



On stochastic mirror descent with interacting particles: Convergence properties and variance reduction

A. Borovykh^{a,*}, N. Kantas^b, P. Parpas^a, G.A. Pavliotis^b

^a Department of Computing, Imperial College London, United Kingdom

^b Department of Mathematics, Imperial College London, United Kingdom

ARTICLE INFO

Article history:

Received 1 July 2020

Received in revised form 27 October 2020

Accepted 3 January 2021

Available online 13 January 2021

Keywords:

Mirror descent

Interacting agents

Variance reduction

ABSTRACT

An open problem in optimization with noisy information is the computation of an exact minimizer that is independent of the amount of noise. A standard practice in stochastic approximation algorithms is to use a decreasing step-size. This however leads to a slower convergence. A second alternative is to use a fixed step-size and run independent replicas of the algorithm and average these. A third option is to run replicas of the algorithm and allow them to interact. It is unclear which of these options works best. To address this question, we reduce the problem of the computation of an exact minimizer with noisy gradient information to the study of stochastic mirror descent with interacting particles. We study the convergence of stochastic mirror descent and make explicit the tradeoffs between communication and variance reduction. We provide theoretical and numerical evidence to suggest that interaction helps to improve convergence and reduce the variance of the estimate.

© 2021 Elsevier B.V. All rights reserved.

1. Introduction

Optimization models that arise in artificial intelligence and statistical learning applications often include noisy estimates of the function and its gradient. This is the case when for example the gradient is computed over a subset (or mini-batch) of the data. In such a situation it is known that the optimization algorithm will converge to a neighborhood of the minimizer [1]. The size of the neighborhood depends on the amount of noise. In addition, for constrained optimization problems noise can violate the constraints making the situation even more complex.

In various applications it can be beneficial to be able to control the fluctuations around the true minimum. The conventional way to control the error is to decrease the step size. Theoretical analysis suggests step sizes which are slow in practice, e.g. $O(1/t)$ in [2]. An alternative is to use a vanishing noise variance [3] or heuristics such as to increase the batch size over time (see e.g. [4]); this however increases the computational costs and is difficult to tune. Another option is to run independent replicas of the algorithm. We will refer to each of these runs as a particle. The question we address in this paper is whether it is beneficial to allow these particles to interact with each other. We study this question using the general framework of Stochastic Mirror

Descent (SMD) (see [5]). SMD can be used to solve constrained and unconstrained problems, and is known to be an optimal algorithm for certain classes of optimization problems [6].

In this paper we will consider generic convex optimization problems of the form,

$$\min_{x \in \mathcal{X}} \{f(x)\},$$

where $\mathcal{X} \subset \mathbb{R}^d$ is a closed convex set that describes the constraints. We are interested in investigating the performance and properties when the minimizer x^* is estimated using the following Itô stochastic differential equation (SDE)

$$dz_t^i = -\nabla f(x_t^i) dt + \sum_{j=1}^N A_{ij} (z_t^j - z_t^i) dt + \sigma dB_t^i,$$

$$x_t^i = \nabla \Phi^*(z_t^i), \quad i = 1, \dots, N,$$

where each particle is driven by independent Brownian motions B_t^i and Φ is the mirror map used in Mirror Descent (MD); we will present more details in Section 2. The interesting feature here is that particles interact through the matrix $A = \{A_{ij}\}_{i,j=1}^N$, which is a $N \times N$ doubly-stochastic matrix representing the interaction weights. This interaction will attract particles towards each other. The matrix A represents an interaction graph which imposes communication constraints on the agents: each particle i can communicate directly only with its immediate neighbors, i.e. $j \in \{1, \dots, N\}$ for whom $A_{ij} \neq 0$. In the absence of interactions (i.e. when $A = 0$) the dynamics would correspond to independent replicas of SMD.

* Corresponding author.

E-mail addresses: a.borovykh@imperial.ac.uk (A. Borovykh), n.kantas@imperial.ac.uk (N. Kantas), panos.parpas@imperial.ac.uk (P. Parpas), g.pavliotis@imperial.ac.uk (G.A. Pavliotis).

The aim of this paper is to demonstrate the advantages of using the interacting particle system, in comparison to independent copies of SMD. We will show that interaction reduces the variance of the estimation of x^* and can improve the convergence properties including the convergence rate to the stationary distribution. In particular, we show that for a (strongly) convex objective the distance to the optimum is bounded by a term related to the standard optimization error in which the noise variance is reduced by a factor of N and an interaction term measuring the deviation of each particle from the system average. This latter term is bounded, so that under certain assumptions consensus and consequently convergence to the optimum is achieved. Using logarithmic Sobolev inequalities we can furthermore improve on this rate and show that the convergence to a stationary distribution can be achieved at an exponential rate. In addition, we believe that the ability to impose communication constraints through the interaction A will be beneficial also from a practical point of view as it will decrease communication costs in a parallel implementation. In this paper we will work with the continuous time formulation for SMD and its proposed interacting version. The main reason behind this is that the analysis for continuous time dynamics provides a clear and complete picture for the benefit of interaction. In our numerical experiments we used a simple forward Euler discretization scheme and we expect that the results can only improve with more sophisticated schemes. A detailed error analysis of the discretization scheme is beyond the scope of this paper.

1.1. Related literature

It is a well-known problem that stochastic optimization algorithms converge to a neighborhood of the (local) optimal solution. The size of this neighborhood is proportional to the noise variance (or the second moment of the sample gradient). Traditionally, common strategies for mitigating this include using a decreasing stepsize or attaining vanishing noise variance by increasing the batch size. Both of these come at additional computational costs. In the context of SMD, the authors of [2] study the convergence of SMD under the assumption that either the variance of the noise decreases over time or the step size is reduced exponentially slow. Various strategies have been also proposed to mitigate the effect of this noise variance to enable convergence closer to the optimum. The effect of sampling strategies such as importance sampling [7] can decrease the effect of the noise variance. Various variance reduction methods have been proposed, e.g. [8–10]. The work of [11] and [12] studied an accelerated version of SMD, however the distance to the optimum remains a function of the noise.

In this work we propose to use interacting particles to achieve a similar variance reduction. The work most closely related to ours is that of [13] in continuous time and [14] in discrete time. The authors in [13] study continuous and interacting SMD however only attained convergence rates for linear dynamics in the mirror domain and did not present any numerical results. The work of [14] covers the deterministic and noisy gradient setting in discrete time in a distributed setting; our work is in continuous time with all particles optimizing the same objective and uses different assumptions on the noise dynamics.

Our work also has parallels with the vast distributed optimization literature from which we list some indicative recent references: [15–18]. We remark that our objective is different from that in distributed optimization. Our goal here is to study how interaction can improve the convergence properties of SMD and analyze the variance reduction effects, while in distributed optimization the task is to optimize efficiently an objective function that is distributed across different nodes. For example a

distributed version of the interacting SDE above would require using a different drift $-\nabla f_i(x_i^t)$ in the dynamics. There are also parallels with the work on consensus and synchronization, see e.g. [19–22] and [23]. Our work generalizes some of these results to the SMD setting and combine it with results on convergence to the minimum. There is also parallel work using stochastic gradient descent (SGD) for constrained sampling problems considered in [24–26]. While there are many similarities, the goals of sampling and optimization are different and different dynamics are used in each case [27]; specifically, in optimization one wants to converge to the optimizer, while in sampling the objective is to converge to the correct invariant measure.

1.2. Contributions and organization

Our main focus is on the convergence properties of stochastic mirror descent with interacting particles (ISMD). We derive regret bounds for the proximity of $f(x_i^t)$ to the optimal value using Lyapunov-based arguments. The application of Lyapunov techniques to optimization problems is an established approach. Our results have similarities with [13] and [2] bar a number of differences. We use different Lyapunov functions and analyze a general convex cost function. We show that in the case of ISMD there is a tradeoff between communication cost, i.e. how many particles interact, and variance reduction. In particular, with a fixed learning rate and non-vanishing noise variance, interaction between the particles can reduce the size of the neighborhood around the optimum to which the algorithm converges. Furthermore, we show that the particles converge to an area around the optimizer at an exponential rate using log-Sobolev inequalities and Bakry–Emery theory (see [28] and [29]). Although such results are standard for the analysis of interacting SDEs with convex potentials (e.g. [30,31]), they have not been considered previously in the context of SMD. In addition, in our setting *the particles are not restricted to all interact simultaneously*, so our results deviate from the mean field type analysis like in [30,31].

Our contributions can be summarized as follows:

- We propose interaction between N particles as a way of controlling the convergence of a stochastic optimization algorithm – and in particular the distance to the minimum. In the presence of noise in the gradients interacting particles are an effective alternative to vanishing learning rates or noise variances.
- We establish that interaction leads to variance reduction. In particular under strong convexity and smoothness assumptions (to be specified later) we show that there exist positive constants K_1 and K_2 such that,

$$\int_0^T e^{K_1\kappa(t-T)} \mathbb{E}[f(x_i^t) - f(x^*)] dt \leq K_2 \left(\frac{1}{2} e^{-K_1\kappa T} + \frac{\sigma^2}{2N} + \frac{N-1}{N} \frac{d\sigma^2}{\underline{\lambda} + \kappa} \right),$$

where $\underline{\lambda}$ is the first nonzero eigenvalue of the graph Laplacian and κ is the strong convexity constant related to the mirrored objective. A similar average regret bound can be established for the convex case (and $\kappa = 0$). This result, together with the technical assumptions, proofs and extensions will be presented in [Propositions 15–16](#) and [17–18](#). The distance to optimality is thus bounded by three terms: (1) a term which decays exponentially with time, (2) a noise term which decreases as the number of particles increases (3) a term which arises from the distance between the particle values which is smaller for a connected graph, a stronger interaction and a larger strong convexity constant.

- We provide explicit rates of convergence and concentration inequalities for the particle system using log Sobolev inequalities. This will be presented in [Section 4.2](#). These results establish exponential rate of convergence in time for

the law of the particle system, so that at equilibrium the samples z_t^i will oscillate around the optimum z^* . The strong convexity constant plays a similar role as in the previous result, with a higher κ resulting in a smaller distance to the optimum.

- Finally, we show the benefits of using interaction in convergence speed and variance reduction through different numerical experiments. In a distributed optimization setup we manage to achieve comparable performance and convergence to the full batch gradient descent, when using a “mini-batch” approach that uses N particles with a mini-batch size that is $1/N$ times the total number of data-points or summands in f (see Fig. 8 in Section 5 for more details).

The organization of this paper is as follows: Section 2 presents background material on MD and SMD together with some basic convergence results. Most of the material in this section is known, but we use this opportunity to set our framework and to relate the continuous time formulation with common discrete time implementations. We then propose ISMD in Section 3 and provide a detailed convergence analysis on how the particles approach the minimum as time increases. In Section 4 we apply the Bakry–Emery theory for the law of the corresponding particle system. In Section 5 we demonstrate the performance of ISMD in practice and demonstrate the effect of interaction on decreasing the variance in a variety of examples. Finally, we provide some concluding remarks in Section 6.

1.2.1. Notation

We will use the following notations: we denote by I_d the d -dimensional identity matrix and $\mathbf{1}_d$ the d -dimensional vector of ones. Let $\text{Diag}(a)$ with $a \in \mathbb{R}^d$ refer to a matrix with diagonal elements $[a_1, \dots, a_d]$. The Kronecker product is denoted by \otimes . Given an arbitrary norm $\|\cdot\|$ on \mathbb{R}^d , we will define $B_{\|\cdot\|} := \{v \in \mathbb{R}^d : \|v\| \leq 1\}$. The dual norm $\|\cdot\|_*$ is defined as $\|z\|_* := \sup\{z^T v : v \in B_{\|\cdot\|}\}$. Denote with $\|A\|_2$ the spectral norm if A is a matrix. Similarly, let $\|A\|_F$ denote the Frobenius norm given by $\|A\|_F^2 = \sum_{i=1}^N \lambda_i^2$, where λ_i denote the eigenvalues of A . Assume that the dual norm is compatible with the spectral norm, i.e. $\|Az\|_* \leq \|A\|_2 \|z\|_*$. Unless specified otherwise K, K' etc. denote generic constants whose value may change according to context.

We will define a μ strongly convex function f (w.r.t. a norm $\|\cdot\|$) if there exists $m > 0$ such that for all x, y :

$$f(y) \geq f(x) + \nabla f(x)^T (y - x) + \frac{m}{2} \|y - x\|^2.$$

The case $m = 0$ corresponds to the convex case. Note that equivalently strong convexity can be defined using

$$(x - y)^T (\nabla f(x) - \nabla f(y)) \geq m \|x - y\|^2.$$

Furthermore, for a Lipschitz function we let

$$\|f\|_{Lip} = \sup_{x \neq y} \frac{|f(x) - f(y)|}{\|x - y\|} < \infty.$$

When looking at the Hessian of a convex function f we will denote with $\text{Hess}(f) \geq \lambda I_d$ that $\text{Hess}(f) - \lambda I_d$ is positive semidefinite with $\lambda \geq 0$.

For random variables X, Y we say $X \stackrel{d}{=} Y$ to denote equality in distribution. For a measure π and measurable function f we use $\pi(f) = \mathbb{E}_\mu[f] = \int f(x)\pi(dx)$. For any measurable space $(\mathcal{Z}, \mathcal{B}(\mathcal{Z}))$ we use $\mathcal{P}(\mathcal{Z})$ to denote the space of all probability measures on \mathcal{Z} and $\mathcal{P}_2(\mathcal{Z})$ the one for finite second moments. A probability measure ν satisfies a Log-Sobolev inequality with constant C is for any smooth function f we have,

$$\text{Ent}_\nu(f^2) \leq C \nu(|\nabla f|^2), \tag{1}$$

where the entropy is defined as,

$$\text{Ent}_\nu(f^2) = \nu(f^2 \log f^2) - \nu(f^2) \log(\nu(f^2)).$$

For $\mu, \nu \in \mathcal{P}(\mathcal{Z})$ we will denote the Kullback–Leibler divergence or relative entropy as,

$$H(\mu|\nu) = \begin{cases} \int \log \frac{d\mu}{d\nu} d\mu & \text{if } \mu \ll \nu, \\ +\infty & \text{otherwise.} \end{cases}$$

The 2-Wasserstein distance is defined as

$$W_2(\mu, \nu) = \left(\inf_{\Pi(\mu, \nu)} \left(\frac{1}{2} \int |z - z'|^2 \Pi(dz, dz') \right) \right)^{1/2},$$

where $\Pi \in \mathcal{P}_2(\mathcal{Z} \times \mathcal{Z})$ is the coupling of μ, ν with $\Pi(\mathcal{Z}, \cdot) = \nu$ and $\Pi(\cdot, \mathcal{Z}) = \mu$. We will assume that the optimization algorithm is started from a fixed deterministic point, and let $\mathbb{E}[\cdot]$ refer to the expected value conditional on the initial value.

2. Background

2.1. Preliminaries

We are interested in computing $x^* \in \arg \min_{x \in \mathcal{X}} \{f(x)\}$ under the assumption of smoothness and convexity for f and \mathcal{X} . Throughout the paper we will assume continuity and smoothness of f :

Assumption 1. We let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be L -Lipschitz continuous with L -Lipschitz continuous gradients. $\mathcal{X} \subset \mathbb{R}^d$ is a closed convex set.

We allow f and ∇f to be Lipschitz w.r.t arbitrary norms. As all norms are equivalent up to a proportionality constant, a natural choice without loss of generality is the Euclidean norm, which is also self dual, i.e. $\|\cdot\|_* = \|\cdot\|_2$ (see [32, Lemma 2.1.2]). Despite this in the remainder we will distinguish between the chosen norm and its dual mainly in order to emphasize the difference when operating in the primal or mirror domain.

We will use a *mirror map* $\Phi : \mathcal{X} \rightarrow \mathbb{R}^d$ to convert the constrained optimization problem to an unconstrained one and adopt the following standard assumptions (e.g. [6, Assumption 9.3]):

Assumption 2. $\Phi : \mathcal{X} \rightarrow \mathbb{R}^d$ is μ strongly convex and continuously differentiable.

The mirror map and its conjugate will be used to pass between the constrained and unconstrained space. Define

$$\Phi^*(z) := \max_{x \in \mathcal{X}} (x^T z - \Phi(x))$$

to be the Legendre–Fenchel convex conjugate, so that

$$\nabla \Phi^*(z) := \arg \max_{x \in \mathcal{X}} (z^T x - \Phi(x)).$$

Note that one has

$$\nabla \Phi \circ \nabla \Phi^*(z) = z.$$

Note when Φ is strongly convex with constant μ , the conjugate correspondence theorem (see Section 5.3.1 in [6]) gives that the gradient of Φ^* is Lipschitz continuous:

$$\|\nabla \Phi^*(z) - \nabla \Phi^*(z')\| \leq \mu^{-1} \|z - z'\|_*.$$

We furthermore make the additional assumption that the conjugate of the mirror map maps \mathbb{R}^d to \mathcal{X} :

Assumption 3. $\nabla \Phi^*(\mathbb{R}^d) = \mathcal{X}$ and $\|\Delta \Phi^*\|_\infty < \infty$.

This simplifies the analysis as $\nabla\Phi^*$ maps directly to \mathcal{X} and this avoids the need for projections. Extending our results without using [Assumption 3](#) is possible by following a route similar to [2].

We will often use the notation z^* for $z^* \in \mathbb{R}^n$ such that $x^* = \nabla\Phi^*(z^*) \in \arg \min_{x \in \mathcal{X}} f(x)$.

The *Bregman divergence* is defined as,

$$D_\Phi(x, y) = \Phi(x) - \Phi(y) - \nabla\Phi(y)^T(x - y).$$

The Bregman divergence is meant to quantify how far a point x is from y and Φ can be thought as a distance generating function that adapts to the geometry or structure of \mathcal{X} . Two well-known examples of the Bregman divergence are the Euclidean distance with $\Phi(x) = \frac{1}{2}\|x\|_2^2$, and the simplex constraint $\Phi(x) = \sum_{i=1}^d x(i) \log x(i) - x(i)$.

