

Convergence of Heterogeneous Distributed Learning in Stochastic Routing Games

Syrine Krichene

Walid Krichene

Roy Dong

Alexandre Bayen

Abstract—We study convergence properties of distributed learning dynamics in repeated stochastic routing games. The game is stochastic in that each player observes a stochastic vector, the conditional expectation of which is equal to the true loss (almost surely). In particular, we propose a model in which every player m follows a stochastic mirror descent dynamics with Bregman divergence D_{ψ_m} and learning rates $\eta_t^m = \theta_m t^{-\alpha_m}$. We prove that if all players use the same sequence of learning rates, then their joint strategy converges almost surely to the equilibrium set. If the learning dynamics are heterogeneous, that is, different players use different learning rates, then the joint strategy converges to equilibrium in expectation, and we give upper bounds on the convergence rate. This result holds for general routing games (no smoothness or strong convexity assumptions are required).

These results provide a distributed learning model that is robust to measurement noise and other stochastic perturbations, and allows flexibility in the choice of learning algorithm of each player. The results also provide estimates of convergence rates, which are confirmed in simulation.

I. INTRODUCTION

The routing game is a model for congestion on networks shared by selfish players, and dates back to the seminal work of Wardrop [24] and Beckmann et al. [3], who modeled congestion on transportation networks. The formulation has since been extended, and used to model routing and congestion in transportation and communication networks, see for example Ozdaglar and Srikant [19], Roughgarden [22] and the references therein. In particular, the routing game is known to be a potential game, which makes the computation of its equilibrium set tractable.

Beyond computing the equilibrium set of the game, different models have been proposed for how players reach equilibrium using learning dynamics, or repeated play. These learning dynamics can be used to model how players adjust their strategies, for example by changing the route a driver chooses for his/her daily commute, or by changing the flow distribution of a router in a communication network. Some models propose player dynamics in continuous time. For example, Fischer and Vöcking [8] study the replicator dynamics, and Hofbauer and Sandholm [10] study general dynamics which satisfy a positive correlation condition with the gradient field of the potential. Other studies propose discrete-time models, i.e. repeated play. Such discrete models

are to some extent a more natural model for decision dynamics, especially in transportation networks, where the time scale of the adjustment is day-to-day. Blum et al. [4] study no-regret dynamics for the routing game, and derive convergence rates for the time-averaged strategies. Kleinberg et al. [12] study multiplicative weight updates, a particular class of no-regret algorithms, and give estimates of the convergence rates. In [15], we propose another family of no-regret learning algorithms, based on distributed mirror descent, and we prove convergence rates.

In this article, we seek to design learning algorithms which are robust to measurement noise and other stochastic perturbations. More precisely, we extend the previous models by assuming that, at each iteration, instead of observing the exact loss vector, the player rather observes a stochastic vector, the conditional expectation of which is (almost surely) equal to the loss of the routing game. This is a natural extension for two reasons. First, the routing game model may not capture all sources of congestion, some random variables (such as weather and incidents) may affect the delay. The stochastic perturbation vector can be used to model these effects. Second, the delay measurements can be inherently noisy, both in transportation and communication networks.

Stochastic learning models have been studied in online learning theory, e.g. Bubeck and Cesa-Bianchi [6], adaptive control theory, e.g. Kumar and Varaiya [16], as well as convex optimization, e.g. Nemirovski et al. [17] and Juditsky et al. [11]. Adapting ideas from these works, we propose a family of stochastic distributed learning algorithms and study their convergence. In our model, we assume that every population m follows a stochastic mirror descent algorithm, with Bregman divergence D_{ψ_m} and learning rates (η_t^m) . We first prove that if the learning is homogeneous, that is, all players use the same sequence of learning rates (but not necessarily the same Bregman divergences), then the joint strategy converges almost surely to the set of equilibria. This result holds for general convex potential functions, without additional assumptions on regularity or strong convexity, in particular, convergence holds even when the equilibrium is not unique, an assumption which was usually made when proving similar almost sure convergence, e.g. [5], and which we manage to relax. We also prove that under an additional strong convexity assumption, the variance converges to zero and we give a bound on the convergence rate. Then we show that in the heterogeneous case, if $\eta_t^m = \theta_m t^{-\alpha_m}$, then the joint strategy converges in expectation to the set of equilibria, at a $\mathcal{O}(\ln t / t^{\min(\alpha_{\min}, 1 - \alpha_{\max})})$ rate (the fastest corresponding rate is $\mathcal{O}(\ln t / \sqrt{t})$). This result also holds

Syrine Krichene is with the ENSIMAG school of Computer Sciences and Applied Mathematics of Grenoble. syrine.krichene@ensimag.grenoble-inp.fr

Walid Krichene, Roy Dong and Alexandre Bayen are with the department of Electrical Engineering and Computer Sciences at the University of California at Berkeley. [{krichene, roy.dong, bayen}@eecs.berkeley.edu"> {krichene, roy.dong, bayen}@eecs.berkeley.edu](mailto)

for general convex potentials, and concerns the convergence of the actual sequence of joint strategies rather than the sequence of averages (a standard but weaker result, which holds for any sublinear regret algorithm, see e.g. Blum et al. [4]). The proof uses a combination of standard regret analysis and a recent induction technique proposed by Shamir and Zhang [23].

These results provide a model of distributed learning that is robust to stochastic perturbations, and which is rather flexible since it allows different populations to use different learning algorithms with different learning rates. And while our study is motivated by the routing game, these results also apply to a more general distributed learning setting, in which agents optimize over a product of convex compact feasible sets. In the routing game, the feasible sets are the probability simplexes over the set of paths available to each population.

In Section II, we give a formal definition of the routing game and its equilibrium set, and review some basic properties. In Section III, we define the distributed stochastic mirror descent algorithm and give a basic regret bound. In Section IV, we prove the convergence results and derive convergence rates. These results are formulated in the general distributed learning setting, and applied to the routing game. Finally, we illustrate these results on numerical examples both in the strongly and non-strongly convex cases, which show that the numerical convergence rates are consistent with the rates predicted by our theorems.

II. THE ROUTING GAME AND THE LEARNING MODEL

In this section, we review definitions and basic properties. The routing game is given by a directed graph (V, E) , and a finite set of populations indexed by $m \in \mathcal{M}$. Each population is given by an origin vertex and a destination vertex on the graph, a set of paths \mathcal{P}_m connecting them, and a total population mass R_m (corresponding to the rate of drivers or the rate of packets to be sent from the origin to the destination). At each iteration, every population chooses a flow distribution $x_m^{(t)}$ in the scaled simplex $\mathcal{X}_m = \{x \in \mathbb{R}_+^{|\mathcal{P}_m|} : \sum_{i \in \mathcal{P}_m} x_i = R_m\}$, that is, a population distributes its total mass on its available paths.

