

Distributed Mirror Descent With Integral Feedback: Asymptotic Convergence Analysis of Continuous-Time Dynamics

Youbang Sun^{ID} and Shahin Shahrampour^{ID}, *Senior Member, IEEE*

Abstract—This letter addresses distributed optimization, where a network of agents wants to minimize a global strongly convex objective function. The global function can be written as a sum of local convex functions, each of which is associated with an agent. We propose a continuous-time distributed mirror descent algorithm that uses purely local information to converge to the global optimum. Unlike previous work on distributed mirror descent, we incorporate an integral feedback in the update, allowing the algorithm to converge with a constant step-size when discretized. We establish the asymptotic convergence of the algorithm using Lyapunov stability analysis. We further illustrate numerical experiments that verify the advantage of adopting integral feedback for improving the convergence rate of distributed mirror descent.

Index Terms—Optimization algorithms, decentralized control, Lyapunov methods, mirror descent.

I. INTRODUCTION

THE MIRROR descent (MD) algorithm [1] is a primal-dual method that has been successfully used for large-scale convex optimization problems. MD can be seen as a generalization of gradient descent, which can exploit the geometry of the optimization problem. The algorithm replaces the Euclidean distance with a so-called *Bregman divergence* as the regularizer for projection. This idea provides a significant convergence speed-up for high-dimensional optimization problems [2].

In practice, optimization methods (including MD) are numerically implemented in discrete time, but their continuous-time analysis has always been of major interest to the control and optimization community [3]–[5]. This stems from the fact that many optimization methods can be interpreted as discretization of ordinary differential equations (ODEs), and therefore, their convergence can be established using the theory of control and dynamical systems. The MD

Manuscript received September 14, 2020; revised November 1, 2020; accepted November 17, 2020. Date of publication November 26, 2020; date of current version December 22, 2020. The work of Shahin Shahrampour was supported by NSF under Grant ECCS-1933878. Recommended by Senior Editor J.-F. Zhang. (Corresponding author: Shahin Shahrampour.)

The authors are with the Wm Michael Barnes '64 Department of Industrial and Systems Engineering, Texas A&M University, College Station, TX 77843 USA (e-mail: ybsun@tamu.edu; shahin@tamu.edu).

Digital Object Identifier 10.1109/LCSYS.2020.3040934

algorithm is no exception in this regard, and it can be studied via a system of ODEs [6].

In this letter, we address distributed continuous-time optimization via decentralized mirror descent, inspired by the success of centralized MD in large-scale optimization. In this setup, a network of agents wants to minimize a global strongly convex objective function. The global function can be written as a sum of local convex functions, each of which is associated with an agent. We develop a continuous-time decentralized MD algorithm that uses purely local gradient information to converge to the global minimizer. Contrary to the prior work on (discrete) distributed mirror descent (e.g., [7], [8]), we enforce consensus among agents using the idea of integral feedback, in addition to the standard neighborhood averaging. The integral feedback is particularly useful for implementation purposes, allowing the algorithm to converge with a constant step-size when discretized. We establish the asymptotic convergence using Lyapunov stability analysis, based on a Lyapunov function that relies on both primal and dual variables. Our numerical experiments verify that adopting integral feedback improves the convergence rate of distributed mirror descent.

A. Related Literature

1) *Gradient Tracking in Discrete Distributed Gradient Descent (DGD)*: A natural question in (discrete) distributed optimization is that whether decentralized algorithms are able to perform on par with their centralized counterparts. For purely convex problems (non-strongly convex and non-smooth), this could be done using diminishing step sizes [9], which tends agents to an agreement. However, since centralized gradient descent for strongly convex and/or smooth problems works optimally under the constant step-size setting, its decentralization was challenging. Therefore, a number of works (see e.g., [10]–[13]) have proposed the idea of gradient tracking to overcome this hurdle. The term “tracking” implies that the algorithm uses a variable calculated from past gradients to keep track of the information from the network. It then uses the variable combined with the current local gradient to output a “corrected” gradient such that the network agents are able to reach consensus. These methods accumulate past information and play a similar role as integral feedback, but

they are discrete-time. Such modification enables the decentralized algorithm to match its centralized counterpart in terms of convergence rate.