An important property of using a mirror map with the Bregman divergence is that

$$D_{\Phi^*}(z_t, z^*) = D_\Phi(x_t^*, x_t^*),$$

for some values z_t and z^* (x_t and x^*) in unconstrained (constrained) space. In addition, the Φ -diameter of \mathcal{X} is defined as

$$D_{\Phi, \mathcal{X}} := \sup_{x \in \mathcal{X}, x' \in \mathcal{X}^o} \sqrt{2D_\Phi(x, x')}, \quad (2)$$

and is meant to quantify the size of the constrained space \mathcal{X} when $D_\Phi(x, y)$ is used as a measure of distance.

We will make the following assumption on the Lipschitz-continuity on the Bregman divergence.

Assumption 4. The Bregman divergence satisfies a Lipschitz condition of the form,

$$|D_{\Phi^*}(x, z) - D_{\Phi^*}(y, z)| \leq K\|x - y\|_*.$$

for some constant K .

This is a common technical condition (e.g. see [16]), which is automatically satisfied when the function Φ^* is Lipschitz. Here it is used only for the strongly convex case in [Proposition 16](#).

We will assume convexity of the objective function f .

Assumption 5. $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is convex.

This further implies that the gradient of f is uniformly bounded on \mathcal{X} by the constant L , i.e. $\|\nabla f(x)\|_* \leq L$. Alternatively, at times we will work under a strong convexity assumption.

Assumption 6. The function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is a μ -strongly convex function with respect to Φ . This means

$$D_f(x, y) \geq \mu D_\Phi(x, y).$$

We further let the anti-derivative \mathcal{V} be defined as $\nabla\mathcal{V} = \nabla f \circ \nabla\Phi^*$. We note that knowledge of \mathcal{V} here is mainly used to improve exposition, convey intuition and in practice we work only with $\nabla\mathcal{V}$. The main use of \mathcal{V} itself will be to describe invariant distributions of the mirror descent SDEs as defined in [Section 2.2](#). For the analysis, we note that [Assumptions 1 and 2](#) imply $\nabla\mathcal{V}$ is Lipschitz and we will use the following convexity assumptions in [Section 3.3](#) to bound the fluctuations of the particle system and in [Section 4](#) to establish rates of convergence. We will distinguish between the convex and strongly convex cases for \mathcal{V} .

Assumption 7. The function \mathcal{V} is convex.

Assumption 8. The function \mathcal{V} is κ -strongly convex with $\kappa > 0$.

2.2. Mirror descent

We begin the presentation of mirror descent (MD) with some preliminaries on gradient descent. We provide a brief account of the main approaches in discrete time and later discuss the continuous time setting, which is of interest in this paper.

2.2.1. Projected gradient descent

One way of finding the solution of a constrained optimization problem is through projected gradient descent (GD),

$$x_{k+1} = \Pi_{\mathcal{X}}(x_k - \eta\nabla f(x_k)), \quad (3)$$

where k here denotes discrete time, η is the learning rate and we define the projection using the Euclidean norm, i.e. $\Pi_{\mathcal{X}}(y) = \arg \min_{x \in \mathcal{X}} \|y - x\|_2^2$. The main drawback of this method is its slow convergence: when d , the dimension of x , increases the convergence can slow down. This is due to the fact that this scheme is tied to the Euclidean geometry of \mathbb{R}^d through the projection operator $\Pi_{\mathcal{X}}$. Consider a setting in which $\|\nabla f(x)\|_\infty \leq 1$. This implies $\|\nabla f(x)\|_2 < \sqrt{d}$, so projected gradient descent will converge at a rate \sqrt{d}/k (see e.g. [Example 9.17](#) from [6]).

2.2.2. Mirror descent in discrete and continuous time

A generalization of the projected gradient descent method is mirror descent, which is given by:

$$\begin{aligned} \nabla\Phi(y_{k+1}) &= \nabla\Phi(x_k) - \eta\nabla f(x_k), \\ x_{k+1} &= \Pi_{\mathcal{X}}^\Phi(y_{k+1}) := \arg \min_{x \in \mathcal{X}} D_\Phi(x, y_{k+1}). \end{aligned}$$

The last step can be rewritten as,

$$x_{k+1} = \arg \min_{x \in \mathcal{X}} \{\Phi(x) - \nabla\Phi(y_{k+1})^T x\}.$$

The current point x_k is thus mapped into the dual space $\nabla\Phi(x_k)$; this is then updated to $\nabla\Phi(y_{k+1})$ by stepping in the direction of the negative gradient $-\eta\nabla f(x_k)$ and then mapped back to x_{k+1} . Observe that for $\Phi(x) = \frac{1}{2}\|x\|_2^2$, we have $\nabla\Phi(x) = x$ and the Bregman divergence is simply given by $D_\Phi(x, y) = \frac{1}{2}\|x - y\|_2^2$, so that the algorithm is equivalent to projected gradient descent. Consider $\Phi(x) = \sum_{i=1}^n x(i) \log x(i)$ and let $\mathcal{X} := \Delta_d$, the d -dimensional simplex. Observe that $\nabla_i\Phi(x) = \log x(i) + 1$, and the projection of y with respect to this Bregman divergence onto the simplex Δ_d amounts to a scaling with the L^1 -norm. The benefit of MD is that it can adapt to the structure or geometry of \mathcal{X} and hence result in a much better dependence of the rate of convergence on the problem dimension d (see [Chapter 9](#) in [6]).

Using [Assumption 3](#) the continuous time version of mirror descent is given by,

$$dz_t = -\nabla f(x_t)dt, \quad x_t = \nabla\Phi^*(z_t), \quad (4)$$

where we have set $\eta = 1$ for simplicity.

2.2.3. Convergence of MD in continuous time

We will proceed with presenting two well-known convergence results for deterministic mirror descent and the dynamics of (4); see [5,33] for more details. These results are based on using Bregman divergence in a Lyapunov function V_t and are included to ease exposition and to facilitate the comparison to the stochastic case. Note that $D_{\Phi, \mathcal{X}}$ is defined in (2).

Lemma 9 (Convergence of MD for a Convex Objective). Assume that [Assumptions 1–3](#) and [5](#) hold. We then have the following result,

$$\frac{1}{T} \int_0^T (f(x_t) - f(x^*))dt \leq \frac{D_{\Phi, \mathcal{X}}^2}{2T}.$$

Proof. Consider,

$$V(z_t) := \Phi^*(z_t) - \Phi^*(z^*) - \nabla \Phi^*(z^*)^T(z_t - z^*).$$

Observe that $V(z^*) = 0$ and $V(z_t) \geq 0$ for $z^* \neq z_t$. Furthermore,

$$\begin{aligned} \frac{dV(z_t)}{dt} &= (\nabla \Phi^*(z_t))^T - \nabla \Phi^*(z^*)^T \frac{dz_t}{dt} \\ &= -(x_t - x^*)^T \nabla f(x_t) \leq f(x^*) - f(x_t) \leq 0, \end{aligned}$$

where in the second equality we used $x_t = \nabla \Phi^*(z_t)$ and in last inequality we have used the convexity of f . Integrating $dV(z_t)/dt$ from $t = 0$ to $t = T$, using the convexity of f and the relationship $D_{\Phi^*}(z_t, z^*) = D_{\Phi}(x_t, x^*)$ so that $V(z_0) \leq \frac{1}{2}D_{\Phi, \mathcal{X}}^2$ we obtain,

$$\int_0^T (f(x_t) - f(x^*))dt \leq \frac{1}{2}D_{\Phi, \mathcal{X}}^2,$$

and the claim follows. \square

Observe that furthermore since $\inf_{0 \leq t \leq T} f(x_t) \leq \frac{1}{T} \int_0^T f(x_t)dt$, we can obtain a result for the best value of x_t ,

$$\inf_{0 \leq t \leq T} f(x_t) - f(x^*) \leq \frac{D_{\Phi, \mathcal{X}}^2}{2T}.$$

Observe that for projected gradient descent we have $\Phi(x) = \frac{1}{2}\|x\|_2^2$. In this case, $D_{\|\cdot\|_2, \mathcal{X}} = \sup_{x, x' \in \mathcal{X}} \|x - x'\|_2$. From Lemma 9 we observe that if $D_{\Phi, \mathcal{X}}^2$ is smaller than $D_{\|\cdot\|_2, \mathcal{X}}$ then mirror descent algorithm could achieve a faster convergence rate than gradient descent due to the smaller constant.

In the strongly convex case the convergence speed can be increased to an exponential rate.

Lemma 10 (Convergence of MD for a Strongly Convex Objective). *Assume that Assumptions 1–3 and 6 hold. We then have the following result,*

$$D_{\Phi}(x^*, x_t) \leq e^{-\mu t} D_{\Phi}(x^*, x_0).$$

Proof. Since f is μ -strongly convex with respect to Φ , it holds,

$$D_f(x, y) \geq \mu D_{\Phi}(x, y). \tag{5}$$

Observe that,

$$\begin{aligned} \frac{dD_{\Phi^*}(z_t, z^*)}{dt} &= -(x_t - x^*)^T \nabla f(x_t) \\ &= -D_f(x^*, x_t) - f(x_t) + f(x^*) \\ &\leq -\mu D_{\Phi^*}(z_t, z^*) - f(x_t) + f(x^*) \\ &\leq -\mu D_{\Phi^*}(z_t, z^*), \end{aligned} \tag{6}$$

where in the second equality we have used the definition of the Bregman divergence, in the second-to-last inequality we have used (5) and the fact that $D_{\Phi}(x^*, x_t) = D_{\Phi^*}(z_t, z^*)$, and in the last inequality we have used the convexity of f , namely that $f(x^*) - f(x_t) < 0$. Let now,

$$V(t, z_t) := e^{\mu t} D_{\Phi^*}(z_t, z^*).$$

Then,

$$\begin{aligned} \frac{dV(t, z_t)}{dt} &= \mu e^{\mu t} D_{\Phi^*}(z_t, z^*) + e^{\mu t} \frac{dD_{\Phi^*}(z_t, z^*)}{dt} \\ &\leq \mu e^{\mu t} D_{\Phi^*}(z_t, z^*) - \mu e^{\mu t} D_{\Phi^*}(z_t, z^*) \leq 0. \end{aligned}$$

where in the second-to-last inequality we have used (6). Then, since $\frac{dV(t, z_t)}{dt} \leq 0$ it holds that $V(t, z_t) \leq V(0, z_0)$. Using this and $D_{\Phi}(x^*, x_t) = D_{\Phi^*}(z_t, z^*)$ we have,

$$\begin{aligned} V(t, z_t) &= e^{\mu t} D_{\Phi^*}(z_t, z^*) = e^{\mu t} D_{\Phi}(x^*, x_t) \\ &\leq V(0, z_0) = D_{\Phi}(x^*, x_0). \end{aligned}$$

This implies,

$$D_{\Phi}(x^*, x_t) \leq e^{-\mu t} D_{\Phi}(x^*, x_0). \quad \square$$

2.3. Stochastic mirror descent

Consider adding noise to (3):

$$x_{k+1} = \Pi_{\mathcal{X}}(x_k - \eta \nabla f(x_k) + \sqrt{\epsilon} \sigma \zeta_k), \quad \zeta_k \sim \mathcal{N}(0, I_d), \text{ i.i.d.}$$

In the limit of the time step $\epsilon \downarrow 0$ one recovers the following Itô stochastic integral equation [34]:

$$\begin{aligned} x_{t+1} &= x_0 - \eta \int_0^t \nabla f(x_s) ds + \sigma B_t + u_t, \\ u_t &= \int_0^t 1_{x_t \in \partial \mathcal{X}} du_s. \end{aligned} \tag{7}$$

Here B_t is a d -dimensional Brownian motion and a local time construction on the boundary of \mathcal{X} is used for u_t to ensure that $X_t \in \mathcal{X}$ almost surely; see [34,35] and the references therein for details. Understanding in detail the ergodicity properties or long time behavior of (7) is largely unexplored in the literature. We note that at the level of the evolution of densities, reflecting boundary conditions are commonly used for the forward or Fokker–Planck equation, e.g. [36, Section 4.2.2], but a detailed study of (7) along these lines goes beyond the scope of this article. In the remainder we will consider the mapping to be on the mirror domain.

The stochastic mirror descent (SMD) is given by the SDE,

$$dz_t = -\nabla f \circ \nabla \Phi^*(z_t) dt + \sigma dB_t, \tag{8}$$

where $x_t = \nabla \Phi^*(z_t)$ and as before in (4) we use $\eta = 1$. Note that when $\sigma = 0$ we obtain a deterministic variant of continuous time mirror descent. The long time behavior of such dynamics are very well understood. The law of z_t will converge to an invariant distribution with density proportional to $\exp(-\frac{2}{\sigma^2} \mathcal{V})$ where recall \mathcal{V} denotes the anti-derivative of $\nabla f \circ \nabla \Phi^*$. We will discuss this in more detail later in Section 4. The downside of this dynamics is that SMD cannot converge to the exact solution, i.e. the noise keeps the algorithm from fully converging exactly to x^* (see [2, 13]). The latter would require the equilibrium distribution being a Dirac/atomic measure, which in turn requires using a decreasing noise in time; see [2] for more details.

Another interesting remark is that the dynamics in the dual space is then equivalent to that of SGD with objective function \mathcal{V} . It is well-known that for a convex objective SGD methods converge in $O(1/t)$ to a neighborhood of the optimum the size of which is proportional to the noise variance [1]. A similar result can be shown to hold for SMD (see e.g. [2,13,37]).

Remark 11 (SGD and Additive Noise). We note that in this work we refer to SGD as the SDE equivalent of gradient descent with additive noise, i.e. we use a constant diffusion coefficient. This is in disagreement with recent works in the areas of machine learning whereby the noise in SGD arises only from sub-sampling a deterministic loss function. The latter can be modeled asymptotically as an SDE with multiplicative noise and state-dependent diffusion; see for instance [38–40].

2.3.1. Convergence of SMD

We proceed with presenting the analogs of Lemmas 9–10 for the SMD case.

Proposition 12 (Convergence of SMD for a Convex Objective). *Assume that Assumptions 1–3 and 5 hold. We then have the following result,*

$$\mathbb{E} \left[\frac{1}{T} \int_0^T (f(x_t) - f(x^*)) dt \right] \leq \frac{1}{2T} D_{\Phi, \mathcal{X}}^2 + \frac{1}{2} \sigma^2 \|\Delta \Phi^*\|_{\infty}.$$

Proof. We define the following Lyapunov function,

$$V(z_t) := \Phi^*(z_t) - \Phi^*(z^*) - \nabla \Phi^*(z^*)^T (z_t - z^*).$$

We then have the following results. By Itô's lemma we obtain,

$$dV(z_t) = \left(\nabla f(x_t)^T (x^* - x_t) + \frac{1}{2} \sigma^2 \Delta \Phi^*(z_t) \right) dt + \sigma (x_t - x^*)^T dB_t,$$

where for the last term we have $\sigma (x_t - x^*)^T dB_t \stackrel{d}{=} \sigma \|x_t - x^*\|_2 dB_t$. Integrating the expression we obtain,

$$V(z_T) = V(z_0) + \int_0^T (x^* - x_t)^T \nabla f(x_t) dt + \frac{1}{2} \sigma^2 \int_0^T \Delta \Phi^*(z_t) dt + \sigma \int_0^T \|x_t - x^*\|_2 dB_t.$$

Observe again that, $(x^* - x_t)^T \nabla f(x_t) \leq f(x^*) - f(x_t)$, by convexity of f . Furthermore, as before $V(z_0) \leq \frac{1}{2} D_{\Phi^*, \mathcal{X}}^2$. Lastly, we have $\frac{1}{2} \sigma^2 \int_0^T \Delta \Phi^*(z_t) dt \leq \frac{1}{2} \sigma^2 T \|\Delta \Phi^*\|_\infty$. Rearranging we then obtain,

$$\int_0^T (f(x_t) - f(x^*)) dt \leq \frac{1}{2} D_{\Phi^*, \mathcal{X}}^2 + \frac{1}{2} \sigma^2 T \|\Delta \Phi^*\|_\infty + \sigma \int_0^T \|x_t - x^*\|_2 dB_t.$$

Using the fact that the process $\{\|x_t - x^*\|_2\}_{t \geq 0}$ is adapted to the filtration of x_t so that

$$\mathbb{E} \left[\sigma \int_0^T \|x_t - x^*\|_2 dB_t \right] = 0,$$

we have,

$$\mathbb{E} \left[\frac{1}{T} \int_0^T (f(x_t) - f(x^*)) dt \right] \leq \frac{1}{2T} D_{\Phi^*, \mathcal{X}}^2 + \frac{1}{2} \sigma^2 \|\Delta \Phi^*\|_\infty \quad \square$$

A similar bound can be obtained for the time average $\mathbb{E} \left[f \left(\frac{1}{T} \int_0^T x_t dt \right) - f(x^*) \right]$ by using Jensen's inequality. Similar to the deterministic case, if $D_{\Phi^*, \mathcal{X}}^2$ is smaller than $D_{\|\cdot\|_2, \mathcal{X}}^2$ then MD achieves a faster convergence rate than gradient descent.

