Generalized Mirror Descents with Non-Convex Potential Functions in Atomic Congestion Games

Po-An Chen

Institute of Information Management
National Chiao Tung University, Taiwan
poanchen@nctu.edu.tw

Abstract. When playing specific classes of no-regret algorithms (especially, multiplicative updates) in atomic congestion games, some previous convergence analyses were done with a standard Rosenthal potential function in terms of mixed strategy profiles (probability distributions on atomic flows), which may not be convex. In several other works, the convergence analysis was done with a convex potential function in terms of nonatomic flows as an approximation of the Rosenthal one in terms of distributions. It can be seen that though with different techniques, the properties from convexity help there, especially for convergence time.

However, it would be always a valid question to ask if convergence can still be guaranteed directly with the Rosenthal potential function, playing mirror descents individually in atomic congestion games. We answer this affirmatively by showing the convergence, individually playing discrete mirror descents with the help of the smoothness property similarly adopted in many previous works for congestion games and Fisher (and some more general) markets and individually playing continuous 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, 7] and various specific no-regret algorithms [8, 9, 4, 5, 10]. Multiplicative updates are played in atomic congestion games to reach pure Nash equilibria with high probability with full information in [8], and in load-balancing games to converge to certain mixed Nash euiqlibria with bulletin-board posting in [9]. The family of mirror descents [1], which include multiplicative updates, gradient descents, and many more classes of algorithms, are generalized and played with bulletin-board posting, and even with only bandit feedbacks in (respectively, nonatomic and atomic) congestion games to guarantee convergence to approximate equilibria [4, 5].

For specific classes of no-regret algorithms (especially, multiplicative updates), the analysis in [8, 10] was done with a standard Rosenthal potential function in terms

In: A. Kononov et al. (eds.): DOOR 2016, Vladivostok, Russia, published at http://ceur-ws.org

of mixed strategy profiles (probability distributions on atomic flows), which may not be convex. In [9, 4, 5] (even with mirror descents, more general than multiplicative updates), the convergence analysis was done with a convex potential function in terms of nonatomic flows as an approximation of the Rosenthal one in terms of distributions.[1] It can be seen that though with different techniques, the properties from convexity and others help there, especially for convergence time.

However, it would be always a valid question to ask *if convergence can still be guaranteed directly with the Rosenthal potential function, playing mirror descents individually in atomic congestion games.* In this paper, we answer this affirmatively by showing the convergence using *discrete* generalized mirror descents with the help of the *smoothness* property similarly adopted in [2, 6, 4, 5] and using *continuous* generalized mirror descents with the *separability* of regularization functions as in [6].

2 Preliminaries

We need to formally define the game and potential function before we proceed. We consider the following atomic congestion game, described by $(N, E, (\mathcal{S}_i)_{i\in N}, (c_e)_{e\in E})$ 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 player has at most d allowed paths, each path has length at most m, each allowed path of a player intersects at most k allowed paths (including that path itself) of that player, 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. It turns out to be more convenient for us to represent each player's strategy p_i by an equivalent form $x_i = (1/n)p_i$, where $1/n$ is the amount of flow each player has. That is, for every $i \in N$ and $\gamma \in \mathcal{S}_i$, we have that $x_{i\gamma} = (1/n)p_{i\gamma} \in [0, 1/n]$ and $\sum_{\gamma\in\mathcal{S}_i} x_{i\gamma} = 1/n$. Let $\mathcal{K}_i$ denote the feasible set of all such $x_i \in [0, 1/n]^{|\mathcal{S}_i|}$ for player i, and let $\mathcal{K} = \mathcal{K}_1 \times ... \times \mathcal{K}_n$, which is the feasible set of all such joint strategy profiles $(x_1, ..., x_n)$ of the n players.

Γ_i is a random subset of resources, and Γ_{-i} is a vector $(\Gamma_i)_{i\in N}$ except Γ_i. We consider the following Rosenthal potential function in terms of mixed strategy profile p ([8, 10]).

$$\Psi(p) = \mathbf{E}_{\Gamma_{-i}}[\Phi_{-i}(\Gamma_{-i})] + \mathbf{E}_{(\Gamma_{-i},\Gamma_i)}[c^i(\Gamma_i)] \tag{1}$$

$$\equiv \mathbf{E}_{\Gamma_{-i}}[\sum_{e\in E}\sum_{j=1}^{K_i(e)} c_e(j)] + \mathbf{E}_{(\Gamma_{-i},\Gamma_i)}[c^i(\Gamma_i)], \tag{2}$$

where $K_i(e)$ is a random variable defined as $K_e(\Gamma_{-i}) \equiv |\{j : j \neq i, e \in \Gamma_j\}|$. Let Γ_i have value γ, which is a strategy for i, with probability $p_{i\gamma}$, and $\mathbf{E}_{(\Gamma_{-i},\Gamma_i)}[c^i(\Gamma_i)]$ is

[1] There is a tradeoff of an error from the nonlinearity of cost functions if the implication of reaching approximate equilibria is needed besides the convergence guarantee ([9, 5]).

defined as follows.

$$\mathbf{E}_{(\varGamma_{-i},\varGamma_i)}[c^i(\varGamma_i)] = \sum_{\gamma\in\mathcal{S}_i} p_{i\gamma}c_{i\gamma}, \tag{3}$$

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

We have the following properties.