II) Continuous-time DGD: Of particular relevance to the current work is the literature on *continuous-time* DGD [14]–[19]. Similar to the centralized setup, these works construct ODEs to describe the dynamics of DGD. While the concept of DGD is rather straightforward, for continuous-time analysis in certain cases (e.g., strongly convex problem), no desirable results are obtained by simply combining gradient descent with a standard neighborhood averaging. This is in the similar spirit as the challenge in discrete DGD, overcome by gradient tracking. To tackle the continuous-time problem, multiple works have utilized the integral feedback idea [16]–[19], which introduces another variable to drive the disagreement among agents to zero. Nevertheless, these works are on gradient descent, and the results cannot be applied directly to MD, which is a more general framework, and proposing an MD algorithm to have a comparable performance is still an open research problem.

III) Distributed Mirror Descent (DMD): Decentralizing mirror descent has recently drawn a great deal of attention. While our focus is on the *continuous-time* analysis, DMD has been largely analyzed in *discrete* time in various contexts, such as online optimization [8], [20], stochastic optimization [21], [22], and the effect of delays in distributed optimization [7]. It has also been applied to social learning and belief dynamics [23]. Furthermore, Doan *et al.* [24] study the convergence of iterates for both centralized and decentralized MD. A large subset of these works (e.g., [7], [8], [20], [21], [24]) feature diminishing step-size to ensure consensus, not suitable for the strongly convex setup. Continuous-time DMD has been studied in [25], [26] with the motivation of noise-variance reduction in stochastic optimization. The main distinction of our work with these literature on DMD is adopting and analyzing the idea of integral feedback.

We also remark that the letter of Yu and Açıkmese [27] adopts an RLC-circuit perspective to DMD. The difference between [27] and this letter is that we study strongly convex functions, so our convergence guarantee is on the decision variable rather than the objective value.

B. Paper Organization

The rest of this letter is organized as follows. In Section II, we lay out the problem formulation and develop the continuous-time distributed mirror descent with integral feedback. In Section III, we provide the theoretical convergence analysis of the algorithm using Lyapunov stability analysis. Section IV provides a discretized version of our algorithm and illustrates a numerical simulation to show effectiveness of the proposed algorithm, and Section V concludes.

II. PROBLEM FORMULATION

A. Notation

We use the following notation in this letter:

$[n]$	set $\{1, 2, 3, \dots, n\}$ for any integer n
x^\top	transpose of vector x
I_d	identity matrix of size $d \times d$
$\mathbb{1}_d$	d -dimensional vector of all ones
\emptyset	vector of all zeros
$\ \cdot\ $	Euclidean norm operator
$\langle x, y \rangle$	inner product between x and y
$[x]_i$	the i -th element of the vector x
$[A]_{ij}$	the ij -th element of the matrix A
A^\dagger	pseudo inverse of matrix A
\otimes	Kronecker product operator

The vectors are all in column format. We denote by $\text{col}\{v_1, \dots, v_n\}$ the vector that stacks all vectors v_i for $i \in [n]$. We use $\text{diag}\{a_1, \dots, a_n\}$ to represent an $n \times n$ diagonal matrix that has the scalar a_i in its i -th diagonal element.

B. Network Setting

In distributed optimization, we often consider a network of n agents modeled with a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where the agents are represented by nodes $\mathcal{V} = [n]$ and the connection between two agents i and j is captured by the edge $\{i, j\} \in \mathcal{E}$. Each agent is associated with a local cost function, and if the link $\{i, j\} \in \mathcal{E}$ exists, that implies agents i and j can exchange information about their respective cost functions. Then, agent j is in the neighborhood of agent i , denoted by $\mathcal{N}_i \triangleq \{j \in \mathcal{V} : \{i, j\} \in \mathcal{E}\}$.

The agents work collectively to find the optimum of the global cost function, which is the sum of all cost local functions (to be defined precisely in Section II-C).

Assumption 1: We assume the graph \mathcal{G} is undirected and connected, i.e., there exists a path between any two distinct agents $i, j \in \mathcal{V}$. We use $\mathcal{L} \in \mathbb{R}^{n \times n}$ to represent the Laplacian of the graph \mathcal{G} .

The connectivity assumption implies that \mathcal{L} has a unique null eigenvalue. That is, $\mathcal{L}\mathbb{1}_n = \emptyset$, and $\mathbb{1}_n$ is the only direction (eigenvector) recovering the null eigenvalue.