Proposition 13 (Convergence of SMD for a Strongly Convex Objective). Assume that Assumptions 1–3 and 6 hold. Then,

$$\mathbb{E} [D_\Phi(x^*, x_T)] \leq e^{-\mu T} \frac{1}{2} D_{\Phi^*, \mathcal{X}}^2 + \frac{1}{2\mu} \sigma^2 (1 - e^{-\mu T}) \|\Delta \Phi^*\|_\infty.$$

Proof. Let $V(t, z_t) = e^{\mu t} D_{\Phi^*}(z_t, z^*)$. We have through Itô's lemma,

$$dV(t, z_t) = \left(e^{\mu t} (x^* - x_t)^T \nabla f(x_t) + \mu e^{\mu t} D_{\Phi^*}(z_t, z^*) + \frac{1}{2} \sigma^2 e^{\mu t} \Delta \Phi^*(z_t) \right) dt + \sigma e^{\mu t} (x_t - x^*)^T dB_t.$$

Integrating the expression we obtain,

$$V(T, z_T) = V(0, z_0) + \int_0^T e^{\mu t} \left((x^* - x_t)^T \nabla f(x_t) + \mu D_{\Phi^*}(z_t, z^*) \right) dt + \int_0^T \frac{1}{2} \sigma^2 e^{\mu t} \Delta \Phi^*(z_t) dt + \int_0^T \sigma e^{\mu t} (x_t - x^*)^T dB_t.$$

Observe again that for a μ -strongly convex function $D_f(x^*, x_t) \geq \mu D_\Phi(x^*, x_t)$ and by the properties of the mirror map $D_\Phi(x^*, x_t) = D_{\Phi^*}(z_t, z^*)$. This implies that $\nabla f(x_t)^T (x^* - x_t) \leq -D_f(x^*, x_t) - f(x_t) + f(x^*) \leq -\mu D_\Phi(x^*, x_t) = -\mu D_{\Phi^*}(z_t, z^*)$. Furthermore, $\int_0^T \sigma e^{\mu t} (x_t - x^*)^T dB_t \stackrel{d}{=} \int_0^T \sigma e^{\mu t} \|x_t - x^*\|_2 dB_t$ and it holds that

$$\int_0^T \frac{1}{2} \sigma^2 e^{\mu t} \Delta \Phi^*(z_t) dt \leq \frac{1}{2} (e^{\mu T} - 1) \sigma^2 \|\Delta \Phi^*\|_\infty.$$

$$V(T, z_T) \leq V(0, z_0) + \frac{1}{2} \sigma^2 (e^{\mu T} - 1) \|\Delta \Phi^*\|_\infty + \int_0^T \sigma e^{\mu t} \|x_t - x^*\|_2 dB_t.$$

Taking expected values we obtain,

$$\mathbb{E} [D_\Phi(x^*, x_T)] \leq e^{-\mu T} D_{\Phi^*, \mathcal{X}}^2 + \frac{1}{2} \sigma^2 (1 - e^{-\mu T}) \|\Delta \Phi^*\|_\infty,$$

where we have used that the expected value of an Itô integral is zero and the relationship $D_{\Phi^*}(z_t, z^*) = D_\Phi(x^*, x_t)$. \square

Again in the strongly convex case the convergence speed can be increased. However, as expected for both the convex and strongly convex results, as $T \rightarrow \infty$, the gap to optimality is bounded from above by a quantity proportional to the noise variance σ^2 . The solution will not converge exactly to the minimum, but oscillate around it.

3. Interacting stochastic mirror descent

It is clear that when using a fixed learning rate and a constant noise variance, SGD does not converge to the optimum. In fact, the distance to the optimum is controlled by the amount of noise. Various strategies such as using a vanishing noise variance by increasing the batch size or variance-reduced SGD (see e.g. [8–10]) have been proposed. In this work, we consider an alternative approach for controlling the distance from the optimum; namely by using interactions between the particles [13].

We consider as an alternative to (8) and we consider the following interacting particle dynamics,

$$dz_t^i = -\nabla f \circ \nabla \Phi^*(z_t^i) dt + \sum_{j=1}^N A_{ij} (z_t^j - z_t^i) dt + \sigma dB_t^i, \quad i = 1, \dots, N, \quad (9)$$

where B_t^i are independent Brownian motions and $A = \{A_{ij}\}_{i,j=1}^N$ is an $N \times N$ doubly-stochastic matrix representing the interaction weights. As mentioned in Section 1 the matrix A represents an interaction graph which can also impose communication constraints on the particles. Each particle i will be influenced only from particles j for which $A_{ij} \neq 0$.

The discretized version of (9) is then,

$$x_k^i = \nabla \Phi^*(z_k^i), \quad (10)$$

$$z_{k+1}^i = z_k^i - \epsilon \nabla f(x_k^i) + \epsilon \sum_{j=1}^N A_{ij} (z_k^j - z_k^i) + \sqrt{\epsilon} \sigma \zeta_k^i, \quad (11)$$

where $\zeta_t \sim \mathcal{N}(0, I_d)$ and ϵ is the discretization parameter. In the absence of interactions, i.e. when $A_{ij} = 0$, we obtain a discretized version of SMD.

We remark here that the interacting mirror descent algorithm is equivalent to an algorithm in which the particle interaction is defined using a Bregman divergence instead of the L^2 distance. In particular for the time discretized version (10)–(11) we have the following result.

Lemma 14 (Interacting Mirror Descent as Bregman Distance Interaction in Primal Space). Let $\sigma = 0$ and $\epsilon = 1$, then interacting mirror descent in (10)–(11) can be rewritten as:

$$x_{k+1}^i = \arg \min_{x \in \mathcal{X}} \{ \nabla f(x_k^i)^T x + \sum_{j=1}^N A_{ij} D_\Phi(x, x_k^j) \}.$$

Proof. We have,

$$\begin{aligned} x_{k+1}^i &= \nabla \Phi^*(z_{k+1}^i) \\ &= \arg \min_{x \in \mathcal{X}} \{D_\Phi(x, \nabla \Phi^*(z_{k+1}^i))\} \\ &= \arg \min_{x \in \mathcal{X}} \{\Phi(x) - \Phi(x_{k+1}^i) - \nabla \Phi \circ \nabla \Phi^*(z_{k+1}^i)^T(x - \nabla \Phi^*(z_{k+1}^i))\} \\ &= \arg \min_{x \in \mathcal{X}} \{\Phi(x) - \Phi(x_{k+1}^i) - (z_{k+1}^i)^T x\}, \end{aligned}$$

where in the last equality we have used $\nabla \Phi(\nabla \Phi^*(x)) = x$. Continuing,

$$\begin{aligned} x_{k+1}^i &= \arg \min_{x \in \mathcal{X}} \{\Phi(x) - \Phi(x_{k+1}^i) + \nabla f(x_k^i)^T x - \sum_{j=1}^N A_{ij}(z_k^j)^T x\} \\ &= \arg \min_{x \in \mathcal{X}} \{\nabla f(x_k^i)^T x + \sum_{j=1}^N A_{ij} (\Phi(x) - \Phi(x_{k+1}^i) - (z_k^j)^T x)\}, \end{aligned}$$

where in the first equality we have used the evolution of z_{t+1}^i (note: the z_t^i terms cancel out due to the double stochasticity of A_{ij}). Then,

$$\begin{aligned} x_{k+1}^i &= \arg \min_{x \in \mathcal{X}} \{\nabla f(x_k^i)^T x \\ &\quad + \sum_{j=1}^N A_{ij} (\Phi(x) - \Phi(x_k^j) - (\nabla \Phi(x_k^j))^T(x - x_k^j))\} \\ &= \arg \min_{x \in \mathcal{X}} \{\nabla f(x_k^i)^T x + \sum_{j=1}^N A_{ij} D_\Phi(x, x_k^j)\}, \end{aligned}$$

where in the first equality we have used $z_k^j = \nabla \Phi(x_k^j)$, and – due to the minimum being taken over x – replaced $\Phi(x_{k+1}^i)$ with $\Phi(x_k^j)$ and added the term $\nabla \Phi(x_k^j)^T x_k^j$ and in the last equality we have used the definition of the Bregman divergence. \square

Unlike the standard L_2 consensus algorithm where $x_{t+1}^i = \arg \min_{x \in \mathcal{X}} \nabla f(x_i^t)^T x + \sum_{j=1}^N A_{ij} \|x - x_t^j\|_2^2$, here the particles interact in the Bregman distance. Lemma 14 can be combined with the analysis in [34] to extend (7) to an interacting reflected SDE like (9) that uses $\sum_{j=1}^N A_{ij} D_\Phi(x_i^t, x_t^j) dt$ as an interaction term. This is useful for conveying intuition, but the analysis in the mirror domain is much simpler. One does not have to deal with terms like u_t in (7), and the interaction term in (9) is symmetric, which simplifies the analysis, see [30] for a more detailed discussion.

3.1. A useful reparameterization

It will often be useful to represent the dynamics of the particle system presented in (9) as a vector SDE with all particles stacked in a single vector variable $\mathbf{z}_t = ((z_t^1)^T, \dots, (z_t^N)^T)^T$. We define the graph Laplacian as $L := \text{Diag}(A\mathbf{1}_N) - A$, and let $\mathcal{L} := L \otimes I_d$, where \otimes is the Kronecker product. Using the Laplacian, we can rewrite the evolution of the z_t^i -s in vector form as,

$$d\mathbf{z}_t = (-\nabla \mathbf{V}(\mathbf{z}_t) - \mathcal{L}\mathbf{z}_t) dt + \sigma d\mathbf{B}_t, \tag{12}$$

where $\mathbf{B}_t := ((B_t^1)^T, \dots, (B_t^N)^T)^T$ is the stacked variable of Brownian motions and $\nabla \mathbf{V}(\mathbf{z}_t) = (\nabla v(z_t^1)^T, \dots, \nabla v(z_t^N)^T)^T$. Note that the interaction is linear and the distinctive properties of ISMD compared to N independent copies of SMD are contained \mathcal{L} . Note that given A is doubly stochastic we have

$$\sum_{i=1}^N \sum_{j=1}^N A_{ij} z^j = \sum_{i=1}^N z^i$$

so the smallest eigenvalue of L is zero and $\mathcal{L}1_{dN} = 0$. We will return to this point in Section 4. We assume throughout that the network graph corresponding to the Laplacian is connected, which in turn implies that \mathcal{L} has eigenvalues [41]:

$$\lambda_0 = 0 < \underline{\lambda} \leq \lambda_3 \leq \dots \leq \lambda_{dN}.$$

We thus denote by $\underline{\lambda}$ the smallest non-zero eigenvalue of the Laplacian.

3.2. Convergence of ISMD: a general bound

As we saw earlier, the SMD dynamics does not converge to the minimum due to the noise in the optimization algorithm. Here we show that introducing interaction between the particles can reduce the effect of noise.

We will decompose each particle as a sum of the particle average and a fluctuation term, $z_t^i := \bar{z}_t^N + \tilde{z}_t^i$, where we let

$$\bar{z}_t^N := \frac{1}{N} \sum_{i=1}^N z_t^i, \quad \tilde{z}_t^i := z_t^i - \bar{z}_t^N,$$

and also define $y_t^N := \nabla \Phi^*(\bar{z}_t^N)$.

We begin by deriving the evolution of the average of the particles \bar{z}_t^N . Observe,

$$\begin{aligned} d\bar{z}_t^N &= -\frac{1}{N} \sum_{i=1}^N \nabla f(x_t^i) dt + \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^N A_{ij} (z_t^j - z_t^i) dt + \frac{\sigma}{N} \sum_{i=1}^N dB_t^i \\ &= -\frac{1}{N} \sum_{i=1}^N \nabla f(x_t^i) dt + \frac{\sigma}{N} \sum_{i=1}^N dB_t^i, \end{aligned}$$

where we have used the fact that $\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^N A_{ij} z_t^j = \bar{z}_t^N$ due to the matrix A being doubly stochastic. The particle average \bar{z}_t moves along the gradient of f towards the minimum x^* . At the same time the interaction aims to control fluctuations of each particle z_t^i around \bar{z}_t . This will appear clearly in the bounds in the subsequent results. More specifically, we obtain a result that the distance of a particle to the optimum is bounded by terms standard to optimization with the variance reduced by a factor of N and terms related to the fluctuation, which we will show is bounded.

We proceed by presenting the analog of Proposition 12 for the ISMD.

Proposition 15 (Convergence of ISMD for a Convex Objective). Assume that Assumptions 1–3 and 5 hold. We then have,

$$\begin{aligned} \frac{1}{T} \int_0^T \mathbb{E} [(f(x_t^i) - f(x^*))] dt &\leq \frac{1}{2T} D_{\Phi, \mathcal{X}}^2 + \frac{\sigma^2}{2N} \|\Delta \Phi^*\|_\infty \\ &\quad + \int_0^T \frac{L}{\mu T} \mathbb{E} [\|\tilde{z}_t^i\|_*] dt \\ &\quad + \int_0^T \frac{2L}{\mu NT} \sum_{i=1}^N \mathbb{E} [\|\tilde{z}_t^i\|_*] dt. \end{aligned}$$

Proof. Observe that,

$$\begin{aligned} \int_0^T (f(x_t^i) - f(x^*)) dt &= \int_0^T (f(y_t^N) - f(x^*)) dt + \int_0^T (f(x_t^i) - f(y_t^N)) dt \\ &\leq \int_0^T (f(y_t^N) - f(x^*)) dt + \int_0^T L \|x_t^i - y_t^N\| dt \\ &\leq \int_0^T (f(y_t^N) - f(x^*)) dt + \int_0^T \frac{L}{\mu} \|\tilde{z}_t^i\|_* dt, \tag{13} \end{aligned}$$

where we have used the L -Lipschitz continuity of f and the conjugate correspondence theorem for $\nabla\Phi^*$. Furthermore,

$$\begin{aligned} \int_0^T (f(y_t^N) - f(x^*)) dt &= \int_0^T \frac{1}{N} \sum_{i=1}^N (f(x_t^i) - f(x^*)) dt \\ &+ \int_0^T \frac{1}{N} \sum_{i=1}^N (f(y_t^N) - f(x_t^i)) dt \\ &\leq \int_0^T \frac{1}{N} \sum_{i=1}^N \left(f(x_t^i) - f(x^*) + \frac{L}{\mu} \|\bar{z}_t^i\|_* \right) dt, \end{aligned} \quad (14)$$

where we have again used the Lipschitz continuity of both f and $\nabla\Phi^*$. We furthermore have, by convexity of f ,

$$\begin{aligned} \int_0^T \frac{1}{N} \sum_{i=1}^N f(x_t^i) - f(x^*) dt &\leq \int_0^T \frac{1}{N} \sum_{i=1}^N (x_t^i - x^*)^T \nabla f(x_t^i) dt \\ &\leq \int_0^T \frac{1}{N} \sum_{i=1}^N \nabla f(x_t^i)^T (x_t^i - y_t^N) dt + \int_0^T \frac{1}{N} \sum_{i=1}^N \nabla f(x_t^i)^T (y_t^N - x^*) dt \\ &\leq \int_0^T \frac{L}{\mu N} \sum_{i=1}^N \|\bar{z}_t^i\|_* dt + \int_0^T \frac{1}{N} \sum_{i=1}^N \nabla f(x_t^i)^T (y_t^N - x^*) dt, \end{aligned} \quad (15)$$

where we used in the last inequality that from [Assumption 5](#) it holds $\|\nabla f(x)\|_* \leq L$ and the Lipschitz continuity of $\nabla\Phi^*$. Using [Lemma 24](#) (proved in the Appendix [Appendix A](#)), combining inequalities (13), (14) and (15), and taking expected values the result follows. \square

A similar result can be obtained for a strongly convex f .

Proposition 16 (Convergence of ISMD for a Strongly Convex Objective). Assume that [Assumptions 1–4](#) and [6](#) hold. Then it holds,

$$\begin{aligned} \int_0^T e^{\mu(t-T)} \mathbb{E}[(f(x_t^i) - f(x^*))] dt &\leq \frac{1}{2} e^{-\mu T} D_{\Phi, x}^2 + \frac{\sigma^2}{2N\mu} (1 - e^{-\mu T}) \|\Delta\Phi^*\|_\infty \\ &+ \int_0^T e^{\mu(t-T)} \frac{L}{\mu} \mathbb{E}[\|\bar{z}_t^i\|_*] dt + \int_0^T e^{\mu(t-T)} \frac{2L + \mu^2}{\mu N} \sum_{i=1}^N \mathbb{E}[\|\bar{z}_t^i\|_*] dt. \end{aligned}$$

Proof. Observe that, using strong convexity, we have,

$$\begin{aligned} \frac{1}{N} \sum_{i=1}^N (f(x_t^i) - f(x^*)) &\leq \frac{1}{N} \sum_{i=1}^N (-D_f(x^*, x_t^i) + (x_t^i - x^*)^T \nabla f(x_t^i)) \\ &\leq \frac{1}{N} \sum_{i=1}^N (-\mu D_{\Phi}(x^*, x_t^i) + (x_t^i - x^*)^T \nabla f(x_t^i)) \\ &= \frac{1}{N} \sum_{i=1}^N (-\mu D_{\Phi^*}(z_t^i, z^*) + \mu D_{\Phi^*}(\bar{z}_t^N, z^*)) \\ &+ \frac{1}{N} \sum_{i=1}^N \nabla f(x_t^i)^T (x_t^i - y_t^N) \\ &+ \frac{1}{N} \sum_{i=1}^N \nabla f(x_t^i)^T (y_t^N - x^*) - \mu D_{\Phi^*}(\bar{z}_t^N, z^*). \end{aligned}$$

Therefore,

$$\begin{aligned} \int_0^T e^{\mu t} \frac{1}{N} \sum_{i=1}^N (f(x_t^i) - f(x^*)) dt \\ \leq \int_0^T e^{\mu t} \frac{1}{N} \sum_{i=1}^N \mu \|\bar{z}_t^i\|_* dt + \int_0^T e^{\mu t} \frac{1}{N} \sum_{i=1}^N \frac{L}{\mu} \|\bar{z}_t^i\|_* dt \end{aligned}$$

$$\begin{aligned} &+ \frac{1}{2} D_{\Phi, x}^2 + \frac{\sigma^2}{2N} (e^{\mu T} - 1) \|\Delta\Phi^*(\bar{z}_t^N)\|_\infty \\ &+ \int_0^T e^{\mu t} \frac{\sigma}{\sqrt{N}} \|y_t^N - x^*\|_2 dW_t, \end{aligned} \quad (16)$$

where we have used the Lipschitz continuity of $\nabla\Phi^*$, the assumption that $\|\nabla f(x)\|_* \leq L$, the assumption that

$$D_{\Phi^*}(\bar{z}_t^N, z^*) - D_{\Phi^*}(z_t^i, z^*) \leq \|\bar{z}_t^N - z_t^i\|_*,$$

and [Lemma 25](#) (proved in [Appendix A](#) of the Appendix). Multiplying by $e^{-\mu T}$, using the assumption that $\|\nabla f(x)\|_* \leq L$ and the Lipschitz continuity of $\nabla\Phi^*$, combining inequalities (13), (14) and (16), and taking expected values the result follows. \square

The deviation from the minimum is upper-bounded by four terms. The first two terms are the standard optimization errors, where we observe that the noise variance is reduced by a factor of N . The third and fourth terms are penalties incurred due to each of the particles having different values. These two terms measure the deviation of each individual particle from the particle average. Loosely speaking, if the fluctuation term $\|\bar{z}_t^i\|_*$ is bounded and small then the interaction will drive the particle system to a state near consensus and the gradient terms will direct the particles towards the minimum. There is furthermore a tradeoff between the interaction and the variance. If the interaction term is bounded and not increasing with N , the more particles, the smaller the distance to the optimum, as is witnessed by the term $\frac{\sigma^2}{2N} \|\nabla\Phi^*\|_\infty$. This implies that the variance of the particles x_t^i is also smaller, and they lie closer around the optimal value. Comparing [Proposition 15](#) in the convex case and [Proposition 16](#) in the μ -strongly convex case we observe that the strong convexity of f can speed up convergence, mainly due to the additional factor $\frac{1}{\mu}$. We remark that the results presented in this section thus show that the expected value of the distance of the objective function evaluated at the time average to the objective function evaluated in the optimum decreases as N increases. In addition to this, as we will present in our numerical results later in [Section 5](#) that both the variance of the cost function as well as the variance of the samples are smaller when using interaction.