The product distribution $(x_m^{(t)})_{m \in \mathcal{M}}$ is denoted by $x^{(t)}$, and determines the loss of each population m , given by $\langle x_m^{(t)}, \ell_m(x^{(t)}) \rangle$, where $\langle \cdot, \cdot \rangle$ is the canonical inner product on $\mathbb{R}^{|\mathcal{P}_m|}$ and $\ell_m(x^{(t)}) \in \mathbb{R}_+^{|\mathcal{P}_m|}$ is the vector of losses on paths in \mathcal{P}_m , given as follows: for a given product distribution x , the loss on a path $p \in \mathcal{P}_m$ is the sum of edge losses along the path, $\ell_{m,p}(x) = \sum_{e \in p} c_e(M_e x)$ where $c_e : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is an increasing function, called the edge congestion function, and $M \in \{0, 1\}^{|E| \times \sum_m |\mathcal{P}_m|}$ is an incidence matrix such that $M_{e,p} = 1$ if $e \in p$ and 0 otherwise, so that $M_e x$ is exactly the total flow on edge e . Note in particular that the distribution of one population may affect the loss of other populations because of shared edges on the network.

A. Wardrop equilibria and the Rosenthal potential

The network is said to be at a Wardrop equilibrium¹ if no population m can decrease its loss by unilaterally changing its distribution x_m . The set of Wardrop equilibria will be denoted by \mathcal{W} . Then we have the following simple variational characterization of \mathcal{W} :

$$\begin{aligned} x^* \in \mathcal{W} &\Leftrightarrow \forall m, \forall x_m \in \mathcal{X}_m, \langle \ell_m(x^*), x_m - x_m^* \rangle \geq 0 \\ &\Leftrightarrow \forall x \in \times_m \mathcal{X}_m, \langle \ell(x^*), x - x^* \rangle \geq 0. \end{aligned} \quad (1)$$

Rosenthal [21] proposed a convex potential function f defined on the product of simplexes $\mathcal{X} = \times_{m \in \mathcal{M}} \mathcal{X}_m$, that satisfies the following property: for all $x \in \mathcal{X}$, $\nabla_{x_m} f(x) = \ell_m(x)$. In other words, the loss vector field $\ell(\cdot)$ is exactly the gradient of the potential f . As a consequence, the set of minimizers of f over \mathcal{X} is exactly the set of equilibria \mathcal{W} (the variational characterization (1) is equivalent to the first order optimality condition of f). The potential function plays a central role in the analysis of distributed learning dynamics: indeed, a sequence $(x^{(t)})$ converges to the equilibrium set \mathcal{W} if and only if $f(x^{(t)})$ converges to f^* , the minimum of f on \mathcal{X} . Thus, the joint dynamics of the population can be viewed as a distributed descent of the potential function.

The potential function is defined as follows: $f(x) = \sum_{e \in E} \int_0^{M_e x} c_e(u) du$, and can be viewed as the composition of the convex function $\tilde{f} : y \mapsto \sum_{e \in E} \int_0^{y_e} c_e(u) du$, and the linear function $x \mapsto Mx$. Note that when c_e are differentiable, \tilde{f} is strongly convex if and only if the derivatives $c'_e(\cdot)$ are bounded below by a positive constant on \mathbb{R}_+ , however, even when \tilde{f} is strongly convex, f may not be, due to the composition with the incidence matrix M . For this reason, it will be important to derive convergence guarantees that do not rely on strong convexity of the potential function.

Algorithm 1 Distributed Stochastic Mirror Descent (DSMD) with Bregman divergences D_{ψ_m} and learning rates (η_t^m) .

```

for  $t \in \mathbb{N}$  do
  for each population  $m \in \mathcal{M}$  do
    Play  $x_m^{(t)}$ 
    Observe  $\hat{\ell}_m^{(t)}$  with  $\ell^{(t)} \triangleq \mathbb{E}[\hat{\ell}^{(t)} | \mathcal{F}_{t-1}] \stackrel{a.s.}{=} \nabla f(x^{(t)})$ 
    Update
      
$$x_m^{(t+1)} = \arg \min_{x_m \in \mathcal{X}_m} \langle \hat{\ell}_m^{(t)}, x_m \rangle + \frac{1}{\eta_t^m} D_{\psi_m}(x_m, x_m^{(t)}) \quad (2)$$


```

B. Learning Model

Previous models for discrete-time population dynamics, such as the models used by Blum et al. [4] and Krichene et al. [15], assume that at iteration t , each population chooses a flow distribution $x_m^{(t)}$, then at the end of the iteration, population m observes the loss vector $\ell_m(x^{(t)})$.

We extend this learning model to allow stochastic perturbations of the loss vectors. That is, we now suppose that

¹The notion of Wardrop equilibrium is due to Wardrop [24]. If every population is identified with a measurable set of players with no atoms, then the Wardrop equilibrium corresponds to a Nash equilibrium up to a null set of players (see Krichene et al. [13]).

at iteration t , population m observes a stochastic vector $\hat{\ell}_m^{(t)}$, which is unbiased in the sense that $\mathbb{E}[\hat{\ell}_m^{(t)} | \mathcal{F}_{t-1}] = \ell_m(x^{(t)})$ a.s., where (\mathcal{F}_t) is the natural filtration of the process $(\hat{\ell}^{(t)})$.

Given this setting, we propose the following model for distributed learning: we suppose that each population m applies a stochastic mirror descent algorithm with Bregman divergence D_{ψ_m} and learning rates (η_t^m) . The sequence (η_t^m) is assumed to be positive decreasing. The learning model is summarized in Algorithm 1. In the next Section, we will give a motivation and geometric interpretation of the model, then analyze its convergence in Section IV.

III. DISTRIBUTED STOCHASTIC MIRROR DESCENT

We start by giving a brief review of the mirror descent method (MD). MD is a general method for constrained convex optimization, proposed by Nemirovsky and Yudin [18]. Consider the problem

$$\begin{aligned} & \text{minimize} && f(x) \\ & \text{subject to} && x \in \mathcal{X} \end{aligned}$$

where $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is a convex function defined on a convex, compact set $\mathcal{X} \subset \mathbb{R}^d$, and call f^* the minimum value of f on \mathcal{X} .




Fig. 1. Mirror Descent iteration

MD can be interpreted, as observed by Beck and Teboulle [2], as minimizing, at each iteration t , a local approximation of the objective function around the current iterate, as follows:

$$\begin{aligned} x^{(t+1)} &= \arg \min_{x \in \mathcal{X}} \left\langle \nabla f(x^{(t)}), x \right\rangle + \frac{1}{\eta_t} D_{\psi_t}(x, x^{(t)}) \\ &= \arg \min_{x \in \mathcal{X}} f(x^{(t)}) + \left\langle \nabla f(x^{(t)}), x - x^{(t)} \right\rangle + \frac{1}{\eta_t} D_{\psi_t}(x, x^{(t)}). \end{aligned}$$

The first term, $f(x^{(t)}) + \langle \nabla f(x^{(t)}), x - x^{(t)} \rangle$ is the first order Taylor approximation of the function around the current iterate, and the second term, $D_{\psi_t}(x, x^{(t)})$, is the Bregman divergence of x with respect to $x^{(t)}$, and penalizes deviations from $x^{(t)}$. The parameter η_t is a generalized step size which determines the tradeoff between the two terms. This is illustrated in Figure 1.