C. Distributed Optimization Problem

In this letter, we consider a distributed (or decentralized) optimization problem in an unconstrained setting. Let us denote by $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$, the cost function associated with agent $i \in [n]$. Then, the goal is to find the optimal solution of the global cost function F , which can be written as a sum of local cost functions as follows,

$$\underset{x \in \mathbb{R}^d}{\text{minimize}} \quad F(x) = \sum_{i=1}^n f_i(x). \quad (1)$$

Since individual agents do not have knowledge of F , they cannot find the global solution on their own, and they must communicate with each other to augment their incomplete information with that of their neighborhood.

Assumption 2: For any agent $i \in \mathcal{V}$, we assume that the local cost function $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$ is convex and differentiable.

While Assumption 2 implies that the global objective function F is also convex and differentiable, we impose an additional assumption on the global cost as follows.

Assumption 3: The global function $F : \mathbb{R}^d \rightarrow \mathbb{R}$ is strongly convex. The optimal value denoted by F^* exists, and the unique solution that achieves F^* is denoted by x^* .

The assumption above will be used later in the analysis to prove the uniqueness of equilibrium for our proposed distributed continuous-time algorithm.

D. Centralized Mirror Descent

Since the focus of this letter is on the mirror descent algorithm, we provide some background on the centralized algorithm in this section, before developing the distributed algorithm in Section II-E.

Gradient descent methods iteratively minimize a first order approximation of a function plus a Euclidean regularizer. Mirror descent generalizes this idea to a non-Euclidean setup by using the notion of *Bregman divergence*, which replaces the Euclidean distance as the regularizer. The Bregman divergence is defined with respect to a generating function $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$, as follows

$$\mathcal{D}_\phi(x, x') \triangleq \phi(x) - \phi(x') - \langle \nabla \phi(x'), x - x' \rangle. \quad (2)$$

It can be immediately seen from above that the Bregman divergence is not generally symmetric, thereby it is not a distance.

Assumption 4: The generating function ϕ is closed, differentiable and μ_ϕ -strongly convex.

The assumption above is standard. For example, $\phi(x) = \frac{1}{2}\|x\|^2$ (the generator for the Euclidean distance), as well as the negative entropy function $\phi(x) = \sum_{j=1}^d [x_j] \log([x_j])$ (the generator for the Kullback–Leibler divergence) both satisfy the assumption [8].

In discrete time, the unconstrained mirror descent algorithm with learning rate η is written as

$$\begin{aligned} x^{(k+1)} = \operatorname{argmin}_x & \left\{ F(x^{(k)}) + \eta \nabla F(x^{(k)})^\top (x - x^{(k)}) \right. \\ & \left. + \mathcal{D}_\phi(x, x^{(k)}) \right\}, \end{aligned} \quad (3)$$

where using the Euclidean distance in lieu of the Bregman divergence (i.e., $\mathcal{D}_\phi(x, x^{(k)}) = \frac{1}{2}\|x - x^{(k)}\|^2$) reduces the algorithm to a gradient descent.

For writing the continuous-time dynamics of mirror descent, an equivalent form of the update above is more convenient to use. This equivalent form is based on the *convex conjugate* or *Fenchel dual* of function ϕ , which is denoted by ϕ^* and defined as follows

$$\phi^*(z) \triangleq \sup_{x \in \mathbb{R}^d} \{ \langle x, z \rangle - \phi(x) \}.$$

The definition above entails the subsequent equivalence

$$z' = \nabla \phi(x') \iff x' = \nabla \phi^*(z'),$$

and Assumption 4 guarantees that ϕ^* is μ_ϕ^{-1} -smooth. More details can be found in [28].

With the definition of ϕ^* in place, the update (3) can be rewritten in the following equivalent form

$$\begin{aligned} z^{(k+1)} &= z^{(k)} - \eta \nabla F(x^{(k)}) \\ x^{(k+1)} &= \nabla \phi^*(z^{(k+1)}). \end{aligned} \quad (4)$$

Then, taking the learning rate η to be infinitesimally small, the centralized mirror descent ODE takes the following form

$$\begin{aligned} \dot{z} &= -\nabla F(x), \\ x &= \nabla \phi^*(z), \\ x(0) &= x_0, z(0) = z_0 \text{ with } x_0 = \nabla \phi^*(z_0), \end{aligned} \quad (5)$$

which has been studied in [6, Sec. 2.1]. It is easy to see that when $\phi(x) = \frac{1}{2}\|x\|^2$, since $\phi^*(z) = \frac{1}{2}\|z\|^2$, we have that $x = \nabla \phi^*(z) = z$, and the mirror descent ODE reduces to the gradient descent ODE.