3.2.1. A comparison with averaging

Consider the case of the convex objective. The interaction in the previous section is introduced in order to obtain a way to control the effects of noise. In particular we observed that if the term $\|\bar{z}_t^i\|_*$ was bounded and not an increasing function of N , the variance could be reduced by having interacting particles. We can compare the setting with interaction to one in which we simply average the trajectories of the particles.

If the following holds,

$$\frac{1}{N} \sum_{i=1}^N \nabla f(x_t^i) = \frac{1}{N} \sum_{i=1}^N \nabla f \circ \nabla\Phi^*(z_t^i) = \nabla f \circ \nabla\Phi^* \left(\frac{1}{N} \sum_{i=1}^N z_t^i \right) := \nabla f(y_t^N), \quad (17)$$

then we have from [Lemma 24](#) (in [Appendix A](#)) and the convexity of f that

$$\mathbb{E} \left[\frac{1}{T} \int_0^T (f(y_t^N) - f(x^*)) dt \right] \leq \frac{1}{2T} D_{\Phi, x}^2 + \frac{\sigma^2}{2N} \|\Delta\Phi^*(\bar{z}_t^N)\|_\infty.$$

Notice that the Assumption in (17) holds for a function of the form $f(x) = a^T x + b + \Phi(x)$ which is the particular loss function considered in [13]. Thus in the setup of this kind structure in the loss function and gradient, averaging the particles can decrease the effect of the noise.

In a general setting where the loss function is nonlinear, one can decrease the effects of noise by averaging the particles trajectories. Combining (14), (15) and Lemma 24 we have,

$$\mathbb{E} \left[\frac{1}{T} \int_0^T (f(y_t^N) - f(x^*)) dt \right] \leq \frac{1}{2T} D_{\phi, \mathcal{X}}^2 + \frac{\sigma^2}{2N} \|\Delta \Phi^*(\bar{z}_t^N)\|_\infty + \int_0^T \frac{2L}{\alpha TN} \sum_{i=1}^N \|\bar{z}_t^i\|_* dt.$$

In the nonlinear case the average of the gradients is not the gradient of the average and therefore the additional term representing the deviations of each individual particle from the particle average plays a role.

3.3. Bounding the fluctuation term

In this section we present a bound on the fluctuation term. When the fluctuation term is sufficiently small all particles have approximately the same value; we refer to this as the particles having achieved consensus. The dynamics of the fluctuation term in a vectorized form is given by,

$$d\bar{z}_t = \left[\left(\frac{1}{N} \mathbf{1}_N \mathbf{1}_N^T \otimes I_d - I_{dN} \right) \nabla \mathbf{V}(\mathbf{z}_t) - \mathcal{L} \mathbf{z}_t \right] dt + \sigma \sqrt{\frac{(N-1)}{N}} d\mathbf{B}_t,$$

where we have used that $(I_{dN} - \frac{1}{N} \mathbf{1}_N \mathbf{1}_N^T \otimes I_d) d\mathbf{B}_t \stackrel{d}{=} \sqrt{\frac{(N-1)}{N}} d\mathbf{B}_t$.

Proposition 17 (Bounding the Fluctuation Term: a General Result for the Convex Case). Assume that Assumptions 1–3 and 7 hold. Then,

$$\frac{1}{TN} \int_0^T \sum_{i=1}^N \mathbb{E} [\nu(z_t^i) - \nu(\bar{z}_t^N)] dt \leq \frac{L}{2TN} \sum_{i=1}^N \|\bar{z}_0^i\|_2^2 + \frac{1}{2} d \sigma^2 \frac{(N-1)}{N}.$$

Proof. Define the Lyapunov function,

$$V_t(\bar{\mathbf{z}}) := \frac{1}{2} \bar{\mathbf{z}}^T \bar{\mathbf{z}}.$$

By Itô's lemma it then follows that,

$$dV_t = \bar{\mathbf{z}}^T \left[\left(\frac{1}{N} \mathbf{1} \mathbf{1}^T \otimes I_d - I_{dN} \right) \nabla \mathbf{V}(\mathbf{z}_t) - \mathcal{L} \mathbf{z}_t \right] dt + dN \frac{1}{2} \sigma^2 \frac{(N-1)}{N} dt + \bar{\mathbf{z}}^T \sigma \sqrt{\frac{(N-1)}{N}} d\mathbf{B}_t.$$

Define $\bar{z}_t^N = \mathbf{1}_N \otimes z_t^N$. We can rewrite the drift term of dV_t as,

$$\mu_t := \bar{\mathbf{z}}^T \left(\frac{1}{N} \mathbf{1} \mathbf{1}^T \otimes I_d \right) \nabla \mathbf{V}(\mathbf{z}_t) - \bar{\mathbf{z}}^T \nabla \mathbf{V}(\mathbf{z}_t) - \bar{\mathbf{z}}^T \mathcal{L} \mathbf{z}_t + dN \frac{1}{2} \sigma^2 \frac{(N-1)}{N}.$$

Note that,

$$\bar{\mathbf{z}}^T \left(\frac{1}{N} \mathbf{1} \mathbf{1}^T \otimes I_d \right) \nabla \mathbf{V}(\mathbf{z}_t) = 0.$$

By Assumption 5,

$$-\bar{\mathbf{z}}^T \nabla \mathbf{V}(\mathbf{z}_t) \leq \sum_{i=1}^N (\nu(\bar{z}_t^N) - \nu(z_t^i)).$$

Furthermore, by the properties of the graph Laplacian, it holds $\mathcal{L} \mathbf{1}_{dN} = 0$, so that,

$$-\bar{\mathbf{z}}^T \mathcal{L} \mathbf{z}_t = -\bar{\mathbf{z}}^T \mathcal{L} \bar{\mathbf{z}}_t \leq -\underline{\lambda} \bar{\mathbf{z}}_t^T \bar{\mathbf{z}}_t \leq 0.$$

Then,

$$V_T \leq V_0 + \int_0^T \sum_{i=1}^N (\nu(\bar{z}_t^N) - \nu(z_t^i)) dt + TdN \frac{1}{2} \sigma^2 \frac{(N-1)}{N}$$

$$+ \int_0^T \sigma \sqrt{\frac{(N-1)}{N}} \|\bar{\mathbf{z}}_t\|_2 d\mathbf{B}_t,$$

where we have additionally used that $\bar{\mathbf{z}}^T d\mathbf{B}_t \stackrel{d}{=} \|\bar{\mathbf{z}}\|_2 d\mathbf{B}_t$. Rearranging, using the property of the Itô integral, and using $V_T \geq 0$ we find,

$$\frac{1}{TN} \int_0^T \sum_{i=1}^N \mathbb{E} [\nu(z_t^i) - \nu(\bar{z}_t^N)] dt \leq \frac{1}{T} V_0 + d \frac{1}{2} \sigma^2 \frac{(N-1)}{N}. \quad \square$$

We note that in the above result we used the convexity of ν but we did not exploit the interaction between the particles. From the above statement we observe that in a deterministic setting, i.e. if $\sigma = 0$, consensus can be achieved even without interaction. Each particle is driven towards the optimum x^* by the gradient term. Consensus is in this case achieved exactly at optimality. In the stochastic setting exact consensus can no longer be achieved and the distance to optimality remains bounded by a function of the noise.

Proposition 18 (Bounding the Fluctuation Term: the Strongly Convex Case). Let Assumptions 1–3 and 8 hold. Then,

$$\mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N \|\bar{z}_t^i\|_*^2 \right] \leq \frac{K}{N} e^{-(\kappa+\underline{\lambda})t} \sum_{i=1}^N \|\bar{z}_0^i\|_2^2 + \frac{dK}{(\kappa+\underline{\lambda})} \sigma^2 \frac{(N-1)}{N} (1 - e^{-(\kappa+\underline{\lambda})t}).$$

Proof. Define again the Lyapunov function,

$$V_t(\bar{\mathbf{z}}) := \frac{1}{2} \bar{\mathbf{z}}^T \bar{\mathbf{z}}.$$

It is clear that $V_t = 0$ if $z_t^i = z_t^j$, or in other words consensus has been achieved and the fluctuation term is zero. By Itô's lemma it then follows that,

$$dV_t = \bar{\mathbf{z}}^T \left[\left(\frac{1}{N} \mathbf{1} \mathbf{1}^T \otimes I_d - I_{dN} \right) \nabla \mathbf{V}(\mathbf{z}_t) - \mathcal{L} \mathbf{z}_t \right] dt + dN \frac{1}{2} \sigma^2 \frac{(N-1)}{N} dt + \bar{\mathbf{z}}^T \sigma \sqrt{\frac{(N-1)}{N}} d\mathbf{B}_t.$$

Define $\bar{z}_t^N = \mathbf{1}_N \otimes z_t^N$. We can rewrite the drift term of dV_t as,

$$\mu_t := \bar{\mathbf{z}}^T \left(\frac{1}{N} \mathbf{1} \mathbf{1}^T \otimes I_d \right) \nabla \mathbf{V}(\mathbf{z}_t) + \bar{\mathbf{z}}^T (\nabla \mathbf{V}(\bar{z}_t^N) - \nabla \mathbf{V}(\mathbf{z}_t)) - \bar{\mathbf{z}}^T \nabla \mathbf{V}(\bar{z}_t^N) - \bar{\mathbf{z}}^T \mathcal{L} \mathbf{z}_t + dN \frac{1}{2} \sigma^2 \frac{(N-1)}{N}.$$

Note that, since $\sum_{i=1}^N (z_t^i - \bar{z}_t^N) = 0$,

$$\bar{\mathbf{z}}^T \left(\frac{1}{N} \mathbf{1} \mathbf{1}^T \otimes I_d \right) \nabla \mathbf{V}(\mathbf{z}_t) = 0, \quad \bar{\mathbf{z}}^T \nabla \mathbf{V}(\bar{z}_t^N) = 0.$$

By Assumption 8,

$$\bar{\mathbf{z}}^T (\nabla \mathbf{V}(\bar{z}_t^N) - \nabla \mathbf{V}(\mathbf{z}_t)) \leq -\kappa \bar{\mathbf{z}}^T \bar{\mathbf{z}}.$$

Furthermore, by the properties of the Laplacian, it holds $\mathcal{L} \mathbf{1}_{dN} = 0$, so that,

$$\mathcal{L} \mathbf{z}_t = \mathcal{L} \bar{\mathbf{z}}_t.$$

Therefore,

$$\mu_t \leq -\kappa \bar{\mathbf{z}}^T \bar{\mathbf{z}} - \bar{\mathbf{z}}^T \mathcal{L} \bar{\mathbf{z}} + dN \frac{1}{2} \sigma^2 \frac{(N-1)}{N} \leq -(\kappa + \underline{\lambda}) \bar{\mathbf{z}}^T \bar{\mathbf{z}} + dN \frac{1}{2} \sigma^2 \frac{(N-1)}{N}.$$

Integrating the expression and using the above bound we obtain,

$$V_t \leq e^{-(\kappa+\underline{\lambda})t} V_0 + \frac{1}{2} \int_0^t e^{-(\kappa+\underline{\lambda})(t-s)} dN \sigma^2 \frac{(N-1)}{N} ds + \int_0^t e^{-(\kappa+\underline{\lambda})(t-s)} \sigma \sqrt{\frac{(N-1)}{N}} \|\tilde{z}_s\|_2 d\mathbf{B}_s,$$

where we have additionally used that $\tilde{z}^T d\mathbf{B}_t \stackrel{d}{=} \|\tilde{z}\|_2 d\mathbf{B}_t$. Then, taking expectations we have:

$$\mathbb{E}[V_t] \leq e^{-(\kappa+\underline{\lambda})t} V_0 + \frac{1}{2(\kappa+\underline{\lambda})} dN \sigma^2 \frac{(N-1)}{N} (1 - e^{-(\kappa+\underline{\lambda})t}).$$

The statement follows from norm equivalence, i.e. $\|\cdot\|_* \leq K \|\cdot\|_2$. \square

Remark 19 (Constant in the Norm Equivalence). If the Euclidean norm is used above in Assumptions 1–2, then $K = 1$. In certain cases however, K can be dimension-dependent.

We furthermore note, by the bound on $\mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N \|\tilde{z}_t^i\|_*^2 \right]$ implies a bound on $\mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N \|\tilde{z}_t^i\|_* \right]$ with the dependence on N remaining the same. This is formalized in the following lemma.

Lemma 20. *It holds for some constant K ,*

$$\mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N \|\tilde{z}_t^i\|_* \right] \leq \left(\frac{K}{N} e^{-(\kappa+\underline{\lambda})t} \sum_{i=1}^N \|\tilde{z}_0^i\|_2^2 + \frac{dK}{(\kappa+\underline{\lambda})} \sigma^2 \frac{(N-1)}{N} (1 - e^{-(\kappa+\underline{\lambda})t}) \right)^{\frac{1}{2}}.$$

Proof. Let $\|\cdot\|$ be an arbitrary norm. Observe that by Hölder’s inequality $\mathbb{E}[|x|] \leq (\mathbb{E}[|x|^2])^{\frac{1}{2}}$. Let $x = \sum_{i=1}^N \|z_t^i\|$. Observe that,

$$\begin{aligned} |x|^2 &= \sum_{i=1}^N \|z_t^i\|^2 + \sum_{i=1}^N \sum_{j \neq i}^N \|z_t^i\| \|z_t^j\| \\ &\leq \sum_{i=1}^N \|z_t^i\|^2 + \frac{1}{2} \sum_{i=1}^N \sum_{j \neq i}^N (\|z_t^i\|^2 + \|z_t^j\|^2) \\ &\leq (1+N) \sum_{i=1}^N \|z_t^i\|^2, \end{aligned}$$

where we have applied Young’s inequality in the second inequality. We obtain,

$$\mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N \|z_t^i\| \right] \leq \mathbb{E} \left[\frac{1+N}{N^2} \sum_{i=1}^N \|z_t^i\|^2 \right]^{\frac{1}{2}}.$$

The statement follows using Proposition 18. \square

The above statements show that for a strongly convex objective approximate consensus can be achieved for a sufficiently high t and for a sufficiently large $\kappa + \underline{\lambda}$. In the strongly convex case approximate consensus can be achieved even in the case of no interaction if κ is sufficiently large. For the convex case we have $\kappa = 0$. In this case approximate consensus can be achieved if $\underline{\lambda}$ is sufficiently high. For a convex objective without interaction we would be left with the noise term similar to the result in Proposition 17. Without interaction the strong convexity of the objective thus determines in how much the effect of noise can be reduced. With interaction the interaction type itself, i.e. the $\underline{\lambda}$, also plays a role and for a sufficiently high $\underline{\lambda}$ the effects of noise on consensus can be minimized. We remark furthermore that we

bound the dual norm by the L_2 norm using norm equivalence; the downside of this is that an additional dimensionality-dependence can be obtained. This could be mitigated by working with directly the dual norm as in [13] but in this case the obtained result is rather limited to a specific set of objective functions; alternatively one can work in the primal space directly. The latter idea will be addressed in future work.

3.3.1. Extensions to distributed optimization

A common setting is to consider minimizing $f(x) = \sum_{i=1}^N f_i(x)$ and each particle corresponds to a computing worker or processor having access only to f_i and thus $\nabla f_i \circ \nabla \Phi^*(z_t^i)$ would be used in (9) (instead of $\nabla f \circ \nabla \Phi^*(z_t^i)$). The proof of Propositions 17 and 18 is based the fact that each particle has access to the full objective function f . As a result terms that would be relevant in a distributed setup related to variability of each f_i vanish. Clearly, this results in a better bound and thus a better convergence result, but the current results can be extended for the distributed case by increasing the interaction strength and multiplying A in (9) by a sufficiently high constant to ensure consensus arises. The details are left for future work.