The Bregman divergence associated to the convex, differentiable function $\psi : \mathcal{X} \rightarrow \mathbb{R}$ is given by

$$D_{\psi}(x, y) = \psi(x) - \psi(y) - \langle \nabla \psi(y), x - y \rangle.$$

By convexity of ψ , the Bregman divergence is non-negative and convex in its first argument. The function ψ is said to be μ -strongly convex w.r.t. a reference norm $\|\cdot\|$ (not necessarily the Euclidean norm) if for all x, y , $D_{\psi}(x, y) \geq \frac{\mu}{2} \|x - y\|^2$. In particular, if we take $\psi(x) = \frac{1}{2} \|x\|_2^2$, then the Bregman divergence is $D_{\psi}(x, y) = \frac{1}{2} \|x - y\|_2^2$, and the mirror descent update becomes a projected gradient descent update. In this sense, MD is a generalization of gradient descent. For a more detailed discussion on the properties of Bregman divergences, see for example Banerjee et al. [1].

We now adapt this convex optimization setting to the stochastic routing model proposed in Section II. First, since the set of joint strategies in the routing game is the Cartesian product of simplexes, we will assume that the feasible set \mathcal{X} is a Cartesian product of convex feasible sets, $\mathcal{X} = \times_{m \in \mathcal{M}} \mathcal{X}_m$.

A. Stochastic optimization

Next, since the loss vectors are assumed stochastic (and the loss vector field coincides with the gradient field of the Rosenthal potential), we will assume that, at each iteration t , we have access to a stochastic vector $\hat{\ell}^{(t)}$, such that the conditional expectation $\ell^{(t)} \triangleq \mathbb{E}[\hat{\ell}^{(t)} | \mathcal{F}_{t-1}]$ is equal to $\nabla f(x^{(t)})$ almost surely, where (\mathcal{F}_t) is the natural filtration of the stochastic process $(\hat{\ell}^{(t)})$.

This stochastic framework is motivated by the stochastic routing game, however, it is also useful in modeling problems where computing the exact gradient can be prohibitively expensive, such as large-scale optimization problems where the objective function is a sum of individual convex loss terms over a large set of samples, that is, $f(x) = \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} \ell(x, z_i)$. A cheap estimate of the subgradient of f can then be obtained by randomly drawing a small subset of samples $\mathcal{I}^{(t)} \subset \mathcal{I}$, and defining $\hat{\ell}^{(t)}$ to be $\hat{\ell}^{(t)} = \frac{1}{|\mathcal{I}^{(t)}|} \sum_{i \in \mathcal{I}^{(t)}} \nabla_x \ell(x^{(t)}, z_i)$.

In the Distributed Stochastic Mirror Descent model (DSMD) described in Algorithm 1, the next iterate $x^{(t+1)}$ is obtained by minimizing, over each feasible set \mathcal{X}_m , the sum of terms $\langle \hat{\ell}_m^{(t)}, x_m \rangle + \frac{1}{\eta_t^m} D_{\psi_m}(x_m, x_m^{(t)})$. In this case, the sequence of iterates $(x^{(t)})$ also forms a stochastic process, such that $x^{(t)}$ is \mathcal{F}_{t-1} measurable. We further assume that the first iterate $x^{(1)}$ is deterministic, i.e. \mathcal{F}_0 is trivial.

We will make the following important assumptions:

- Assumption 1:* (i) For each m , the Bregman divergence D_{ψ_m} is strongly convex w.r.t. a reference norm $\|\cdot\|$, and bounded on \mathcal{X}_m , that is, there exists $\mu_m > 0$ and $D_m > 0$ such that for all $x, y \in \mathcal{X}_m$, $\frac{\mu_m}{2} \|x - y\|^2 \leq D_{\psi_m}(x, y) \leq D_m$.
- (ii) The noisy gradient vectors are uniformly square integrable in the dual norm, that is, there exists $G > 0$ such that for all t , $\mathbb{E}[\|\hat{\ell}^{(t)}\|_*^2] \leq G^2$.

B. A fundamental lemma

The following lemma is an essential step in proving the convergence results in the next section. It is a straightforward generalization of Lemma 2.1 in Nemirovski et al. [17].

Proposition 1: Consider the DSMD algorithm with Bregman divergences D_{ψ_m} and decreasing learning rates (η_t^m) and let $(x^{(t)})$ be the resulting stochastic process. Then for all m , all for all τ , and all $\mathcal{F}_{\tau-1}$ measurable x_m ,

$$D_{\psi_m}(x_m, x_m^{(\tau+1)}) \leq D_{\psi_m}(x_m, x_m^{(\tau)}) - \eta_\tau \langle \hat{\ell}_m^{(\tau)}, x_m^{(\tau)} - x_m \rangle + \frac{\eta_\tau^2}{2\mu_m} \|\hat{\ell}_m^{(\tau)}\|_*^2, \quad (3)$$

additionally, for all $t_2 > t_1 \geq 1$, and all \mathcal{F}_{t_1-1} -measurable x_m ,

$$\sum_{\tau=t_1}^{t_2} \mathbb{E} \left[\langle \hat{\ell}_m^{(\tau)}, x_m^{(\tau)} - x_m \rangle \right] \leq \frac{\mathbb{E} \left[D_{\psi_m}(x_m, x_m^{(t_1)}) \right]}{\eta_{t_1}^m} + D_m \left(\frac{1}{\eta_{t_2}^m} - \frac{1}{\eta_{t_1}^m} \right) + \frac{G^2}{2\mu_m} \sum_{\tau=t_1}^{t_2} \eta_\tau^m. \quad (4)$$

This bound can be interpreted as a regret bound, if we adopt an online learning point of view, as in Zinkevich [25], Hazan et al. [9], Cesa-Bianchi and Lugosi [7]. The sum $\sum_{\tau=t_1}^{t_2} \langle \hat{\ell}_m^{(\tau)}, x_m^{(\tau)} - x_m \rangle$ is the cumulative regret of population m with respect to the stationary process x_m .

IV. CONVERGENCE GUARANTEES OF HOMOGENEOUS AND HETEROGENEOUS DSMD

In this section, we study convergence properties of the DSMD model proposed in Algorithm 1. The results will be formulated for general convex functions over a Cartesian product set $\mathcal{X} = \times_{m \in \mathcal{M}} \mathcal{X}_m$. These results can then be applied to the stochastic routing game by taking \mathcal{X}_m to be the scaled simplex over paths \mathcal{P}_m , and the objective function f to be the Rosenthal potential.