E. Distributed Mirror Descent With Integral Feedback

We now develop the distributed version of mirror descent algorithm. Motivated by the use of *integral feedback* [16]–[19] to enforce consensus among agents, we propose the following continuous-time algorithm

$$\begin{aligned} \dot{z}_i &= -\nabla f_i(x_i) + \sum_{j \in \mathcal{N}_i} (x_j - x_i) + \int_0^t \sum_{j \in \mathcal{N}_i} (x_j - x_i) \\ x_i &= \nabla \phi^*(z_i), \end{aligned} \quad (6)$$

initialized with $x_i(0) = x_{i0}$, $z_i(0) = z_{i0}$, where $x_{i0} = \nabla \phi^*(z_{i0})$.

The dual update z_i for agent $i \in [n]$ uses only private gradient information. It also enforces the primal variables in the neighborhood of i to get close to each other by using both a *consensus* term and an *integral feedback*. The *integral feedback* uses an additional state in each agent to store the information of past *consensus* term. Then, the second update maps the variable z_i back to the primal space using ϕ^* .

To analyze (6), it is more convenient to stack all the local vectors as follows

$$\begin{aligned} \mathbf{x} &\triangleq \operatorname{col}\{x_1, x_2, \dots, x_n\} \\ \mathbf{z} &\triangleq \operatorname{col}\{z_1, z_2, \dots, z_n\}, \end{aligned} \quad (7)$$

and define the following notation

$$\begin{aligned} \mathbf{L} &\triangleq \mathcal{L} \otimes I_d \\ \nabla \phi^*(\mathbf{z}) &\triangleq \operatorname{col}\{\nabla \phi^*(z_1), \nabla \phi^*(z_2), \dots, \nabla \phi^*(z_n)\} \\ \nabla f(\mathbf{x}) &\triangleq \operatorname{col}\{\nabla f_1(x_1), \nabla f_2(x_2), \dots, \nabla f_n(x_n)\}. \end{aligned} \quad (8)$$

Then, by introducing the variable \mathbf{y} to replace the integral, the dynamics given in (6) can be rewritten as follows,

$$\begin{aligned} \dot{\mathbf{z}} &= -(\nabla f(\mathbf{x}) + \mathbf{Lx} + \mathbf{y}), \\ \dot{\mathbf{y}} &= \mathbf{Lx}, \\ \mathbf{x} &= \nabla \phi^*(\mathbf{z}), \end{aligned} \quad (9)$$

where $\mathbf{y} \in \mathbb{R}^{nd}$ and $\mathbf{y}(0) = \mathbf{0}$.

III. MAIN RESULTS

In this section, we establish the theoretical convergence of the distributed mirror descent algorithm with integral feedback, proposed in (6). We prove that all agents will converge asymptotically to the minimizer of the global function F , defined in (1).

First, we propose a lemma to show that the unique equilibrium of (6) for primal variables coincides with the minimizer of problem (1), which will later be used in the proof for the asymptotic convergence to the equilibrium.

Lemma 1: Given Assumptions 1-4, an equilibrium point for the continuous-time dynamics (6) exists, and it is unique. In the equilibrium, $x_i^* = x^* = \nabla\phi^*(z_i^*)$ for all $i \in [n]$, i.e., the equilibrium point has the consensus property, and at equilibrium, the primal variable for each agent is the solution to problem (1).

Proof: Since the continuous-time dynamics (9) is equivalent to (6), to prove Lemma 1, it is sufficient to show that there exists a unique point $(\mathbf{x}^*, \mathbf{y}^*, \mathbf{z}^*)$ satisfying equilibrium conditions for (9):

– To have $\dot{\mathbf{y}} = \mathbf{0}$, we need \mathbf{x}^* to be in the null space of $\mathbf{L} = \mathcal{L} \otimes I_d$, which together with the connectivity assumption (Assumption 1), it implies that $\mathbf{x}^* = \mathbb{1}_n \otimes a$ for some vector $a \in \mathbb{R}^d$.

Next, we show that indeed $a = x^*$, where x^* is the minimizer of F .