3.3.2. On the tradeoff between noise reduction and interaction

From the results in Section 3.2 we observed that for interacting particles the convergence can be improved since the effect of noise in the term $\frac{1}{N} \sigma \|\Delta \Phi(z_t^i)\|_\infty$ is reduced by a factor $1/N$. We however remark that this will only be achieved if the term $\|\tilde{z}_t^i\|_*$ is bounded and non-increasing with N . As we showed in Section 3.3 we can present a bound on the interaction term. In particular, we showed that the fluctuation is bounded and non-increasing with N . The fluctuation is bounded in the L_2 norm, and a noise term does remain; as discussed, this noise term can be controlled by the strong convexity of the objective or by imposing an interaction strength, which means replacing A with ϑA for $\vartheta > 1$. In other words, when $\underline{\lambda} + \kappa$ is sufficiently high the fluctuation term is sufficiently small. Specifically, in the strongly convex case for a large N, T the fluctuation term is bounded by $\frac{dK}{(\kappa+\underline{\lambda})} \sigma^2$ for some constant K . We conclude this section by saying that as long as the decrease in the value of $\frac{1}{N} \sigma \|\Delta \Phi(z_t^i)\|_\infty$ is larger than the increase in the value of $\|\tilde{z}_t^i\|_*$, interaction with N particles achieves a closer convergence to the optimum. In other words, the amount of particles can be seen as an alternative to a decreasing learning rate or vanishing noise variance (e.g. the latter can be achieved by increasing the mini-batch size).

4. Understanding performance from convergence to stationarity

We proceed with exploiting tools from the analysis of SDEs and in particular the rate of convergence to their invariant distribution. Studying the convergence of the $Law(z_t)$ or $Law(\mathbf{z}_t)$ is a cornerstone in the analysis of sampling schemes and can also provide valuable insights into the optimization problems; see [42,43] for recent works in this direction. In particular, we will discuss the convergence of (8) or (9) to the corresponding stationary distributions. Note that in the spirit of dual (or Nesterov) averaging the SDE in (9) is fairly standard as it evolves only in the mirror space. In the analysis so far we were investigating how close a time average $\frac{1}{T} \int_0^t f(x_t^i) dt$ (or $\frac{1}{T} \int_0^t f \circ \nabla \Phi^*(z_t^i) dt$) is to $f(x^*)$ using a Lyapunov method based on Bregman divergence. We will complement these results with rates of convergence based on logarithmic Sobolev inequalities based on the celebrated Bakry-Emery approach, see [28,29] for details.

4.1. Stationary distributions and particle correlations

We first consider the case when particles do not interact. One approach for finding x^* is to run N independent copies of (8) for a long time and then select best particle i^* as $\arg \min_{i \in \{1, \dots, N\}} f(x_t^i)$. This intuitive approach is based on $\min f$ being equivalent to $\max \exp(-\frac{2}{\sigma^2} \mathcal{V})$ and on the convergence of $\frac{d \text{Law}(z_t)}{dz}(z)$ to $\frac{1}{Z} \exp(-\frac{2}{\sigma^2} \mathcal{V})$. When stacking the particles together in \mathbf{z} , then one can postpone choosing the best particle and consider $\min_{\mathbf{x}} \sum_{i=1}^N f(x^i)$ or $\min_{\mathbf{z}} \sum_{i=1}^N V(z^i)$ instead, with each particle i here sharing the same properties and dynamics.

To make this a bit clearer for our setup, denote with $\eta_t = \text{Law}(\mathbf{z}_t)$ the law of N independent particles each following (8); or equivalently each z_t^i evolving as (9) with $A = 0$. It is well-known (see e.g. [36]) that under appropriate assumptions on the objective \mathcal{V} , the stationary distribution of this SDE is given by,

$$\eta_\infty(d\mathbf{z}) = \frac{1}{Z} \exp\left(-\frac{2}{\sigma^2} \sum_{i=1}^N \mathcal{V}(z^i)\right) d\mathbf{z},$$

with Z being a finite normalizing constant. It is clear that finding the mode of η_∞ is equivalent to solving the following optimization problem:

$$\mathbf{z}^* = \arg \min_{\mathbf{z} \in \mathbb{R}^{dN}} \sum_{i=1}^N \mathcal{V}(z^i). \tag{18}$$

An important observation here is that independence of each i is not crucial and one can modify the cost $\mathcal{V}(z^i)$ by adding terms that do not affect the minimizer. This is true when using instead the following cost function for \mathbf{z} :

$$\mathcal{W}(\mathbf{z}) = \sum_{i=1}^N \mathcal{V}(z^i) + \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N A_{ij} (z^i - z^j)^2. \tag{19}$$

Note that the quadratic term due to interaction terms acts as a regularizer, which aims to impose consensus i.e. $z^i = \bar{z}^N$. We can make the this claim more precise in the following lemma.

Lemma 21 (Interaction Preserves the Minimum). *Let $\mathbf{z}^* := \arg \min_{\mathbf{z}} \mathcal{V}(\mathbf{z})$, and let $\mathcal{W}(\mathbf{z})$ be as defined in (19). Then $\mathbf{z}^* = (z^{*T}, \dots, z^{*T})^T$ is a minimizer of $\mathcal{W}(\mathbf{z})$.*

Proof. From the definition of \mathbf{z}^* and the optimality conditions we have $\nabla \mathcal{V}(\mathbf{z}^*) = 0$ and that there exists $\beta > 0$ such that $\text{Hess}(\mathcal{V})(\mathbf{z}^*) \succeq \beta I_d$. We have,

$$(\nabla \mathcal{W}(\mathbf{z}))_i = \nabla \mathcal{V}(z^i) + \sum_{j \neq i} A_{ij} (z_i - z_j).$$

Clearly, when $z^i = \mathbf{z}^*$ for every i we then have $\nabla \mathcal{W}(\mathbf{z}) = 0$. Furthermore:

$$(\text{Hess})(\mathcal{W}(\mathbf{z}))_{ii} = \text{Hess}(\mathcal{V}(z^i)) + \sum_{j \neq i} A_{ij}, \tag{20}$$

and

$$(\text{Hess})(\mathcal{W}(\mathbf{z}))_{ij} = -A_{ij}. \tag{21}$$

This implies

$$(\text{Hess})(\mathcal{W}(\mathbf{z}^*))_{ii} \geq \beta - \sum_{j \neq i} (\text{Hess})(\mathcal{W}(\mathbf{z}^*))_{ij},$$

so that $\text{Hess}(\mathcal{W})(\mathbf{z}^*) \succeq \beta I_{dN}$ and \mathbf{z}^* minimizes \mathcal{W} since $\beta > 0$. \square

We will now extend the previous discussion on invariant distributions and \mathbf{z}^* for ISMD and (9) with A being nonzero. We

will denote the distribution of the joint particle system and each marginal as follows:

$$\eta_t^N = \text{Law}(\mathbf{z}_t) \quad \text{and} \quad \eta_t^{i,N} = \text{Law}(z_t^i).$$

Recall \mathbf{z}_t evolves as in (12), whose the invariant distribution is given as

$$\eta_\infty^N(d\mathbf{z}) = \frac{1}{Z_N} \exp\left(-\frac{2}{\sigma^2} \left(\sum_{i=1}^N \mathcal{V}(z^i) + \frac{1}{2} \mathbf{z}^T \mathcal{L} \mathbf{z}\right)\right) d\mathbf{z},$$

where Z_N is the normalization constant. Similar as in the non-interacting setting, finding the mode of η_∞^N is equivalent to solving the following optimization problem:

$$\bar{\mathbf{z}} = \arg \min_{\mathbf{z} \in \mathbb{R}^{dN}} \mathcal{W}(\mathbf{z}),$$

which we saw earlier in Lemma 21 is equivalent to solving (18).

Despite introducing correlations between the particles through \mathcal{L} , the mode of the invariant distribution of the interacting particle system is exactly the minimizer of the objective \mathcal{V} and is achieved when all particles are at consensus. The presence of the noise means that exact consensus or synchronization of all particles cannot be achieved. Similar to SMD stochasticity implies that at stationarity the z_t^i will behave as samples from the invariant distribution and thus oscillating together around the optimum \mathbf{z}^* . In this context it is convenient to use convergence results for $\eta_t^{i,N}$ for deriving performance bounds for optimization or assessing the level of consensus based on differences between $\eta_t^{i,N}$ and $\eta_t^{j,N}$.

In addition to establishing positive curvature at \mathbf{z}^* , using diagonal dominance in the final steps in the proof of Lemma 21 implies positive definiteness for \mathcal{W} . This is not surprising given the sum of quadratics in the interaction term acts as a convex regularizer. We summarize this in the corollary below:

Corollary 22. *Let Assumptions 1–3 and 8 hold. Then \mathcal{W} is strongly convex.*

Proof. Assumption 8 implies that $\text{Hess}(\mathcal{V})(z) \succeq \kappa I_d$ for some $\kappa > 0$. Note that (20)–(21) hold for any \mathbf{z} and

$$\text{Hess}(\mathcal{W}) = \text{Hess}\left(\sum_{i=1}^N \mathcal{V}(z^i)\right) + \frac{1}{2} \text{Hess}(\mathbf{z}^T \mathcal{L} \mathbf{z}) \succeq \kappa I_{dN}. \quad \square$$

The result is somewhat negative in that to obtain strict convexity of \mathcal{W} one requires strong convexity of \mathcal{V} , that is adding $\mathbf{z}^T \mathcal{L} \mathbf{z}$ is not sufficient to obtain strong convexity when \mathcal{V} is only convex. This is particularly relevant later when we will apply Bakry–Emery theory and derive log Sobolev inequalities.

4.2. Log-Sobolev inequalities for the particle system

We proceed to present the exponential convergence properties for η_t^N . Our first result in this section is to apply Bakry–Emery theory to derive the Log-Sobolev inequality for the particle system. Similar to the result in Lemma 3.5, Corollary 3.7 in [30] we have the following result:

Proposition 23 (Convergence of ISMD using Bakry–Emery Theory). *Let Assumptions 1–3 and 7 hold (w.r.t the Euclidean norm). Assume also η_0^N satisfies a log-Sobolev inequality with constant C_0 . Then η_t^N satisfies a log-Sobolev inequality for any smooth function f*

$$\text{Ent}_{\eta_t^N}(f^2) \leq C_t \mathbb{E}_{\eta_t^N} [|\nabla f|^2]$$

with constant $C_t = \frac{2}{\rho} (1 - e^{-\rho t}) + C_0 e^{-\rho t}$ where $\rho = \frac{\sigma^2}{2} \kappa$. If in addition [Assumption 8](#) holds, we have

$$H(\eta_t^N | \eta_\infty^N) \leq Ke^{-2\rho t} \quad \text{and} \quad W_2(\eta_t^N, \eta_\infty^N) \leq \sqrt{\frac{K}{2}} e^{-\rho t}, \quad (22)$$

and

$$|\mathbb{E}_{\eta_t^N} \mathcal{W} - \mathcal{W}(\mathbf{z}^*)| \leq \sqrt{\frac{K}{2}} e^{-\rho t} + \frac{\sigma^2}{2} \left(\frac{2dN}{\rho} - \frac{1}{2} \log \left(\frac{\sigma^2}{2L_N} \right) \right), \quad (23)$$

with $L_N = \frac{1}{\mu} + \|\mathcal{L}\|$ where $\|\mathcal{L}\|$ is any matrix norm on \mathcal{L} .

Proof. For the proof we will rely on the Bakry–Emery theory. To state the Bakry–Emery criterion for (12), we need to define the following differential operators:

$$\mathcal{L}(f) = -\nabla \mathcal{W}^T \nabla f + \frac{\sigma^2}{2} \text{Hess}(f),$$

$$\Gamma(f) = \frac{1}{2} (\mathcal{L}(fg) - f\mathcal{L}(g) - g\mathcal{L}(f)),$$

$$\Gamma_2(f) = \frac{1}{2} (\mathcal{L}(\Gamma(f, f)) - 2\Gamma(\mathcal{L}(f), f)).$$

The Bakry–Emery criterion requires to verify $\Gamma_2(f) \geq \rho \Gamma(f, f)$. We have:

$$\Gamma_2(f) - \rho \Gamma(f, f) = \left(\frac{\sigma^2}{2} \right)^2 \|\text{Hess}(f)\|^2 + \frac{\sigma^2}{2} \|\nabla f\|^2 \left(\frac{\sigma^2}{2} \text{Hess}(\mathcal{W}) - \rho I \right),$$

so need to verify

$$\text{Hess}(\mathcal{W}) \geq \frac{2}{\sigma^2} \rho I_{dN},$$

that holds for $\rho = \frac{\sigma^2}{2} \kappa$ from [Assumption 7](#) or [Assumption 8](#), [Corollary 22](#) and $\mathcal{L}\mathbf{1} = 0$. For the second part we explicitly require $\rho > 0$, which follows straightforwardly from [Assumption 8](#). This gives directly C_t and we note that the second term in C_t is due to the random initialization (e.g. see [30, Corollary 3.7]). When [Assumption 8](#) holds and $\rho > 0$ we have

$$H(\eta_t^N | \eta_\infty^N) \leq Ke^{-2\rho t},$$

for $K = H(\eta_0^N | \eta_\infty^N)$; see [Theorem 27](#) in [Appendix B.1](#). [Theorem 26](#) gives $W_2(\eta_t^N, \eta_\infty^N) \leq \sqrt{\frac{K}{2}} e^{-\rho t}$.

For the mean-mode result, we first need to establish $\eta_\infty^N \in \mathcal{P}_2(\mathbb{R}^{dN})$. This can be shown using [Theorem 26](#) in [Appendix B.2](#) with the log Sobolev constant $C_\infty = \frac{2}{\rho}$ implies

$$\eta_\infty^N(e^{sf}) \leq e^{s\eta_\infty^N(f) + C_\infty s^2},$$

for any 1-Lipschitz function f , so consider $f(\mathbf{z}) = \mathbf{z}$ and the moment generating function to get the following moment bounds for $s = 0$

$$\mathbb{E}_{\eta_\infty^N} [\mathbf{z}\mathbf{z}^T] \leq (2C_\infty I_{dN} + E_{\eta_\infty^N}[\mathbf{z}] E_{\eta_\infty^N}[\mathbf{z}]^T),$$

so moments can be bound using exponential integrability and η_∞^N has finite second moment. Combining this with

$$\log \det(\Sigma^N) \leq \text{Tr}(\Sigma^N - I)$$

gives

$$\log \det(\Sigma^N) \leq dN(2C_\infty - 1).$$

We can then gather [Corollary 29](#) and [Proposition 30](#) in [Appendix B.2](#) we get

$$\begin{aligned} \mathbb{E}_{\eta_t^N} \mathcal{W} - \mathcal{W}(\mathbf{z}^*) &\leq \frac{\sigma^2}{2} \left(\frac{dN}{2} + \frac{1}{2} \log \det(\Sigma^N) - \frac{1}{2} \log \left(\frac{\sigma^2}{2L_N} \right) \right) \\ &\leq \frac{\sigma^2}{2} \left(\frac{2dN}{\rho} - \frac{1}{2} \log \left(\frac{\sigma^2}{2L_N} \right) \right) \end{aligned}$$

The statement then follows by noting that,

$$|\mathbb{E}_{\eta_t^N} \mathcal{W}(\mathbf{z}_t) - \mathcal{W}(\mathbf{z}^*)| \leq |\mathbb{E}_{\eta_t^N} \mathcal{W}(\mathbf{z}_t) - \mathbb{E}_{\eta_\infty^N} \mathcal{W}(\mathbf{z}_t)| + |\mathbb{E}_{\eta_\infty^N} \mathcal{W}(\mathbf{z}_t) - \mathcal{W}(\mathbf{z}^*)|,$$

and using the convergence rate from the result derived above in (22) for the first term and the mean-mode bound for the second. \square

This result shows that the rate of convergence to the invariant measure is exponential in time with factor ρ . Using the relationship between the mode of the invariant measure and the minimizer of the objective function, this result implies that with exponential convergence the samples \mathbf{z}_t^i will oscillate around the optimum \mathbf{z}^* for a strongly convex objective. Note that the discrepancy between the mean and the mode at stationarity seen in (23) means this oscillation is not centered around \mathbf{z}^* . Eq. (23) also shows that when the width of η_∞^N is small, i.e. σ^2 is low, then \mathbf{z}_t will lie closer to the minimizer. To identify \mathbf{z}^* in practice, one needs to look at the histogram of \mathbf{z}_t^i (w.r.t i) or the occupation measure of $\mathcal{W}(\mathbf{z}_t)$ after some burn-in time t_0 or a combination of the two. The amount of required burn-in can be lower when fast convergence to stationarity occurs. From the expression for ρ we can see that the rate of convergence ρ increases with the noise variance and the strong convexity coefficient of the objective function. Finally, the dependence on N arrives from dimensional dependence of certain bounds involving Gaussian integrals; see [Proposition 30](#) in [Appendix B.2](#).

At the level of the marginal $\eta_\infty^{i,N}$ a similar bound to (23) can be obtained, where dN is replaced with d . Convergence of η_t^N to η_∞^N implies convergence of the marginals so one can deduce a similar result to [Proposition 23](#) for the marginals and a similar concentration inequality for $\mathcal{V}(\mathbf{z}_t^i)$ showing that the dimension dependence of the optimization is with d and not dN . The details are omitted as the proof is very similar to [Proposition 23](#).