A. Convergence of averages

We begin with a relatively weak result, which concerns convergence of the sequence of averages, $\bar{x}^{(t)} = \frac{\sum_{\tau=1}^t x^{(\tau)}}{t}$. We show that if the algorithm has a sublinear regret in expectation, that is, $\frac{\sum_{\tau=1}^t \mathbb{E}[\langle \hat{\ell}^{(\tau)}, x^{(\tau)} - x \rangle]}{t}$ converges to 0 as $t \rightarrow \infty$, then $f(\mathbb{E}[\bar{x}^{(t)}])$ converges to f^* . This can be guaranteed when (η_t^m) have appropriate decay rates, as in the following Corollary.

Corollary 1: Consider the DSMD method with $\eta_t^m = \theta_m t^{-\alpha_m}$, with $\theta_m > 0$ and $\alpha_m \in (0, 1)$. Then

$$f(\mathbb{E}[\bar{x}^{(t)}]) - f^* \leq \sum_{m \in \mathcal{M}} \left(\frac{D_m}{\theta_m t^{1-\alpha_m}} + \frac{\theta_m}{1-\alpha_m} \frac{G^2}{2\mu_m} \frac{1}{t^{\alpha_m}} \right).$$

The bound is $\mathcal{O}(t^{-\min(\alpha_{\min}, 1-\alpha_{\max})})$, where α_{\min} and α_{\max} are, respectively, the smallest and largest rate α_m .

Proof: Let x^* be a minimizer of f over \mathcal{X} . We have by convexity of f and the fact that $\ell^{(\tau)} \stackrel{a.s.}{=} \nabla f(x^{(\tau)})$,

$$\begin{aligned} f(\mathbb{E}[\bar{x}^{(t)}]) - f^* &\leq \frac{\sum_{\tau=1}^t \mathbb{E}[f(x^{(\tau)}) - f^*]}{t} \\ &\leq \sum_{m \in \mathcal{M}} \frac{\sum_{\tau=1}^t \mathbb{E}[\langle \hat{\ell}_m^{(\tau)}, x_m^{(\tau)} - x_m^* \rangle]}{t}. \end{aligned}$$

Then by Proposition 1, and since x^* is \mathcal{F}_0 -measurable (deterministic),

$$\begin{aligned} f(\mathbb{E}[\bar{x}^{(t)}]) - f^* &\leq \frac{\sum_{\tau=1}^t \mathbb{E}[\langle \hat{\ell}^{(\tau)}, x^{(\tau)} - x \rangle]}{t} \\ &\leq \sum_{m \in \mathcal{M}} \frac{\mathbb{E}[D_{\psi_m}(x_m^*, x_m^{(1)})]}{\eta_1^m t} + \frac{D_m}{t} \left(\frac{1}{\eta_t^m} - \frac{1}{\eta_1^m} \right) + \frac{G^2}{2\mu_m} \frac{\sum_{\tau=1}^t \eta_\tau^m}{t} \\ &\leq \sum_{m \in \mathcal{M}} \frac{D_m}{t \eta_t^m} + \frac{G^2}{2\mu_m} \frac{\sum_{\tau=1}^t \eta_\tau^m}{t}. \end{aligned}$$

Finally, since $u \mapsto u^{-\alpha}$ is a decreasing function over \mathbb{R}^+ , $\sum_{\tau=1}^t \eta_\tau^m \leq \theta_m \int_0^t u^{-\alpha_m} du = \frac{\theta_m}{1-\alpha_m} t^{1-\alpha_m}$, which concludes the proof. \blacksquare

B. Almost sure convergence to \mathcal{X}^* in the homogeneous case

Let us denote the set of minimizers by $\mathcal{X}^* \triangleq \arg \min_{x \in \mathcal{X}} f(x)$. We say that a sequence $x^{(t)}$ converges to \mathcal{X}^* , and write $x^{(t)} \rightarrow \mathcal{X}^*$, if $d(x^{(t)}, \mathcal{X}^*) \rightarrow 0$ as $t \rightarrow \infty$ where d is the distance to the set defined as $d(x, \mathcal{X}^*) = \inf_{y \in \mathcal{X}^*} \|x - y\|$.

The bound of Proposition 1 can be used to show that $x^{(t)} \xrightarrow{a.s.} \mathcal{X}^*$, if the learning is homogeneous, i.e. all populations use the same learning rates (but not necessarily the same Bregman divergence).

Theorem 1: Consider the DSMD method, and suppose that the learning rates (η_t) do not depend on the population m . Suppose further that $\sum_{t=1}^{\infty} \eta_t = \infty$ and $\sum_{t=1}^{\infty} \eta_t^2 < \infty$. Then

$$x^{(t)} \xrightarrow{a.s.} \mathcal{X}^*.$$

Note that a similar almost sure convergence result is known in the stochastic optimization literature, see for example Bottou [5]. However, such results assume uniqueness of the minimizer. We relax this uniqueness assumption by analyzing the sequence of Bregman divergences from the set of minimizers to the current iterate, as follows.

Proof: First, we define a Bregman divergence on the Cartesian product $\mathcal{X} = \times_m \mathcal{X}_m$ as follows: for all $x \in \mathcal{X}$, let

$$\psi(x) = \sum_m \psi_m(x_m). \quad (5)$$

Then the corresponding Bregman divergence is $D_\psi(x, y) = \sum_m D_{\psi_m}(x_m, y_m)$. Now let $D_\psi(\mathcal{X}^*, x) = \inf_{x^* \in \mathcal{X}^*} D_\psi(x^*, x)$. Since D_ψ is continuous and \mathcal{X}^* is compact (it is a closed subset of the compact set \mathcal{X}), we have that the infimum is attained and $D_\psi(\mathcal{X}^*, \cdot)$ is continuous. By continuity of $D_\psi(\mathcal{X}^*, \cdot)$ and compactness of \mathcal{X} , we have $x^{(t)} \rightarrow \mathcal{X}^*$ if and only if $D_\psi(\mathcal{X}^*, x^{(t)}) \rightarrow 0$.