– To have $\dot{\mathbf{z}} = \mathbf{0}$, we need

$$\nabla f(\mathbf{x}^*) + \mathbf{y}^* = \mathbf{0}. \quad (10)$$

Due to the initialization $\mathbf{y}(0) = \mathbf{0}$, we have that

$$\mathbf{y}(t) = \mathbf{L} \int_0^t \mathbf{x}(\tau) d\tau, \quad (11)$$

which implies $(\mathbb{1}_n \otimes I_d)^\top \mathbf{y}(t) = (\mathbb{1}_n \otimes I_d)^\top \mathbf{L} \int_0^t \mathbf{x}(\tau) d\tau = \mathbf{0}$. Therefore, $(\mathbb{1}_n \otimes I_d)^\top \mathbf{y}^* = \mathbf{0}$, and combining this with (10), we get

$$(\mathbb{1}_n \otimes I_d)^\top \nabla f(\mathbf{x}^*) = \mathbf{0} \implies \sum_{i=1}^n \nabla f_i(a) = \nabla F(a) = \mathbf{0}.$$

Due to the strong convexity of F in Assumption 3, the solution for a is unique and $a = x^*$ is the minimizer of F . Hence, the following point is the unique equilibrium

$$\mathbf{x}^* = \mathbb{1}_n \otimes x^*, \quad \mathbf{y}^* = -\nabla f(\mathbf{x}^*), \quad \mathbf{z}^* = \nabla\phi(\mathbf{x}^*),$$

thereby completing the proof. \blacksquare

In order to better capture the dynamics of the variables, without loss of generality, we shift the equilibrium of the dynamics to zero by defining a set of new variables

$$\tilde{\mathbf{x}} \triangleq \mathbf{x} - \mathbf{x}^*, \quad \tilde{\mathbf{y}} \triangleq \mathbf{y} - \mathbf{y}^*, \quad \tilde{\mathbf{z}} \triangleq \mathbf{z} - \mathbf{z}^*, \quad (12)$$

where $(\mathbf{x}^*, \mathbf{y}^*, \mathbf{z}^*)$ is the unique equilibrium point given in Lemma 1. We can then rewrite the first two equations of (9) as follows

$$\begin{aligned} \dot{\tilde{\mathbf{z}}} &= -(\nabla f(\tilde{\mathbf{x}} + \mathbf{x}^*) + \mathbf{L}(\tilde{\mathbf{x}} + \mathbf{x}^*) + \tilde{\mathbf{y}} + \mathbf{y}^*) \\ &= -(\nabla f(\tilde{\mathbf{x}} + \mathbf{x}^*) - \nabla f(\mathbf{x}^*)) - \mathbf{L}\tilde{\mathbf{x}} - \tilde{\mathbf{y}}, \\ \dot{\tilde{\mathbf{y}}} &= \mathbf{L}\tilde{\mathbf{x}}, \end{aligned} \quad (13)$$

where we used the fact that $\mathbf{x}^* = \mathbb{1}_n \otimes x^*$ and $\mathbf{y}^* = -\nabla f(\mathbf{x}^*)$.

Now, as the matrix $\mathbf{L} = \mathcal{L} \otimes I_d$ is symmetric and positive semi-definite, there exists a decomposition $\mathbf{L} = \mathbf{Q}\Lambda\mathbf{Q}^\top$, where \mathbf{Q} is an orthogonal matrix and $\Lambda = \text{diag}\{\lambda_1, \dots, \lambda_{nd}\}$ is a diagonal matrix. Let

$$\mathbf{S} = \mathbf{L}^{\frac{1}{2}} = \mathbf{Q}\Lambda^{\frac{1}{2}}\mathbf{Q}^\top,$$

where $\Lambda^{\frac{1}{2}} = \text{diag}\{\sqrt{\lambda_1}, \dots, \sqrt{\lambda_{nd}}\}$. Given (11), there exists a variable $\mathbf{w}(t) = \mathbf{S} \int_0^t \mathbf{x}(\tau) d\tau$ and its centered version $\tilde{\mathbf{w}} = \mathbf{w} - \mathbf{w}^*$, such that

$$\mathbf{y} = \mathbf{S}\mathbf{w}, \quad \text{and} \quad \tilde{\mathbf{y}} = \mathbf{S}\tilde{\mathbf{w}}. \quad (14)$$

Replacing $\tilde{\mathbf{y}}$ in (13) with $\mathbf{S}\tilde{\mathbf{w}}$, we have that

$$\begin{aligned} \dot{\tilde{\mathbf{z}}} &= -(\nabla f(\tilde{\mathbf{x}} + \mathbf{x}^*) - \nabla f(\mathbf{x}^*)) - \mathbf{L}\tilde{\mathbf{x}} - \mathbf{S}\tilde{\mathbf{w}}, \\ \dot{\tilde{\mathbf{w}}} &= \mathbf{S}\tilde{\mathbf{x}}. \end{aligned} \quad (15)$$

Following the proof of Lemma 1, it is straightforward to show that the dynamics above at equilibrium satisfy $\tilde{\mathbf{x}} = \mathbf{0}$. We state our main theorem below to show that the system converges asymptotically to the equilibrium provided by Lemma 1.