4.2.1. On the tradeoff between noise reduction and interaction: a sampling perspective

One could attempt to compare [Proposition 23](#) with the results obtained in [Section 3.2](#), although they are quite different in nature. In [Propositions 15](#) and [16](#) we showed that the expected time average converges to a neighborhood of the optimizer at rate $e^{-\mu t}$ for the strongly convex objective (with μ being the convexity constant for f). [Proposition 23](#) and specifically the result in (23) looks directly at $\mathcal{W}(\mathbf{z}_t) - \mathcal{W}(\mathbf{z}^*)$ and uses instead κ , the convexity constant of \mathcal{V} , which was also used to bound the fluctuation terms in [Section 3.3](#). While this means direct comparisons are case specific and depend on the choice of the mirror map Φ , the convergence rates are in both settings exponential. In other words, for the strongly convex case both [Proposition 16](#) and [Proposition 23](#) lead to exponential rate albeit with different constants due to working with f and \mathcal{V} respectively.

Furthermore, in terms of the tradeoffs between interaction and convergence, both bounds give a similar result. Observe that from [Proposition 16](#) combined with [Proposition 18](#) we obtain,

$$\frac{1}{N} \sum_{i=1}^N \mathbb{E} [f(x_\infty^i) - f(x^*)] \leq K \left(\frac{\sigma^2}{2N} + \frac{d(N-1)}{(\kappa + \underline{\lambda})N} \sigma^2 \right)^{1/2},$$

where we remark that we have simplified notation to shorten the presentation and we assume $T \gg 0$ and denote with x_∞^i the corresponding exponential time average. The expectation is furthermore taken with respect to the initial condition. Furthermore, from (23),

$$\frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\eta_\infty^N} [\mathcal{V}(z^i) - \mathcal{V}(z^*)] \leq K \left(\frac{2d}{\kappa} + \frac{\sigma^2}{4N} \log \left(\frac{2L_N}{\sigma^2} \right) \right).$$

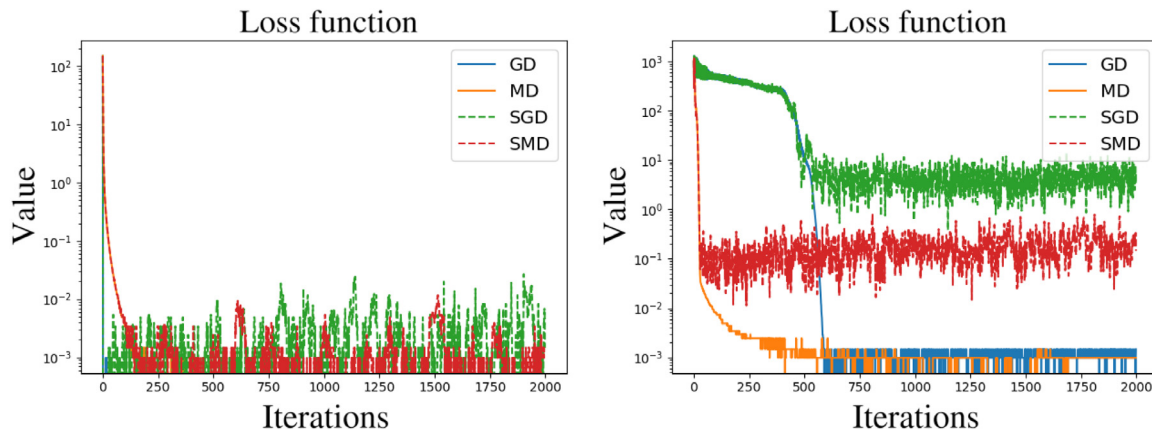


Fig. 1. A comparison between GD, MD, SGD and SMD for the linear system with condition number 10 with $\eta = 0.3/\sqrt{t}$ (L) and condition number 100 with $\eta = 0.1/\sqrt{t}$ (R). MD attains a better performance, both in the deterministic as stochastic setting for a high condition number.

From the above bounds – while as mentioned different in nature – we observe that both consist of a term which decreases with N , and a term which is of $\mathcal{O}(1)$ in N . Both bounds additionally decrease with an increase in κ , or the strong convexity constant of the objective function.

4.3. Mapping back to the primal space

Convergence in the mirror domain and in z_t^i should imply convergence in the primal domain and x_t^i . Given Φ is invertible the law of X_t follows from standard change of variable/coordinates of continuous random variables. This is we have for all $t \geq 0$:

$$\frac{dLaw(x_t^i)}{dx}(x) = \frac{d\eta_t^{i,N}}{dx} \circ \nabla \Phi^*(x) |\det(\nabla^2 \Phi^*(x))|.$$

Recall an implication of Proposition 23 is that $W_2(\eta_t^N, \eta_\infty^N) \leq K'e^{-\rho t}$, so this is equivalent to

$$W_2(Law(x_t^i), Law(x_\infty^i)) \leq K'e^{-\rho t}.$$

In the context of Bakry–Emery theory and Markov semigroups we refer the interested reader to Section 1.15.1 of [29] and to [24] for a rigorous treatment related to sampling methods including equivalence in total variation norm.

5. Numerical results

In this section we report numerical experiments using two standard benchmark problems. The first problem is based on linear regression and is an adapted version of the experimental setup that appeared in [6]. The second problem is based on a variation of the classic traffic assignment problem [2]. We consider compare between (Stochastic) Gradient Descent with Euclidean projections ((S)GD), (Stochastic) Mirror Descent with a Bregman divergence generated by the entropy function. In our numerical implementations of the optimization algorithm we use the Euler discretization of the continuous SMD dynamics:

$$z_{t+1}^i = z_t^i - \eta \epsilon \nabla f(x_t^i) + \epsilon \sum_{j=1}^N A_{ij}(z_t^j - z_t^i) + \sigma \sqrt{\epsilon} \mathcal{N}(0, 1).$$

The parameters that we have to choose are the time step ϵ , the learning rate η , the number of training iterations T_N and the noise level σ (set to zero in case of GD or MD). Unless otherwise mentioned we set $A_{ij} = \frac{1}{N}$, i.e. a mean-field interaction matrix.

5.1. Linear system with simplex constraints

We consider a first similar set-up as in Example 9.19 from [6] to compare mirror descent with projected gradient descent. Consider the problem,

$$\min_{x \in \mathcal{X}} \|Wx - b\|_2^2,$$

where $\mathcal{X} = \Delta_n$, the unit simplex, $W \in \mathbb{R}^{m \times d}$ and $b \in \mathbb{R}^m$. Unless otherwise mentioned, we set $\epsilon = 0.1$ and $T_N = 2000$. We generate A randomly with some given condition number κ , in particular a well-conditioned problem with $\kappa(W) = 1$ and an ill-conditioned one with $\kappa(W) = 100$. We let $b_i \sim \mathcal{N}(0, 1)$. We set $m = 100$ and $d = 100$.

Single particle optimization. We begin with a comparison of the convergence speed of GD, MD and SGD and SMD, to show that MD/SMD can outperform GD/SGD, as shown in the theoretical results in Lemmas 9–10 and Propositions 12–13. We observe from Fig. 1 that mirror descent algorithms attain a much faster convergence than gradient descent in the system where the matrix W has a high condition number. For a low condition number the performance of mirror descent is similar to that of gradient descent. Mirror descent seems more robust to a higher noise level, and is able to converge even in a high noise setting. We set $\sigma = 0.05$.

Interacting particle optimization. The theoretical results in Section 3.2 show that the expected value of the distance between the objective evaluated at the time average and the objective evaluated at the optimum is smaller for a larger number of particles. In this section we analyze this numerically and observe that ISMD is indeed able to converge closer to the minimum. Here we consider the linear system optimized with one, 10 and 100 particles. We set $\sigma = 0.05$ and $\eta = 0.01$ for $\kappa(W) = 10$ and $\eta = 0.001$ for $\kappa(W) = 100$. We remark that we thus use fixed learning rates. In Fig. 2 we show the initial convergence speed. In the setting with a high condition number the convergence speed using interaction can be significantly faster than when considering just a single particle.

In Fig. 3 we show the convergence speed as well as the distribution of the loss $f(x_t^i) - f(x^*)$ for the linear system with a condition number 100. We observe from the plots and histograms that the noise in the interacting sampler is smaller so that it able to converge to closer to the global minimum than single-particle optimization. More specifically, the more particles used, then the smaller the variance of the samples is, so that the values are closer to the true minimizer.




Fig. 2. A comparison between the initial convergence of SMD and ISMD for the linear system with condition number 10 (L) and condition number 100 (R). We observe a speedup in convergence using interacting particles.

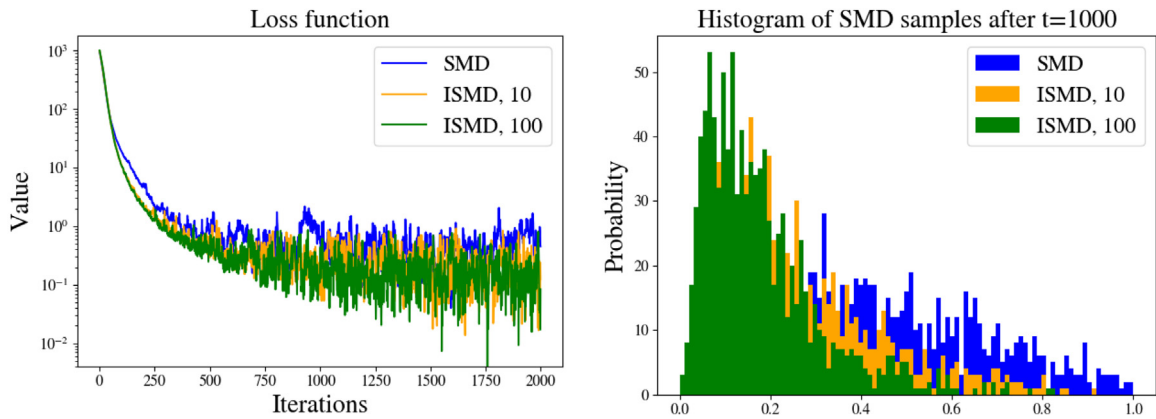


Fig. 3. A comparison between SMD and ISMD for the linear system with condition number 100: (L) the loss function, (R) the histogram which shows the distribution of the losses $f(x_t^i) - f(x^*)$ after convergence ($t > 1000$).

Convergence: vanishing noise against interaction. In this example we study how using a vanishing noise compares to using a fixed learning rate with and without interaction. Often in practice, a vanishing noise is achieved by using a larger batch size, so that the gradient noise is smaller. In Fig. 4 the results for the distance to the optimum, $f(x_t^i) - f(x^*)$, are shown. As usual, the convergence speed using interaction is higher and using interaction can achieve the same effect as using a vanishing noise in terms of the distance to the optimum. We remark here that using an interaction strength, i.e. letting the interaction term be given by $\theta \sum_{i=1}^N A_{ij}(z_t^j - z_t^i)$ can further help in controlling the noise variance. The effects and necessity of such an interaction term will be studied in future work. An additional benefit of having noise is that noise allows to escape local minima in a non-convex problem. Being able to control this noise without decreasing the convergence speed is a significant benefit.

Variance reduction. As implied by the theory in Section 3.2 and the numerical results in the last section, the expected value of the distance between the objective evaluated at the time average and the objective evaluated at the optimum is smaller for a larger number of particles. We keep the learning rate fixed at $\lambda = 0.001$ and set the condition number to 100. To make a fair comparison we sample N i.i.d. copies of the SMD algorithm, and use N particles in the ISMD algorithm. We first plot the variance of $f(x_t^i) - f(x^*)$ after convergence ($t > 1000$) for the i.i.d. an interacting setting in the left-hand side of Fig. 5. Clearly, in the non-interacting setting the variance is not decreased when increasing the number of particles and due to the algorithm not

having converged properly the variance is fluctuating. Even more so, the evolution of the variance of the loss is chaotic since the number of particles does not influence it. In the interacting setting more particles results in a lower loss variance. To understand the effects of interaction and the number of particles on the fluctuation term we compute $\|x_t^i - \bar{x}_t^N\|_2$, whose results are shown in the right-hand side of Fig. 5. As expected from the theoretical analysis the deviation from the mean is smaller in the interacting case than in the i.i.d. case. Furthermore, as N increases the deviation remains constant, validating the theoretical results in Section 3.3 that the fluctuation is bounded and a non-increasing function of N .

5.2. Traffic assignment problem

The objective of the traffic assignment problem is to compute the optimal path between two nodes in a graph. To save space we refer the reader to [2] for a precise description of the problem. We only mention that the problem is a convex optimization problem with a simplex constraint and therefore fits our framework.

Our experiments below are based on a random geometric graph $G(n, r)$ i.e. we placed n points uniformly at random in $[0, 1]^2$ and connect any two points whose distance is at most some w . The weights of the edges in the graph also act as the weight of each edge. We randomly chose two nodes to act as the origin and destination nodes. To set up the optimization problem we also need to compute all the simple paths between the origin and destination node. We denote the length of the maximum path considered in this phase as r_{\max} . The size of r_{\max} decides

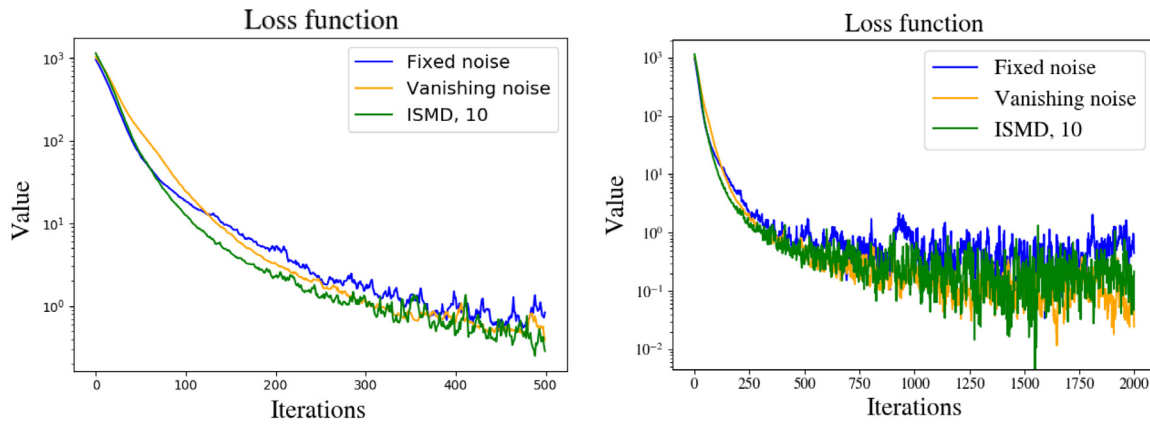


Fig. 4. The initial distance to the optimum (L) and the distance to the optimum after convergence (R) for the linear problem with condition number 100 for SMD with a fixed noise $\sigma = 0.05$, SMD with a vanishing noise $\sigma = \sigma/(t + 1)^{1/10}$ and ISMD with 10 particles.




Fig. 5. The variance of the loss trajectory as a function of the amount of particles (L) and the mean of the L_2 -norm of the fluctuation term $\|x_t^i - \bar{x}_t^N\|_2$ for the i.i.d. and interacting case. We see that in the interacting case increasing the number of particles decreases the variance. The deviation from the mean is smaller in the interacting case, as expected.

the dimensionality of the problem. For the stochastic version of the problem we added Gaussian $N(0, \sigma_g)$ noise on each edge. We used a constant step-size for all algorithms. For mirror descent (all variations) we used $\eta = 0.2$ and for gradient descent we manually selected the best step-size for the problem (see below for additional remarks regarding GD for this problem).

Comparison between MD and GD. An interesting difference between the regression problem above and the traffic assignment problem is that the solution of the latter is sparse. As a result the solution is on the boundary with many paths having zero load. In this case the gradient descent algorithm performed considerably worse than all variations of mirror descent (i.e. with many or single particles) and in all variations of the problem we considered (small/large problems, with and without noise). In order to illustrate the difference in performance between mirror descent and gradient descent we plot the results for one example in Fig. 6. The graph used to generate the problem above has 100 nodes, and the resulting optimization problem has 5831 possible paths i.e. $d = 5831$. We begin with an illustration of the performance of standard MD against GD in Fig. 6; GD is significantly slower to converge than MD. Given the large difference in performance between MD and GD we will not report any more results from GD for the traffic assignment problem.

Comparison between SMD and ISMD. Similar to the experiments run for the linear regression we compare the iterations of the single particle mirror descent i.e. (SMD) and its interacting variant

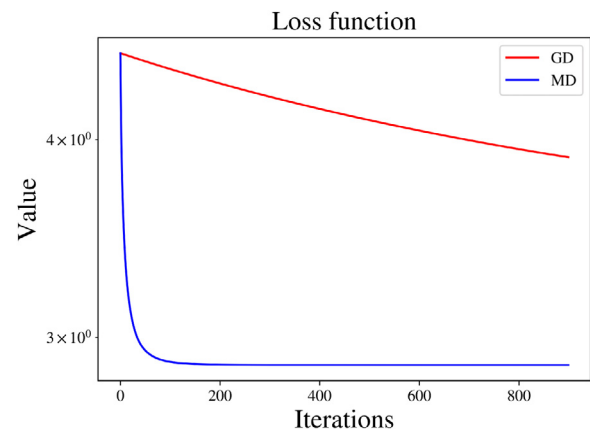


Fig. 6. A comparison between MD and GD. GD with Euclidean projections is considerably slower than MD for this problem because the solution is sparse. IMD performs the same as MD for deterministic problems.

(ISMD) with a different number of particles ($N = 10, 100, 1000$). In the left of Fig. 7 we show the iteration history of (SMD) and (ISMD) with $N = 100$ particles. To visualize the differences between (SMD) and (ISMD) and the impact of the number of particles we plot a histogram in the right plot in Fig. 7 of the iterations after convergence (typically $t > 100$). Clearly there




Fig. 7. Comparison between SMD and ISMD for a single path (L) and a histogram of the samples for the comparison between SMD and ISMD with 10, 100 and 1000 particles (R). With more particles the variance of the samples is lower.