We start by showing that $D_\psi(\mathcal{X}^*, x^{(t)})$ converges almost surely, using a semi martingale convergence theorem. From inequality (3) in Proposition 1, summing over m and letting μ be the harmonic mean of $(\mu_m)_{m \in \mathcal{M}}$, we have

$$D_\psi(x, x^{(\tau+1)}) \leq D_\psi(x, x^{(\tau)}) - \eta_\tau \langle \hat{\ell}^{(\tau)}, x^{(\tau)} - x \rangle + \frac{\eta_\tau^2}{2\mu} \|\hat{\ell}^{(\tau)}\|_*^2.$$

In particular, taking x to be equal to $x^{*\langle\tau\rangle} \triangleq \arg \min_{x^* \in \mathcal{X}^*} D_\psi(x^*, x^{(\tau)})$, we have

$$\begin{aligned} D_\psi(\mathcal{X}^*, x^{(\tau+1)}) &\leq D_\psi(x^{*\langle\tau\rangle}, x^{(\tau+1)}) \\ &\leq D_\psi(x^{*\langle\tau\rangle}, x^{(\tau)}) - \eta_\tau \langle \hat{\ell}^{(\tau)}, x^{(\tau)} - x^{*\langle\tau\rangle} \rangle + \frac{\eta_\tau^2}{2\mu} \|\hat{\ell}^{(\tau)}\|_*^2 \\ &= D_\psi(\mathcal{X}^*, x^{(\tau)}) - \eta_\tau \langle \hat{\ell}^{(\tau)}, x^{(\tau)} - x^{*\langle\tau\rangle} \rangle + \frac{\eta_\tau^2}{2\mu} \|\hat{\ell}^{(\tau)}\|_*^2. \end{aligned}$$

Then, we take conditional expectations with respect to $\mathcal{F}_{\tau-1}$, and observe that since $x^{(\tau)}$ and $x^{*\langle\tau\rangle}$ are $\mathcal{F}_{\tau-1}$ -measurable,

$$\begin{aligned} \mathbb{E} \left[\langle \hat{\ell}^{(\tau)}, x^{(\tau)} - x^{*\langle\tau\rangle} \rangle \middle| \mathcal{F}_{\tau-1} \right] &= \left\langle \mathbb{E} \left[\hat{\ell}^{(\tau)} \middle| \mathcal{F}_{\tau-1} \right], x^{(\tau)} - x^{*\langle\tau\rangle} \right\rangle \\ &\stackrel{\text{a.s.}}{=} \left\langle \nabla f(x^{(\tau)}), x^{(\tau)} - x^{*\langle\tau\rangle} \right\rangle \\ &\geq f(x^{(\tau)}) - f^*. \end{aligned}$$

Therefore, we have a.s.

$$\begin{aligned} \mathbb{E} \left[D_\psi(\mathcal{X}^*, x^{(\tau+1)}) \middle| \mathcal{F}_{\tau-1} \right] &\leq D_\psi(\mathcal{X}^*, x^{(\tau)}) \\ &\quad - \eta_\tau (f(x^{(\tau)}) - f^*) + \frac{\eta_\tau^2}{2\mu} \mathbb{E} \left[\|\hat{\ell}^{(\tau)}\|_*^2 \middle| \mathcal{F}_{\tau-1} \right]. \end{aligned}$$

By the previous inequality, and the fact that

- (i) $\eta_\tau (f(x^{(\tau)}) - f^*) \geq 0$, and
- (ii) $\sum_{\tau=1}^{\infty} \frac{\eta_\tau^2}{2\mu} \|\hat{\ell}^{(\tau)}\|_*^2$ is a.s. finite by assumption on (η_t) and $\mathbb{E}[\|\hat{\ell}^{(\tau)}\|_*^2]$,

the process $(D_\psi(\mathcal{X}^*, x^{(t)}))$ is an almost super-martingale. Therefore, by the Robbins and Siegmund [20] convergence Theorem, $D_\psi(\mathcal{X}^*, x^{(\tau)})$ converges almost surely, and $\sum_{\tau=1}^{\infty} \eta_\tau (f(x^{(\tau)}) - f^*)$ is a.s. finite.

To show that the limit of $D_\psi(\mathcal{X}^*, x^{(t)})$ is almost surely 0, suppose that for some realization, $D_\psi(\mathcal{X}^*, x^{(t)})$ converges to $d > 0$, then there exists $T > 0$ such that for all $t \geq T$, $D_\psi(\mathcal{X}^*, x^{(t)}) > d/2$. Let $\delta \triangleq \inf_{\{x \in \mathcal{X}: D_\psi(\mathcal{X}^*, x) > \frac{d}{2}\}} f(x) - f^*$. By continuity of f , we have that $\delta > 0$, thus $\sum_{\tau=1}^{\infty} \eta_\tau (f(x^{(\tau)}) - f^*) \geq \delta \sum_{t \geq T} \eta_\tau = \infty$. Therefore the event $\lim_{t \rightarrow \infty} D_\psi(\mathcal{X}^*, x^{(t)}) > 0$ is a subset of the event $\sum_{\tau} \eta_\tau (f(x^{(\tau)}) - f^*) = \infty$, which proves that $D_\psi(\mathcal{X}^*, x^{(\tau)}) \xrightarrow{\text{a.s.}} 0$. ■

C. Convergence in the homogeneous, strongly convex case

In this section, we assume that f is μ_f -strongly convex with respect to D_ψ , in the following sense: for all $x, y \in \mathcal{X}$, $f(x) \geq f(y) + \langle \nabla f(y), x - y \rangle + \mu_f \max(D_\psi(x, y), D_\psi(y, x))$.

We show that under this assumption, the variance of the iterates converges to 0. First, we observe that by strong convexity of ψ , we have $\mathbb{E}[\|x - x^*\|^2] \leq \frac{2}{\mu} \mathbb{E}[D_\psi(x^*, x)]$, thus it suffices to show the convergence of $\mathbb{E}[D_\psi(x^*, x^{(t)})]$. First, we show the following Lemma.

Lemma 1: Suppose f is μ_f -strongly convex with respect to D_ψ , and let x^* be the minimizer of f over \mathcal{X} . Then for all $y \in \mathcal{X}$, $\langle \nabla f(y), y - x^* \rangle \geq 2\mu_f D_\psi(x^*, y)$.

Proof: By strong convexity of f , we have

$$\begin{aligned} f(x^*) &\geq f(y) + \langle \nabla f(y), x^* - y \rangle + \mu_f D_\psi(x^*, y) \\ f(y) &\geq f(x^*) + \mu_f D_\psi(x^*, y), \end{aligned}$$

and we conclude by summing the two inequalities. ■

We now show convergence of $\mathbb{E}[D_\psi(x^*, x^{(t)})]$.