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

3 Dynamics and Convergence

Discrete Generalized Mirror Descents

We consider the following discrete update rule for player i.

$$p_i^{t+1} = \arg\min_{z_i\in\mathcal{K}_i}\{\eta_i\langle(c_{i\gamma}^t)_\gamma, z_i\rangle + \mathcal{B}^{R_i}(z_i,p_i^t)\} \tag{4}$$
$$= \arg\min_{z_i\in\mathcal{K}_i}\mathcal{B}^{R_i}(z_i, p_i^t - \eta_i\nabla_i\varPsi(p^t)) \tag{5}$$

The equality in (5) holds because $(c_{i\gamma}^t)_\gamma = \nabla_i\varPsi(p^t)$. Here, $\eta_i > 0$ is some learning rate, $R_i : \mathcal{K}i \to \mathbb{R}$ is some 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 $u_i, v_i \in \mathcal{K}_i$.

We can ask about the actual convergence of the value of $\varPsi$ to some minimum and thus an approximate equilibrium and/or aim for convergence of the mixed strategy profile pt. The difficulty to deal with such $\varPsi$ is that it is not convex. There may be no way to bound the convergence time, and there can be multiple minima to converge to. Nevertheless, with the "smoothness" as defined in [2, 6] and [4, 5], restated in the following, it is still possible to show convergence of $\varPsi$ and approximate equilibria in atomic congestion games.

Definition 1 (Smoothness). *We say that $\varPsi$ is λ-smooth with respect to $(R_1,...,R_n)$ if for any two inputs $p = (p_1,...,p_n)$, $p' = (p'_1,...,p'_n) \in \mathcal{K}$,*

$$\varPsi(p') = \varPsi(p) + \langle\nabla\varPsi(p), p' - p\rangle + \lambda\sum_{i=1}^{n}\mathcal{B}^{R_i}(p'_i, p_i). \tag{6}$$

Note that λ has to be properly set in different games.

The convergence follows from the assumption that for each i, $\eta_i \le 1/\lambda$ and thus $\lambda \le 1/\eta_i$, and the implication from the λ-smoothness condition (7).

Theorem 1. *For any $t \ge 0$, $\varPsi(p^{t+1}) \le \varPsi(p^t)$.*

Continuous Generalized Mirror Descents

We consider the case where R_i is a *separable* function, i.e., it is of the form $\sum_{\gamma \in \mathcal{S}_i} R_i(u_{i\gamma})$. The continuous update rule with respect to $\mathcal{B}^{R_i}(\cdot, \cdot)$ is defined as follows. Let

$$p_i(\epsilon) = \arg\min_{z_i \in \mathcal{K}_i} \{\eta_i \langle (c_{i\gamma}^t)_\gamma, z_i \rangle + \frac{1}{\epsilon}\mathcal{B}^{R_i}(z_i, p_i^t)\}. \quad (7)$$

$$\frac{dp_i^t}{dt} = \lim_{\epsilon \to 0} \frac{p_i(\epsilon) - p_i^t}{\epsilon}. \quad (8)$$

The following claim and the convexity of R_i give the convergence result.

Claim (LEMMA *4.2 of [6]).* For any i and γ, $\frac{dp_{i\gamma}^t}{dt} = \frac{-\nabla_{i\gamma}\Psi(p^t)}{R_i''(p_{i\gamma}^t)}$ where $R_i(p_i^t) = \sum_{\gamma \in \mathcal{S}_i} R_i(p_{i\gamma}^t)$.

Theorem 2. *For any $t \geq 0$, $\frac{d\Psi(p^t)}{dt} \leq 0$.*

4 Discussion and Future Work

More challengingly, we can also consider partial-information models with such trickier potential function Ψ. Does the bandit algorithm in [5] extend here to guarantee convergence? Can the gradient still be estimated accurately (close to the real one with high probability)? Are there other suitable gradient estimation methods as well that work in other less or more stringent partial-information models other than the bandit one?

Another direction is application to the analysis of the average-case price of anarchy as the application of "pointwise" convergence for the average-case price of anarchy in [10].

References

1. Amir Beck and Marc Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3):167–175, 2003.
2. B. Birnbaum, N. Devanur, and L. Xiao. Distributed algorithms via gradient descent for fisher markets. In *Proc. 12th ACM Conference on Electronic Commerce*, 2011.
3. A. Blum, E. Even-Dar, and K. Ligett. Routing without regret: On convergence to nash equilibria of regret-minimizing algorithms in routing games. In *Proc. 25th ACM Symposium on Principles of Distributed Computing*, 2006.
4. P.-A. Chen and C.-J. Lu. Generalized mirror descents in congestion games with splittable flows. In *Proc. 13th International Conference on Autonomous Agents and Multiagent Systems*, 2014.
5. P.-A. Chen and C.-J. Lu. Playing congestion games with bandit feedbacks (extended abstract). In *Proc. 14th International Conference on Autonomous Agents and Multiagent Systems*, 2015.
6. Y. K. Cheung, R. Cole, and N. Devanur. Tatonnement beyond gross substitutes?: gradient descent to the rescue. In *Proc. 45th ACM Symposium on theory of computing*, pages 191–200, 2013.
7. E. Even-dar, Y. Mansour, and U. Nadav. On the convergence of regret minimization dynamics in concave games. In *Proc. 41st Annual ACM Symposium on Theory of Computing*, 2009.

8. R. Kleinberg, G. Piliouras, and E. Tardos. Multiplicative updates outperform generic no-regret learning in congestion games. In *Proc. 40th ACM Symposium on Theory of Computing*, 2009.
9. R. Kleinberg, G. Piliouras, and E. Tardos. Load balancing without regret in the bulletin board model. *Distributed Computing*, 24, 2011.
10. I. Panageas and G. Piliouras. From pointwise convergence of evolutionary dynamics to average case analysis of decentralized algorithms. 2015. https://arxiv.org/abs/1403.3885v3.