Table 1

A comparison of the variance of N i.i.d. copies of SMD and N particles in ISMD. We obtain a variance reduction using interaction.

Nodes	r_{\max}	d	N	$\frac{\sigma_{\text{SMD}}^2 - \sigma_{\text{ISMD}}^2}{\sigma_{\text{SMD}}^2}$
50	5	70	10	0.18
50	5	70	50	0.21
50	5	70	100	0.09

is a considerable reduction in the variance of the iterates between (MD) and (ISMD). We attempt to quantify the reduction in variance in the next set of experiments.

In order to further quantify the variance reduction between ISMD and SMD in a fair way it is important to compare the reduction in variance between N i.i.d. copies of SMD and ISMD with N particles. In Table 1 we present the variance reduction results for different particles and observe that interaction can significantly reduce the variance.

5.3. Mini-batch optimization

In machine learning the optimization objective typically consists of a sum over data samples, i.e. $f(x) = \sum_{i=1}^m f_i(x)$, where m is the sample size. The gradient is typically computed over a subset of the data (a mini-batch), since for large m computing the full gradient is too costly. The downside of computing the gradient over a mini-batch is that the gradient contains noise. In our work we proposed interaction as a way of reducing the effects of this gradient noise. In this section we show how interaction helps in reducing the effect of noise in mini-batch gradient descent, and consider in more detail the computational costs related to interaction against an increased batch size. The ideas presented here can be generalized to popular machine learning techniques such as neural networks.

Consider a similar setup as in Section 5.1,

$$f(x) = \frac{1}{m} \sum_{i=1}^m f_i(x) := \frac{1}{m} \sum_{i=1}^m \|W_i \cdot x - b_i\|_2^2.$$

In every iteration, the gradient is computed over a subset of the data, $f_S(x) = \frac{1}{|S|} \sum_{i \in S} f_i(x)$, where $|S|$ refers to the size of the batch. The optimization algorithm is then given by,

$$z_{t+1}^i = z_t^i - \eta \epsilon \nabla f_S(x_t^i) + \epsilon \sum_{j=1}^N A_{ij} (z_t^j - z_t^i).$$

The noise is thus implicit in the gradient ∇f_S .

Convergence speed. In Fig. 8 the convergence results are presented for different batch sizes $|S|$ with and without interaction. We set $m = 200$ and $d = 100$, and use $\kappa(W) = 10$ with $\eta = 0.1$ and $\kappa(W) = 200$ with $\eta = 0.1$. Using N particles infers a computational cost of N times that of a single particle. However, as observed in the plot, using N particles allows to use a significantly smaller batch, i.e. due to the interaction the optimization can converge even with a higher noise. We conclude that interaction can be a convenient methodology for improving the convergence of the algorithm when using stochastic gradients.

Changing the interaction matrix. In this section we study the trade-offs between interaction matrix sparsity and convergence time. We use a batch size of 10 and 10 nodes, and $\kappa(W) = 200$ with $\eta = 0.01$. Table 2 shows the results for a Erdős-Rényi communication graph, a random graph where each edge is chosen with a certain probability. A connectivity probability of p means that on average each node is connected to $p \times 10$ other nodes. The communication between the nodes is thus determined by the connectivity probability. The weight matrix A is determined by a doubly stochastic version of this communication graph. From the results we see, as expected, if more communication is present the convergence is faster at the cost of more communication. Interestingly, the total communication is lowest for the most sparse graph, showing that fast convergence can even be achieved with minimal communication. We do remark that this depends on the particular setup of the problem: in a same f setting the convergence can be equally fast for sparse and dense communication (not shown here) while for highly different local functions f_i more communication may be needed to achieve convergence. In Fig. 9 we plot the convergence speed for the different interaction matrices. In correspondence with the results in Table 2 we observe a faster initial convergence for the densely connected matrix. Interestingly, even with sparse communications a relatively fast convergence can be obtained, showing the potential of our methods.

6. Conclusion

In this work we analyzed stochastic mirror descent with interacting particles with the aim of understanding the tradeoffs between computation, communication and variance reduction. With a fixed learning rate or a constant noise variance convergence to the exact minimizer is not possible. Prior results showed that decreasing the learning rate [43] or using a vanishing noise variance [2] can result in good convergence properties of stochastic optimization. In our work we showed that using N

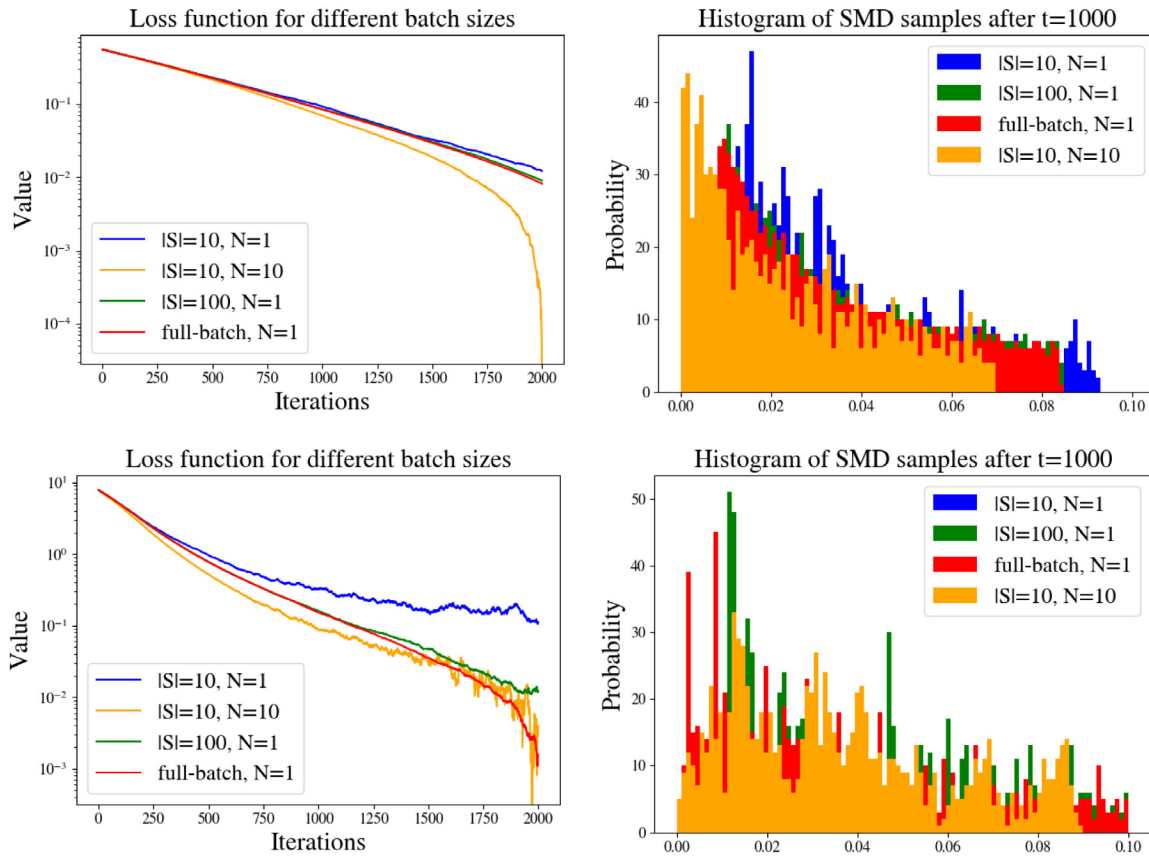


Fig. 8. The loss function (L) and the histogram (R) for ISMD with different batch sizes and different particles for $\kappa(W) = 10$ (top) and $\kappa(W) = 200$ (bottom). The loss function is computed as a sum over samples, and the gradient is computed over a batch of samples. Using interacting particles allows to use a smaller batch size while still attaining convergence. The presented results are averaged over 10 runs.




Fig. 9. The loss function for initial convergence (L) and full convergence (R) for different interaction matrices. Interestingly, even with sparse communication a fast convergence can be achieved.

interacting particles can in a similar fashion decrease the effect of noise and result in a closer convergence to the optimum. In addition using interaction decreases the variance of the samples further than that of N independent replicas of SMD. We therefore argue that interaction between N particles is a viable method and can control the convergence of the algorithm.

We analyzed the convergence from an optimization perspective exploiting Lyapunov functions. We showed that the deviation of an individual particle from the particle mean is bounded and showed that ISMD achieves a linear convergence rate in the convex case and an exponential one in the strongly convex setting. We furthermore presented the relationship between sampling

from the invariant measure and converging to the optimizer of the objective function. Using log-Sobolev inequalities we obtained an explicit convergence rate to the invariant measure, implying that the particles converge to a neighborhood of the optimizer at an exponential rate. Our numerical results show that interaction is beneficial in cases where the solution is sparse and therefore the behavior of the algorithm at the boundary is important. This is a case where the role of the mirror map becomes more critical, which is also reflected in our assumptions.

Interesting extensions include the study of SMD in the non-convex setting where the noise can help escape from local minima (see e.g. [44,45]). In this case the interaction can play the role

Table 2

A comparison of the communication costs and convergence time for different interaction graphs. The results are averaged over 10 runs of the algorithm (both the generation of the graph as the optimization). Time to convergence is determined as the first time the loss value hits below a certain level (here 43.02). The total communication for convergence is determined as the number of iterations needed to converge times the average communication per round.

\mathbb{P} (connectivity)	0.3	0.5	1
Time to convergence	883	712	347
Total communication for convergence	2649	3560	3470
Loss at $t = 2000$	42.93	42.83	41.93
Consensus error at $t = 2000$	0.013	0.012	0.011

of a regularizer due to the convexification of the loss surface [46]. Additionally, even in a non-convex setting using interaction enforces the particles to converge to consensus in a local minimum. The interaction can also act as a pre-conditioner and its statistical properties may help speed up the convergence or minimize the variance. An open question is the study of the optimal interaction matrix beyond just mixing times [47]. Alternatively of interest is to study the choice of the interaction such that communication is minimal, but variance reduction and consensus is still guaranteed.

It is also interesting to study such algorithms in a truly distributed setting, i.e. when each particle has access to a subset of the data. This is the case in learning with privacy constraints or federated learning approaches. Some preliminary results on this have been presented in [48]. In this setting also the effect of an interaction strength will be of interest and we aim to address the necessity of considering an additional interaction strength θ , i.e. using $\theta \sum_{i=1}^N A_{ij}(z_t^j - z_t^i)$, in future work. More specifically, in a distributed setting the particles optimize different objectives, and – unlike in our setting of a same objective – convergence to consensus cannot be guaranteed without imposing a sufficiently high interaction strength.

Lastly, the effect of time discretization was not considered here. This is a well-understood topic. We refer the interested reader to [31,49] for details related to the Euler discretization and [12] for more advanced schemes which also include acceleration.

CRedit authorship contribution statement

A. Borovykh: Conceptualization, Methodology, Writing, Software, Formal analysis. **N. Kantas:** Conceptualization, Methodology, Writing, Formal analysis, Supervision. **P. Parpas:** Conceptualization, Methodology, Writing, Software, Supervision. **G.A. Pavliotis:** Conceptualization, Methodology, Writing, Supervision.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This project was funded by JPMorgan Chase & Co, United Kingdom under a J.P. Morgan A.I. Research Award 2019. Any views or opinions expressed herein are solely those of the authors listed, and may differ from the views and opinions expressed by JPMorgan Chase & Co. or its affiliates. This material is not a product of the Research Department of J.P. Morgan Securities LLC. This material does not constitute a solicitation or offer in any jurisdiction. G.A.P. was partially supported by the EPSRC, United Kingdom through the Grant Number EP/P031587/1.

Appendix A. Auxiliary results for Section 3

Here we prove the Lemmas that were used in Section 3. In Proposition 15 we used the following lemma.

Lemma 24. *Under the continuous stochastic mirror descent dynamics (8) it holds,*

$$\int_0^T \frac{1}{N} \sum_{i=1}^N (y_t^N - x^*)^T \nabla f(x_t^i) dt \leq \frac{1}{2} D_{\Phi, x}^2 + \frac{\sigma^2 T}{2N} \Delta \Phi^*(\bar{z}_t^N) + \int_0^T \frac{\sigma}{\sqrt{N}} \|y_t^N - x^*\|_2 dB_t^i.$$

Proof. Consider the Lyapunov function defined as

$$V(\bar{z}_t^N) = \Phi^*(\bar{z}_t^N) - \Phi^*(z^*) - (x^*)^T (\bar{z}_t^N - z^*).$$

Applying Itô's lemma and using the fact that $\sigma(y_t^N - x^*)^T \frac{1}{N} \sum_{i=1}^N dB_t^i \stackrel{d}{=} \frac{\sigma}{\sqrt{N}} \|y_t^N - x^*\|_2 dB_t$ we obtain,

$$\begin{aligned} dV(\bar{z}_t^N) &= d\Phi^*(\bar{z}_t^N) - (x^*)^T d\bar{z}_t^N \\ &= (\nabla \Phi^*(z^*) - \nabla \Phi^*(\bar{z}_t^N))^T \frac{1}{N} \sum_{i=1}^N \nabla f(x_t^i) dt \\ &\quad + \frac{\sigma^2}{2N} \Delta \Phi^*(\bar{z}_t^N) dt + \frac{\sigma}{\sqrt{N}} \|y_t^N - x^*\|_2 dB_t. \end{aligned}$$

Integrating, and using the standard bounds we obtain,

$$\int_0^T \frac{1}{N} (y_t^N - x^*)^T \sum_{i=1}^N \nabla f(x_t^i) dt \leq \frac{1}{2} D_{\Phi, x}^2 + \frac{\sigma^2 T}{2N} \Delta \Phi^*(\bar{z}_t^N) + \int_0^T \frac{\sigma}{\sqrt{N}} \|y_t^N - x^*\|_2 dB_t. \quad \square$$

Similar to the convex case, in Proposition 16 we used the following lemma.

Lemma 25. *Under the continuous stochastic mirror descent dynamics (8) it holds,*

$$\begin{aligned} \int_0^T e^{\mu t} \left(\frac{1}{N} (y_t^N - x^*)^T \sum_{i=1}^N \nabla f(x_t^i) - \mu D_{\Phi^*}(\bar{z}_t^N, z^*) \right) dt \\ \leq \frac{1}{2} D_{\Phi, x}^2 + \frac{\sigma^2}{2N} (e^{\mu T} - 1) \Delta \Phi^*(\bar{z}_t^N) \\ + \int_0^T e^{\mu t} \frac{\sigma}{\sqrt{N}} \|y_t^N - x^*\|_2 dW_t. \end{aligned}$$

Proof. Consider the Lyapunov function defined as $\bar{V}(t, \bar{z}_t^N) = e^{\mu t} D_{\Phi^*}(\bar{z}_t^N, z^*)$. Applying Itô's lemma we obtain,

$$\begin{aligned} dV(t, \bar{z}_t^N) &= e^{\mu t} (x^* - y_t^N)^T \frac{1}{N} \sum_{i=1}^N \nabla f(x_t^i) dt \\ &\quad + \mu e^{\mu t} D_{\Phi^*}(\bar{z}_t^N, z^*) dt + \frac{\sigma^2}{2N} e^{\mu t} \Delta \Phi^*(\bar{z}_t^N) dt \\ &\quad + \frac{\sigma}{\sqrt{N}} e^{\mu t} \|y_t^N - x^*\|_2 dW_t, \end{aligned}$$

where we have used the evolution of \bar{z}_t^N . Integrating and using the standard bounds we obtain the statement. \square

Appendix B. Auxiliary results for Section 4

B.1. A primer on log Sobolev inequalities and Markov semigroups

Recall, we defined in Section 1.2.1 that a probability measure ν satisfies a Log-Sobolev inequality with constant C is for any smooth function f we have,

$$\text{Ent}_\nu(f^2) \leq C\nu(|\nabla f|^2), \tag{24}$$

where the entropy was defined as,

$$\text{Ent}_\nu(f^2) = \nu(f^2 \log f^2) - \nu(f^2) \log(\nu(f^2)).$$

Satisfying (24) has many important consequences summarized below.

Theorem 26. Let $\nu \in \mathcal{P}(\mathbb{R}^d)$ satisfy a log Sobolev inequality with constant C . Then we have the following:

1. Exponential integrability of Lipschitz functions (Herbst argument): Let f be a 1-Lipschitz function with $\|f\|_{Lip} \leq 1$, then for every $s \in \mathbb{R}$

$$\nu(e^{sf}) \leq e^{s\nu(f) + Cs^2}$$

$$\text{and } \nu(e^{r^2 f}) < \infty \text{ for } r^2 < C^{-1};$$

2. A concentration of measure result:

$$\nu(|f - \mathbb{E}_\nu[f]| \geq r) \leq 2 \exp\left(-\frac{r^2}{C \|f\|_{Lip}^2}\right)$$

for any Lipschitz f ;

3. The Wasserstein distance is bounded by relative entropy: $W_2(\nu, \tilde{\nu})^2 \leq \frac{C}{2} H(\nu|\tilde{\nu})$ for $\nu \ll \tilde{\nu}$;

Proof. We provide references for each point: 1. see Proposition 5.4.1 [29]; 2. see [50]; 3. see [51]. \square

Note point 1. in Theorem 26 implies finite moments for Lipschitz f and $p \geq 2$:

$$\nu(|f|^p)^2 \leq \nu(|f|^2)^2 + 2C(p-2) \|f\|_{Lip}^2,$$

see Proposition 5.4.2 in [29] for details.