Proposition 2: Suppose that f is μ_f -strongly convex with respect to D_ψ , where ψ is defined as the sum of ψ_m , as in equation (5). Then the homogeneous DSMD algorithm with homogeneous learning rates (η_t) guarantees

$$\mathbb{E} \left[D_\psi(x^*, x^{(t+1)}) \right] \leq (1 - 2\mu_f \eta_t) \mathbb{E} \left[D_\psi(x^*, x^{(t)}) \right] + \frac{G^2}{2\mu} \eta_t^2.$$

Proof: We start from inequality (3) in Proposition 1. Taking expectation with $x_m = x_m^*$, and summing over m , it follows that

$$\begin{aligned} \mathbb{E}[D_\psi(x^*, x^{(t+1)})] &\leq \mathbb{E}[D_\psi(x^*, x^{(t)})] \\ &\quad - \eta_t \mathbb{E}[\langle \hat{\ell}^{(t)}, x^{(t)} - x^* \rangle] + \frac{\mathbb{E}[\|\hat{\ell}^{(t)}\|_*^2]}{2\mu} \eta_t^2, \end{aligned}$$

and since $\mathbb{E}[\hat{\ell}^{(t)} | \mathcal{F}_{t-1}] = \nabla f(x^{(t)})$ a.s., we have by Lemma 1

$$- \mathbb{E}[\langle \hat{\ell}^{(t)}, x^{(t)} - x^* \rangle] \leq -2\mu_f \mathbb{E} \left[D_\psi(x^*, x^{(t)}) \right].$$

Combining the two inequalities, we have the claim. ■

Theorem 2 (Convergence of variance for $\eta_t = \Theta(t^{-\alpha})$): Suppose that f is μ_f strongly convex with respect to D_ψ , and consider the homogeneous DSMD with learning rates $\eta_t = \frac{\theta}{2\mu_f t^\alpha}$, $\alpha \in (0, 1)$. Then for all $t \geq t_0$

$$\mathbb{E} \left[D_\psi(x^*, x^{(t)}) \right] \leq \frac{C}{t^\alpha}, \quad (6)$$

where $t_0 = \left\lceil \left(\frac{2\alpha}{\theta} \right)^{\frac{1}{1-\alpha}} \right\rceil$ and $C = \max(Dt_0^\alpha, \frac{G^2\theta}{4\mu\mu_f^2})$.

Proof: We show the claim by induction on $t \geq t_0$. For $t = t_0$, we have by assumption on D_ψ

$$\mathbb{E} \left[D_\psi(x^*, x^{(t_0)}) \right] \leq D \leq \frac{C}{t_0^\alpha}.$$

Now suppose by induction that $\mathbb{E} \left[D_\psi(x^*, x^{(t)}) \right] \leq \frac{C}{t^\alpha}$. Then by Proposition 2,

$$\begin{aligned} \mathbb{E} \left[D_\psi(x^*, x^{(t+1)}) \right] &\leq \left(1 - \frac{\theta}{t^\alpha} \right) \mathbb{E} \left[D_\psi(x^*, x^{(t)}) \right] + \frac{G^2\theta^2}{8\mu\mu_f^2} \frac{1}{t^{2\alpha}} \\ &\leq \left(1 - \frac{\theta}{t^\alpha} \right) \frac{C}{t^\alpha} + \frac{G^2\theta^2}{8\mu\mu_f^2} \frac{1}{t^{2\alpha}} \\ &= \frac{C}{(t+1)^\alpha} \left[\left(\frac{t+1}{t} \right)^\alpha \left(1 + \frac{1}{t^\alpha} \left(-\theta + \frac{G^2\theta^2}{8\mu\mu_f^2 C} \right) \right) \right] \\ &\leq \frac{C}{(t+1)^\alpha} \exp \left[\frac{\alpha}{t} + \frac{1}{t^\alpha} \left(\frac{G^2\theta^2}{8\mu\mu_f^2 C} - \theta \right) \right]. \end{aligned}$$

To conclude, it suffices to prove that the exponential term is less than one. By definition of C , $\frac{G^2\theta^2}{8\mu\mu_f^2 C} - \theta \leq -\frac{\theta}{2}$, thus the exponential term is less than one if $\frac{\alpha}{t} - \frac{\theta}{2t^\alpha} \leq 0$, i.e. $t \geq \left(\frac{2\alpha}{\theta} \right)^{\frac{1}{1-\alpha}}$, which is true if $t \geq t_0$. Therefore we have

$$\mathbb{E} \left[D_\psi(x^*, x^{(t+1)}) \right] \leq \frac{C}{(t+1)^\alpha},$$

which completes the induction. ■

We observe that when $\alpha = 1$, the inequality $\frac{1}{t} - \frac{\theta}{2t} \leq 0$ holds whenever $\theta \geq 2$, in which case $t_0 = 1$, and we recover the

$\mathcal{O}(\frac{1}{t})$ bound of Shamir and Zhang [23] for the Euclidean case with $\eta_t = \frac{1}{\mu_f t}$.

In fact, we can show that $\mathbb{E}[D_\psi(x^*, x^{(t)})]$ converges to 0 for any sequence of learning rates such that $\eta_t \rightarrow 0$ and $\sum_t \eta_t = \infty$.

Lemma 2: Let $(d^{(t)})$ be a sequence of non-negative numbers that satisfy the following inequality

$$d^{(t+1)} \leq (1 - \nu_t)d^{(t)} + \Gamma\nu_t^2,$$

for some $\Gamma > 0$ and a positive decreasing sequence ν_t with $\sum_t \nu_t = \infty$. Then for all T with $\nu_T \leq 1$, and all $t > T$,

$$d^{(t)} \leq \nu_T \Gamma + d^{(T)} e^{-\sum_{\tau=T}^{t-1} \nu_\tau}.$$

Combining Proposition 2 and Lemma 2, we can take $\nu_t = \mu_f \eta_t$ and $\Gamma = \frac{G^2}{2\mu\mu_f^2}$ to obtain

$$\mathbb{E}[D_\psi(x^*, x^{(t)})] \leq \frac{G^2}{2\mu\mu_f} \eta_T + D e^{-\sum_{\tau=T}^{t-1} \mu_f \eta_\tau},$$

for any $t > T$ such that $\mu_f \eta_T \leq 1$. In particular, this proves that $\mathbb{E}[D_\psi(x^*, x^{(t)})] \rightarrow 0$.

D. Convergence in the heterogeneous case

We now analyze the convergence of $\mathbb{E}[f(x^{(t)})]$ for general convex functions, under the heterogeneous DSMD model with learning rates $\eta_t^m = \theta_m t^{-\alpha_m}$, $\alpha_m \in (0, 1)$. Shamir and Zhang [23] prove the convergence of the last iterate in the case of stochastic gradient descent (a special case of SMD) for $\alpha = \frac{1}{2}$. Our analysis uses their technique and extends it to the SMD method with heterogeneous learning rates and general $\alpha_m \in (0, 1)$.

Theorem 3: Consider DSMD with learning rates $\eta_t^m = \theta_m t^{-\alpha_m}$. Then for all $t \geq 1$,

$$\begin{aligned} & \mathbb{E}[f(x^{(t)})] - f^* \\ & \leq \sum_m \left(\frac{D_m}{\theta_m} \frac{1}{t^{1-\alpha_m}} + \frac{\theta_m G^2}{2\mu_m(1-\alpha_m)} \frac{1}{t^{\alpha_m}} \right) (2 + \ln t). \end{aligned} \quad (7)$$

This bound is a $\mathcal{O}(t^{-\min(\alpha_{\min}, 1-\alpha_{\max})} \ln t)$.

