$

$$\frac{\partial^2 \Psi(p)}{\partial p_{i\gamma}\partial p_{j\gamma'}} = \sum_{e\in\gamma\cap\gamma'} \frac{\partial c_{i\gamma}}{\partial p_{j\gamma'}}. \tag{5}$$

3. Dynamics and convergence

We introduce our continuous-time generalized mirror-descent algorithm first and then the discrete-time generalized mirror-descent algorithm, along with the convergence results taking advantage of the separability of regularization functions and the smoothness of the potential function, respectively.

For both continuous- and discrete-time dynamics, $\eta_i > 0$ is player i's learning rate, $R_i : \mathcal{K}_i \to \mathbb{R}$ is player i's regularization function, and $B^{R_i}(\cdot,\cdot)$ is the Bregman divergence with respect to R_i defined as

$$\mathcal{B}^{R_i}(u_i, v_i) = R_i(u_i) - R_i(v_i) - \langle \nabla R_i(v_i), u_i - v_i\rangle$$

for $u_i, v_i \in \mathcal{K}_i$. For example, when choosing $R_i(u_i) = ||u_i||_2^2/2$ for each i, and thus $\mathcal{B}^{R_i}(u_i, v_i) = ||u_i - v_i||_2^2/2$ for $u_i, v_i \in \mathcal{K}_i$, each player plays the gradient-descent algorithm; when choosing $R_i(u_i) = \sum_s u_{i,s}\ln(u_{i,s} - u_{i,s})$ for each i, and thus $\mathcal{B}^{R_i}(u_i, v_i) = \sum_s u_{i,s}\ln(u_{i,s}/v_{i,s})$ for $u_i, v_i \in \mathcal{K}_i$, each player plays the multiplicative update algorithm.

Continuous-time generalized mirror descents

We consider the case where each regularization function is a *separable* function, i.e., it is of the form $\sum_s R_i(u_{i,s})$ for a 1-dimensional function $R_i : \mathbb{R} \to \mathbb{R}$ [7]. We have to point out that the regularization functions that many well-known algorithms use are actually separable. For instance, $||u_i||_2^2/2 = \sum_s ||u_{i,s}||_2^2/2$ for each i, and $\sum_s u_{i,s}\ln(u_{i,s} - u_{i,s})$ is by definition separable for each i. The continuous update rule with respect to $\mathcal{B}^{R_i}(\cdot,\cdot)$ is defined as follows. Note that each player i can have her own η_i in such continuous-time *generalized* mirror descents. Let

$$p_i(\epsilon) = \arg\min_{z_i\in\mathcal{K}_i}\{\eta_i\langle(c_{i\gamma}^t)_\gamma, z_i\rangle + \frac{1}{\epsilon}\mathcal{B}^{R_i}(z_i, p_i^t)\}. \tag{6}$$

$$\frac{dp_i^t}{dt} = \lim_{\epsilon\to 0}\frac{p_i(\epsilon) - p_i^t}{\epsilon}. \tag{7}$$