Proving Theorem 26 uses properties of certain Markov semigroups and the corresponding Fokker Planck equations. The converse is also possible. One can establish the log-Sobolev inequalities for the law of the process one can employ the classical results from [28]. We briefly summarize some useful results for an SDE of the form,

$$dz_t = -\nabla \nu(z_t) dt + \sqrt{2} dB_t, \quad z_0 = z \tag{25}$$

Let $\nu_t = Law(z_t)$. In order to demonstrate ergodicity and exponential convergence to equilibrium, one approach is to require existence of a (L^2) spectral gap. A common requirement $\lim_{|z| \rightarrow \infty} \left(\frac{|\nabla \nu|^2}{2} - \Delta \nu\right) = \infty$ and this ensures that ν_t satisfies a Poincare inequality and the SDE is ergodic with invariant density

$$\nu_\infty(dz) := \frac{1}{Z} \exp(\nu(z)) dz,$$

see Proposition 4.2 and Theorem 4.3 in [36] for more details. One can strengthen these results with log-Sobolev inequalities from [28]:

Theorem 27. Assume

$$\text{Hess} \nu \geq \rho I_d, \quad \rho \in \mathbb{R} \tag{26}$$

Then ν_t satisfies (1) with constant

$$C_t = \frac{2}{\rho} (1 - e^{-\rho t}).$$

If in addition $\rho > 0$, we get,

$$H(\nu_t | \nu_\infty) \leq Ke^{-2\rho t}.$$

Note that requiring ν being strongly convex is crucial to obtain explicit rates of convergence to equilibrium. This is a particular instance of the celebrated (and more general) Bakry–Emery criterion on the curvature of Markov semigroups. To state the more general version for any Markov semigroup we need to define the following differential operators defined on appropriate domains:

$$\mathcal{L}(f) = -\nabla \nu^T \nabla f + \frac{\sigma^2}{2} \text{Hess}(f),$$

$$\Gamma(f) = \frac{1}{2} (\mathcal{L}(fg) - f \mathcal{L}(g) - g \mathcal{L}(f)),$$

$$\Gamma_2(f) = \frac{1}{2} (\mathcal{L}(\Gamma(f, f)) - 2\Gamma(\mathcal{L}(f), f)).$$

The Bakry–Emery criterion consists of verifying,

$$\Gamma_2(f) \geq \rho \Gamma(f, f),$$

for any smooth f , which is equivalent to (26) for the dynamics in (25).

In addition, point 3. in Theorem 26 also implies a convergence in W_2 , i.e.

$$W_2(\nu_t, \nu_\infty) \leq \sqrt{\frac{K}{2}} e^{-\rho t}.$$

and by Pinsker’s inequality we also have convergence in total variation with same rate:

$$\|\nu_t - \nu_\infty\|_{TV} \leq \sqrt{\frac{K}{2}} e^{-\rho t},$$

where we denote $\|\nu\|_{TV} = \sup_{|\varphi| < 1} |\nu(\varphi)|$, for any ν -measurable function φ .

B.2. Empirical risk bounds for Proposition 23

The main purpose of this section is to provide so called empirical risk bounds for $\mathbb{E}_{\eta_\infty^N}(\mathcal{W}) - \mathcal{W}(z^*)$. We will base our derivation on Proposition 3.4 in [42] and specifically focus on certain aspects of the particle system with invariant distribution

$$\eta_\infty^N(dz) = \frac{1}{Z^N} \exp\left(-\frac{2}{\sigma^2} \mathcal{W}(z)\right) dz.$$

We will define the differential entropy for a probability measure $\eta \in \mathcal{P}_2(\mathbb{R}^d)$ as $\mathcal{H}(\eta) = -\int \log\left(\frac{d\eta}{dz}\right)(z) \eta(dz)$. From the definition of \mathcal{H} we have for η_∞^N

$$\mathbb{E}_{\eta_\infty^N} \mathcal{W} = \frac{\sigma^2}{2} (\mathcal{H}(\eta_\infty^N) - \log Z^N).$$

In what follows we will provide upper bounds for $\mathcal{H}(\eta_\infty^N)$ and lower bounds for $\log Z^N$ to reach an empirical risk bound.

Lemma 28. Let $\eta \in \mathcal{P}_2(\mathbb{R}^d)$ and let Σ denote the covariance matrix of η . Then $\mathcal{H}(\eta) \leq \frac{d}{2} (\log 2\pi + 1) + \frac{1}{2} \log \det \Sigma$.

Proof. Consider the Kullback Leibler divergence between η and a Gaussian distribution with the same mean and variance μ, Σ :

$$H(\eta | \mathcal{N}(\mu, \Sigma)) = -\mathcal{H}(\eta) - \mathbb{E}_\eta [\log \phi(z; \mu, \Sigma)],$$

where

$$\mathbb{E}_\eta [\log \phi(z; \mu, \Sigma)] = -\frac{d}{2} \log 2\pi - \frac{1}{2} \log \det(\Sigma) - \frac{1}{2} \mathbb{E}_\eta [(z - \mu)^T \Sigma^{-1} (z - \mu)].$$

Straightforward manipulations show

$$\begin{aligned} \mathbb{E}_\eta [(z - \mu)^T \Sigma^{-1} (z - \mu)] &= \mathbb{E}_\eta [z^T \Sigma^{-1} z] - \mu^T \Sigma^{-1} \mu \\ &= \text{Tr}[\Sigma^{-1} \Sigma] + \mu^T \Sigma^{-1} \mu - \mu^T \Sigma^{-1} \mu \\ &= d \end{aligned}$$

and hence

$$H(\eta | \mathcal{N}(\mu, \Sigma)) = -\mathcal{H}(\eta) + \frac{d}{2} \log 2\pi + \frac{1}{2} \log \det(\Sigma) + \frac{d}{2}.$$

Given $H(\eta | \mathcal{N}(\mu, \Sigma)) \geq 0$ we get the required result. \square

Corollary 29. Let $\eta_\infty^N \in \mathcal{P}_2(\mathbb{R}^{dN})$ and Σ^N denote the covariance matrix of η_∞^N . Then for the particle system it holds that

$$\mathcal{H}(\eta_\infty^N) \leq \frac{dN}{2} (\log 2\pi + 1) + \frac{1}{2} \log \det(\Sigma^N).$$

Next we provide a lower bound for $\log Z^N$.

Proposition 30. Let Assumptions 1, 2 hold w.r.t the Euclidean norm. We have that $\nabla \mathcal{W}$ is L_N -Lipschitz and as a result

$$\log Z^N \geq -\frac{2}{\sigma^2} \mathcal{W}^* + \frac{1}{2} \log \left(\frac{\sigma^2}{2L_N} \right) + \frac{dN}{2} \log 2\pi.$$

Proof. The first claim is trivial and one has $L_N = \frac{L}{\mu} + \|\mathcal{L}\|$ where $\|\mathcal{L}\|$ is any matrix norm on \mathcal{L} . Smoothness of $\nabla \mathcal{W}$ implies the following quadratic upper bounds:

$$\mathcal{W}(\mathbf{z}) \leq \mathcal{W}(\mathbf{z}^*) + \frac{L_N}{2} \|\mathbf{z} - \mathbf{z}^*\|^2.$$

Then we have,

$$\begin{aligned} \log Z^N &= -\frac{2}{\sigma^2} \mathcal{W}(\mathbf{z}^*) + \log \int \exp \left(-\frac{2}{\sigma^2} (\mathcal{W}(\mathbf{z}) - \mathcal{W}(\mathbf{z}^*)) \right) d\mathbf{z} \\ &\geq -\frac{2}{\sigma^2} \mathcal{W}(\mathbf{z}^*) + \log \int \exp \left(-\frac{L_N}{\sigma^2} \|\mathbf{z} - \mathbf{z}^*\|^2 \right) d\mathbf{z} \\ &= -\frac{2}{\sigma^2} \mathcal{W}(\mathbf{z}^*) + \frac{1}{2} \log \left(\frac{\sigma^2}{2L_N} \right) + \frac{dN}{2} \log 2\pi, \end{aligned}$$

where in the last step we used the Gaussian integral $\int \exp \left(-\frac{1}{2} \mathbf{z}^T A \mathbf{z} + b^T \mathbf{z} \right) d\mathbf{z} = \sqrt{\frac{(2\pi)^{dN}}{\det A}} \exp \left(\frac{1}{2} b^T A^{-1} b \right)$ with $A = \frac{2L_N}{\sigma^2} I_{dN}$ and $b = A\mathbf{z}^*$. \square

Remark 31. When non-Euclidean norms are used $L_N = \frac{L}{\mu} K + \|\mathcal{L}\|$, with K arising from norm equivalence.

References

[1] R.M. Gower, N. Loizou, X. Qian, A. Sailanbayev, E. Shulgin, P. Richtárik, SGD: General analysis and improved rates, 2019, arXiv preprint arXiv:1901.09401.
 [2] P. Mertikopoulos, M. Staudigl, On the convergence of gradient-like flows with noisy gradient input, SIAM J. Optim. 28 (1) (2018) 163–197.
 [3] P. Mertikopoulos, M. Staudigl, Stochastic mirror descent dynamics and their convergence in monotone variational inequalities, J. Optim. Theory Appl. 179 (3) (2018) 838–867.
 [4] S.L. Smith, P.-J. Kindermans, C. Ying, Q.V. Le, Don't decay the learning rate, increase the batch size, 2017, arXiv preprint arXiv:1711.00489.
 [5] A.S. Nemirovsky, D.B. Yudin, Problem complexity and method efficiency in optimization, 1983.
 [6] A. Beck, First-Order Methods in Optimization, Vol. 25, SIAM, 2017.

[7] D. Needell, R. Ward, N. Srebro, Stochastic gradient descent, weighted sampling, and the randomized Kaczmarz algorithm, in: Advances in Neural Information Processing Systems, 2014, pp. 1017–1025.
 [8] R. Johnson, T. Zhang, Accelerating stochastic gradient descent using predictive variance reduction, in: Advances in Neural Information Processing Systems, 2013, pp. 315–323.
 [9] A. Defazio, F. Bach, S. Lacoste-Julien, SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives, in: Advances in Neural Information Processing Systems, 2014, pp. 1646–1654.
 [10] E. Gorbunov, F. Hanzely, P. Richtárik, A unified theory of SGD: Variance reduction, sampling, quantization and coordinate descent, 2019, arXiv preprint arXiv:1905.11261.
 [11] W. Krichene, P.L. Bartlett, Acceleration and averaging in stochastic mirror descent dynamics, 2017, arXiv preprint arXiv:1707.06219.
 [12] P. Xu, T. Wang, Q. Gu, Accelerated stochastic mirror descent: From continuous-time dynamics to discrete-time algorithms, in: International Conference on Artificial Intelligence and Statistics (AISTATS), 2018.
 [13] M. Raginsky, J. Brourie, Continuous-time stochastic mirror descent on a network: Variance reduction, consensus, convergence, in: 2012 IEEE 51st IEEE Conference on Decision and Control, CDC, IEEE, 2012, pp. 6793–6800.
 [14] J.C. Duchi, A. Agarwal, M.J. Wainwright, Dual averaging for distributed optimization: Convergence analysis and network scaling, IEEE Trans. Autom. Control 57 (3) (2011) 592–606.
 [15] P. Lin, W. Ren, J.A. Farrell, Distributed continuous-time optimization: nonuniform gradient gains, finite-time convergence, and convex constraint set, IEEE Trans. Automat. Control 62 (5) (2016) 2239–2253.
 [16] S. Shahrampour, A. Jadbabaie, Distributed online optimization in dynamic environments using mirror descent, IEEE Trans. Automat. Control 63 (3) (2017) 714–725.
 [17] A. Koloskova, S.U. Stich, M. Jaggi, Decentralized stochastic optimization and gossip algorithms with compressed communication, 2019, arXiv preprint arXiv:1902.00340.
 [18] K. Seaman, F. Bach, S. Bubeck, Y.T. Lee, L. Massoulié, Optimal algorithms for smooth and strongly convex distributed optimization in networks, in: Proceedings of the 34th International Conference on Machine Learning, Vol. 70, JMLR.org, 2017, pp. 3027–3036.
 [19] V.N. Belykh, I.V. Belykh, M. Hasler, Connection graph stability method for synchronized coupled chaotic systems, Physica D: Nonlinear Phenom. 195 (1–2) (2004) 159–187.
 [20] G. Shi, A. Proutiere, K.H. Johansson, Network synchronization with convexity, SIAM J. Control Optim. 53 (6) (2015) 3562–3583.
 [21] W. Yu, G. Chen, M. Cao, Consensus in directed networks of agents with nonlinear dynamics, IEEE Trans. Automat. Control 56 (6) (2011) 1436–1441.
 [22] W. Yu, J. Cao, G. Chen, J. Lu, J. Han, W. Wei, Local synchronization of a complex network model, IEEE Trans. Syst. Man Cybern. B 39 (1) (2008) 230–241.
 [23] W. Yu, J. Cao, J. Lü, Global synchronization of linearly hybrid coupled networks with time-varying delay, SIAM J. Appl. Dyn. Syst. 7 (1) (2008) 108–133.
 [24] Y.-P. Hsieh, A. Kavis, P. Rolland, V. Cevher, Mirrored Langevin dynamics, in: Advances in Neural Information Processing Systems, 2018, pp. 2878–2887.
 [25] J. Baker, P. Fearnhead, E. Fox, C. Nemeth, Large-scale stochastic sampling from the probability simplex, in: Advances in Neural Information Processing Systems, 2018, pp. 6721–6731.
 [26] S. Yokoi, T. Otsuka, I. Sato, On transformations in stochastic gradient MCMC, 2019, arXiv preprint arXiv:1903.02750.
 [27] Y.-A. Ma, Y. Chen, C. Jin, N. Flammario, M.I. Jordan, Sampling can be faster than optimization, Proc. Natl. Acad. Sci. 116 (42) (2019) 20881–20885.
 [28] D. Bakry, On Sobolev and logarithmic Sobolev inequalities for Markov semigroups, New Trends Stoch. Anal. (1997) 43–75, (Charingworth, 1994).
 [29] D. Bakry, I. Gentil, M. Ledoux, Analysis and Geometry of Markov Diffusion Operators, Vol. 348, Springer Science & Business Media, 2013.
 [30] F. Malrieu, Logarithmic Sobolev inequalities for some nonlinear PDE's, Stoch. Process. Appl. 95 (1) (2001) 109–132.
 [31] A. Veretennikov, On ergodic measures for McKean–Vlasov stochastic equations, in: Monte Carlo and Quasi-Monte Carlo Methods 2004, Springer, 2006, pp. 471–486.
 [32] Y. Nesterov, Lectures on Convex Optimization, Vol. 137, Springer, 2018.
 [33] S. Bubeck, Convex optimization: Algorithms and complexity, 2014, arXiv preprint arXiv:1405.4980.
 [34] R. Pettersson, Projection scheme for stochastic differential equations with convex constraints, Stochastic Process. Appl. 88 (1) (2000) 125–134.
 [35] A. Storm, Stochastic differential equations with a convex constraint, Stoch. Stoch. Rep. 53 (3–4) (1995) 241–274.
 [36] G.A. Pavliotis, Stochastic Processes and Applications: Diffusion Processes, the Fokker-Planck and Langevin Equations, Vol. 60, Springer, 2014.
 [37] A. Wilson, Lyapunov Arguments in Optimization, University of California, Berkeley, 2018.
 [38] Q. Li, C. Tai, E. Weinan, Stochastic modified equations and adaptive stochastic gradient algorithms, in: International Conference on Machine Learning, 2017, pp. 2101–2110.

- [39] Q. Li, C. Tai, E. Weinan, Stochastic modified equations and dynamics of stochastic gradient algorithms I: Mathematical foundations., *J. Mach. Learn. Res.* 20 (2019) 40–1.
- [40] S. Mandt, M.D. Hoffman, D.M. Blei, Stochastic gradient descent as approximate Bayesian inference, *J. Mach. Learn. Res.* 18 (1) (2017) 4873–4907.
- [41] M. Mesbahi, M. Egerstedt, *Graph Theoretic Methods in Multiagent Networks*, Princeton University Press, 2010.
- [42] M. Raginsky, A. Rakhlin, M. Telgarsky, Non-convex learning via stochastic gradient Langevin dynamics: A nonasymptotic analysis, in: *Conference on Learning Theory*, 2017, pp. 1674–1703.
- [43] B. Shi, W.J. Su, M.I. Jordan, On learning rates and Schrödinger operators, 2020, arXiv preprint [arXiv:2004.06977](https://arxiv.org/abs/2004.06977).
- [44] Z. Zhu, J. Wu, B. Yu, L. Wu, J. Ma, The anisotropic noise in stochastic gradient descent: Its behavior of escaping from minima and regularization effects, 2018, arXiv preprint [arXiv:1803.00195](https://arxiv.org/abs/1803.00195).
- [45] B. Helffer, Remarks on decay of correlations and Witten Laplacians Brascamp–Lieb inequalities and semiclassical limit, *J. Funct. Anal.* 155 (2) (1998) 571–586.
- [46] S. Zhang, A.E. Choromanska, Y. LeCun, Deep learning with elastic averaging SGD, in: *Advances in Neural Information Processing Systems*, 2015, pp. 685–693.
- [47] S. Boyd, P. Diaconis, L. Xiao, Fastest mixing Markov chain on a graph, *SIAM Rev.* 46 (4) (2004) 667–689.
- [48] A. Borovykh, P. Parpas, N. Kantas, G. Pavliotis, To interact or not? The convergence properties of interacting stochastic mirror descent, in: *International Conference on Machine Learning (ICML) Workshop on ‘Beyond First order methods in ML Systems’*, 2020.
- [49] F. Malrieu, D. Talay, Concentration inequalities for Euler schemes, in: *Monte Carlo and Quasi-Monte Carlo Methods 2004*, Springer, 2006, pp. 355–371.
- [50] M. Ledoux, Concentration of measure and logarithmic Sobolev inequalities, in: *Seminaire de Probabilites XXXIII*, Springer, 1999, pp. 120–216.
- [51] F. Otto, C. Villani, Generalization of an inequality by Talagrand and links with the logarithmic Sobolev inequality, *J. Funct. Anal.* 173 (2) (2000) 361–400.