Proof: Let t be fixed. Adapting the proof of Shamir and Zhang [23], we define S_k to be

$$S_k = \frac{1}{k+1} \sum_{\tau=t-k}^t \mathbb{E}[f(x^{(\tau)})].$$

We have by convexity of f ,

$$\begin{aligned} \sum_{\tau=t-k}^t \mathbb{E}[f(x^{(\tau)}) - f(x^{(t-k)})] & \leq \sum_{\tau=t-k}^t \mathbb{E}[\langle \ell^{(\tau)}, x^{(\tau)} - x^{(t-k)} \rangle] \\ & = \sum_{m \in \mathcal{M}} \sum_{\tau=t-k}^t \mathbb{E}[\langle \ell_m^{(\tau)}, x_m^{(\tau)} - x_m^{(t-k)} \rangle]. \end{aligned}$$

and applying Proposition 1 with $t_1 = t - k$, $t_2 = t$, and $x_m = x_m^{(t-k)}$, which is \mathcal{F}_{t-k-1} -measurable, we have

$$\begin{aligned} & \sum_{\tau=t-k}^t \mathbb{E}[\langle \ell_m^{(\tau)}, x_m^{(\tau)} - x_m^{(t-k)} \rangle] \\ & \leq D_m \left(\frac{1}{\eta_t^m} - \frac{1}{\eta_{t-k}^m} \right) + \frac{G^2 \theta_m}{2\mu_m} \sum_{\tau=t-k}^t \tau^{-\alpha_m} \\ & \leq \frac{D_m (t^{\alpha_m} - (t-k)^{\alpha_m})}{\theta_m} + \frac{\theta_m G^2 (t^{1-\alpha_m} - (t-k-1)^{1-\alpha_m})}{2\mu_m(1-\alpha_m)} \end{aligned}$$

where we used the integral bound $\sum_{\tau=t-k}^t \tau^{-\alpha_m} \leq \int_{t-k-1}^t u^{-\alpha_m} du$. To simplify this bound, we can use the fact that $-(t-k-1)^{-\alpha_m} \leq -t^{-\alpha_m}$ and write

$$t^{1-\alpha_m} - (t-k-1)^{1-\alpha_m} \leq \frac{t - (t-k-1)}{t^{\alpha_m}} = \frac{k+1}{t^{\alpha_m}}.$$

Similarly, $t^{\alpha_m} - (t-k)^{\alpha_m} \leq \frac{k}{t^{1-\alpha_m}}$. Therefore

$$\begin{aligned} & \sum_{\tau=t-k}^t \mathbb{E}[f(x^{(\tau)}) - f(x^{(t-k)})] \\ & \leq \sum_m \left(\frac{D_m}{\theta_m} \frac{k+1}{t^{1-\alpha_m}} + \frac{\theta_m G^2}{2\mu_m(1-\alpha_m)} \frac{k+1}{t^{\alpha_m}} \right). \end{aligned} \quad (8)$$

Dividing by $k+1$, we have

$$-\mathbb{E}[f(x^{(t-k)})] \leq -S_k + \sum_m \left(\frac{D_m}{\theta_m} \frac{1}{t^{1-\alpha_m}} + \frac{\theta_m G^2}{2\mu_m(1-\alpha_m)} \frac{1}{t^{\alpha_m}} \right).$$

Therefore

$$\begin{aligned} S_{k-1} & = \frac{1}{k} \left((k+1)S_k - \mathbb{E}[f(x^{(t-k)})] \right) \\ & \leq S_k + \sum_m \left(\frac{D_m}{\theta_m} \frac{1}{t^{1-\alpha_m}} + \frac{\theta_m G^2}{2\mu_m(1-\alpha_m)} \frac{1}{t^{\alpha_m}} \right) \frac{1}{k}. \end{aligned} \quad (9)$$

We seek to derive a bound on $\mathbb{E}[f(x^{(t)})] - f^* = S_0 - f^*$, thus, we can sum inequality (9) for $k \in \{1, \dots, t\}$, we have

$$\begin{aligned} S_0 - f^* & \leq S_{t-1} - f^* \\ & + \sum_{m \in \mathcal{M}} \left(\frac{D_m}{\theta_m} \frac{1}{t^{1-\alpha_m}} + \frac{\theta_m G^2}{2\mu_m(1-\alpha_m)} \frac{1}{t^{\alpha_m}} \right) \sum_{k=1}^{t-1} \frac{1}{k}, \end{aligned} \quad (10)$$

and from Corollary 1, we have

$$S_{t-1} - f^* \leq \sum_{m \in \mathcal{M}} \left(\frac{D_m}{\theta_m t^{1-\alpha_m}} + \frac{\theta_m G^2}{2\mu_m(1-\alpha_m)} \frac{1}{t^{\alpha_m}} \right). \quad (11)$$

Finally, combining the inequalities (10) and (11) and using the fact that $\sum_{k=1}^{t-1} \frac{1}{k} \leq 1 + \ln t$, gives the desired bound. ■ In particular, for $\psi(x) = \frac{1}{2} \|x\|_2^2$, the Bregman divergence is $D_\psi(x, y) = \frac{1}{2} \|x - y\|_2^2$, which is strongly convex with respect to the Euclidean norm with constant $\mu = 1$. Then, taking $\alpha = \frac{1}{2}$ yields the same bound obtained by Shamir and Zhang [23], Theorem 2.

V. NUMERICAL EXAMPLES




Fig. 2. Example network with a weakly convex Rosenthal potential.

To illustrate the convergence results of Section IV, we simulate the stochastic distributed routing on an example network given in Figure 2. The resulting Rosenthal potential function f is not strongly convex. The path losses are taken to be bounded by 1. We add, to each path, a centered




Fig. 3. Potential values $f(x^{(\tau)}) - f^*$, averaged across 100 simulations (with a 1 standard deviation in dotted lines), for different choices of learning rate sequences. The dashed lines show the $\mathcal{O}(t^{-\min_m \min(\alpha_m, 1-\alpha_m)} \log t)$ rate predicted by Theorem 3.

Gaussian noise with standard deviation σ , which results in stochastic loss vectors with a bounded second moment, $\mathbb{E}[\|\hat{\ell}\|_\infty^2] \leq 1 + \sigma^2$. For the population dynamics, we implement the DSMD given by Algorithm 1, with a Bregman divergence generated by a regularized entropy function $\psi_\epsilon(x) \triangleq \sum_i (x_i + \epsilon) \ln(x_i + \epsilon)$, for some parameter $\epsilon > 0$. The corresponding Bregman divergence is $D_{\psi_\epsilon}(x, y) = \sum_i (x_i + \epsilon) \ln \frac{x_i + \epsilon}{y_i + \epsilon}$. This choice of function ψ ensures that the Bregman divergence remains bounded on the simplex, and that the update step $x_m^{(t+1)} = \arg \min_{x \in \mathcal{X}_m} \langle \hat{\ell}_m^{(t)}, x \rangle + \frac{1}{\eta_t^m} D_{\psi_\epsilon}(x_m, x_m^{(t)})$ can be solved efficiently: if \mathcal{X}_m is a scaled d -dimensional simplex, then the update can be solved either in $\mathcal{O}(d \ln d)$ time using a deterministic algorithm or in $\mathcal{O}(d)$ expected time using a randomized pivot algorithm, see Krichene et al. [14].

The results of the simulations are given in Figure 3, in which we report, in log-log scale, the potential values averaged over 100 realizations, for two different choices of (heterogeneous) learning rates. The empirical convergence rates observed in simulation are consistent with those predicted by Theorem 3.

In addition to the convergence of $\mathbb{E}[f(x^{(t)})]$, Theorem 2 provides a bound on $\mathbb{E}[D_{\psi_\epsilon}(x^*, x^{(t)})]$ if the potential f is strongly convex and the populations use the same sequence of learning rates. To illustrate this result, we simulate the stochastic routing game on a second network, given in Figure 4, for which the Rosenthal potential is strongly convex. We implement the DSMD dynamics with homogeneous learning rates $\eta_t = \frac{\theta}{t}$. We report the sequence $\mathbb{E}[D_{\psi_\epsilon}(x^*, x^{(t)})]$ in Figure 5. The empirical convergence rate is consistent with the $\mathcal{O}(1/t)$ bound predicted by Theorem 2.




Fig. 4. Example network with a strongly convex Rosenthal potential.




Fig. 5. Bregman divergence to equilibrium, averaged across 100 simulations. The dashed line shows the $\mathcal{O}(t^{-1})$ convergence rate predicted by Theorem 2.

VI. CONCLUSION

We study a model of stochastic routing, in which the observed loss vectors are noisy. We propose a distributed model of learning, based on the stochastic mirror descent method.

We analyzed the convergence of the DSMD dynamics for non-smooth convex potentials. In the homogeneous case, i.e. assuming all populations use learning rates with the same decay, we prove almost sure convergence of $x^{(t)}$ to the set of minimizers \mathcal{X}^* , and convergence of the variance under an additional strong convexity assumption.

Under the more general, heterogeneous case, we prove that with learning rates $\eta_t^m = \theta_m t^{-\alpha_m}$, $\alpha_m \in (0, 1)$, the sequence of expected potentials $\mathbb{E}[f(x^{(t)})]$ converges to f^* at a $\mathcal{O}(t^{-\min(\alpha_{\min}, 1-\alpha_{\max})} \ln t)$ rate. This result holds for general convex functions (non-smooth, non strongly convex), and proves convergence of the actual sequence, as opposed to the weaker result of convergence for the sequence of averages $\bar{x}^{(t)} = \frac{\sum_{\tau=1}^t x^{(\tau)}}{t}$, as discussed in Section IV-A.

These results provide a general model for population dynamics for distributed routing, under which convergence is robust to stochastic perturbations: convergence of the last iterate is guaranteed (with a bound on the convergence rate in expectation), even when different populations use different Bregman divergences and different learning rates.

Beyond convergence of the DSMD learning model, a related problem is to study its privacy guarantees: for example, if populations have parameters which they would like to keep private, such as the total mass R_m , one can show that the presence of noise in the observations provides differential privacy guarantees. We are currently investigating this problem, in order to quantify these guarantees.

REFERENCES

- [1] Arindam Banerjee, Srujana Merugu, Inderjit S. Dhillon, and Joydeep Ghosh. Clustering with bregman divergences. *J. Mach. Learn. Res.*, 6:1705–1749, December 2005. ISSN 1532-4435.
- [2] Amir Beck and Marc Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Oper. Res. Lett.*, 31(3):167–175, May 2003.
- [3] Martin J Beckmann, Charles B McGuire, and Christopher B Winsten. Studies in the economics of transportation. 1955.
- [4] Avrim Blum, Eyal Even-Dar, and Katrina Ligett. Routing without regret: on convergence to nash equilibria of regret-minimizing algorithms in routing games. In *Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*, PODC '06, pages 45–52, New York, NY, USA, 2006. ACM.
- [5] Léon Bottou. Online algorithms and stochastic approximations. 1998. revised, oct 2012.
- [6] Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122, 2012.
- [7] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.
- [8] Simon Fischer and Berthold Vöcking. On the evolution of selfish routing. In *Algorithms–ESA 2004*, pages 323–334. Springer, 2004.
- [9] Elad Hazan, Amit Agarwal, and Satyen Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.
- [10] Josef Hofbauer and William H. Sandholm. Stable games and their dynamics. *Journal of Economic Theory*, 144(4):1665 – 1693.e4, 2009. ISSN 0022-0531.
- [11] Anatoli Juditsky, Arkadi Nemirovski, and Claire Tauvel. Solving variational inequalities with stochastic mirror-prox algorithm. *Stoch. Syst.*, 1(1):17–58, 2011.
- [12] Robert Kleinberg, Georgios Piliouras, and Eva Tardos. Multiplicative updates outperform generic no-regret learning in congestion games. In *Proceedings of the 41st annual ACM symposium on Theory of computing*, pages 533–542. ACM, 2009.
- [13] Walid Krichene, Benjamin Drighès, and Alexandre Bayen. Learning nash equilibria in congestion games. *SIAM Journal on Control and Optimization (SICON)*, 2015.
- [14] Walid Krichene, Syrine Krichene, and Alexandre Bayen. Efficient bregman projections onto the simplex. In *54th IEEE Conference on Decision and Control*, 2015.
- [15] Walid Krichene, Syrine Krichene, and Alexandre Bayen. Convergence of mirror descent dynamics in the routing game. In *European Control Conference (ECC)*, 2015.
- [16] P. R. Kumar and Pravin Varaiya. *Stochastic Systems: Estimation, Identification and Adaptive Control*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1986. ISBN 0-13-846684-X.
- [17] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.
- [18] A. S. Nemirovsky and D. B. Yudin. *Problem complexity and method efficiency in optimization*. Wiley-Interscience series in discrete mathematics. Wiley, 1983.
- [19] Asuman Ozdaglar and R Srikant. Incentives and pricing in communication networks. *Algorithmic Game Theory*, pages 571–591, 2007.
- [20] H. Robbins and D. Siegmund. A convergence theorem for non negative almost supermartingales and some applications. *Optimizing Methods in Statistics*, 1971.
- [21] Robert W Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.
- [22] Tim Roughgarden. Routing games. In *Algorithmic game theory*, chapter 18, pages 461–486. Cambridge University Press, 2007.
- [23] Ohad Shamir and Tong Zhang. Stochastic gradient descent for non-smooth optimization: Convergence results and optimal averaging schemes. In *ICML*, pages 71–79, 2013.
- [24] John Glen Wardrop. Some theoretical aspects of road traffic research. In *ICE Proceedings: Engineering Divisions*, volume 1, pages 325–362. Thomas Telford, 1952.
- [25] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *ICML*, pages 928–936, 2003.