Theorem 1: Given Assumptions 1-4, for any starting point $x_i(0) = x_{i0}, z_i(0) = z_{i0}$ with $x_{i0} = \nabla\phi^*(z_{i0})$, the distributed mirror descent algorithm with integral feedback proposed in (6) will converge to the global optimum asymptotically, i.e., $\lim_{t \rightarrow \infty} x_i(t) = x^*$ for any $i \in [n]$.

Proof: We study the convergence of the dynamics (15). Let us consider the candidate Lyapunov function

$$V(\tilde{\mathbf{z}}, \tilde{\mathbf{w}}) = \sum_{i=1}^n \mathcal{D}_{\phi^*}(z_i, z^*) + \frac{1}{2} \tilde{\mathbf{w}}^\top \tilde{\mathbf{w}}. \quad (16)$$

Notice that the Bregman divergence used in the candidate Lyapunov function is defined with respect to ϕ^* . Since ϕ^* is convex and μ_ϕ^{-1} -smooth, $V(\tilde{\mathbf{z}}, \tilde{\mathbf{w}})$ is non-negative and has Lipschitz-continuous first derivatives. Differentiating V and recalling (6)-(13)-(15), we derive

$$\begin{aligned} \dot{V} &= \frac{d}{dt} \sum_{i=1}^n (\phi^*(z_i) - \phi^*(z^*) - \langle \nabla\phi^*(z^*), z_i - z^* \rangle) + \tilde{\mathbf{w}}^\top \frac{d\tilde{\mathbf{w}}}{dt} \\ &= \sum_{i=1}^n \langle x_i - x^*, \frac{dz_i}{dt} \rangle + \tilde{\mathbf{w}}^\top S\tilde{\mathbf{x}} \\ &= \langle \mathbf{x} - \mathbf{x}^*, \frac{d\mathbf{z}}{dt} \rangle + \tilde{\mathbf{w}}^\top S\tilde{\mathbf{x}} \\ &= \langle \tilde{\mathbf{x}}, -(\nabla f(\tilde{\mathbf{x}} + \mathbf{x}^*) - \nabla f(\mathbf{x}^*) + \mathbf{L}\tilde{\mathbf{x}} + S\tilde{\mathbf{w}}) \rangle + \tilde{\mathbf{w}}^\top S\tilde{\mathbf{x}} \\ &= -\langle \tilde{\mathbf{x}}, \nabla f(\tilde{\mathbf{x}} + \mathbf{x}^*) - \nabla f(\mathbf{x}^*) \rangle - \tilde{\mathbf{x}}^\top \mathbf{L}\tilde{\mathbf{x}}. \end{aligned}$$

It is clear from the convexity of local functions that $\dot{V} \leq 0$ at all times. When local variables do not have consensus, $\dot{V} < 0$ since $\tilde{\mathbf{x}}^\top \mathbf{L}\tilde{\mathbf{x}} < 0$. When consensus is reached, $\tilde{\mathbf{x}} = \mathbb{1}_n \otimes a$, and then the first term $\langle \tilde{\mathbf{x}}, \nabla f(\tilde{\mathbf{x}} + \mathbf{x}^*) - \nabla f(\mathbf{x}^*) \rangle = \sum_{i=1}^n \langle a, \nabla f_i(a + x^*) - \nabla f_i(x^*) \rangle$ is equal to zero if and only if $a = \mathbf{0}$, which implies $\tilde{\mathbf{x}} = \mathbf{0}$. The uniqueness of $a = \mathbf{0}$ is due to the strong convexity in Assumption 3. Therefore, the condition for equality is $\mathbf{x} = \mathbf{x}^*$, which also gives $\mathbf{z} = \mathbf{z}^*$.

The Lyapunov function satisfies $V > 0$ for $\tilde{\mathbf{z}} \neq \mathbf{0}$, and $V = 0$ when $\tilde{\mathbf{z}} = \mathbf{0}, \tilde{\mathbf{w}} = \mathbf{0}$. We also have $\dot{V} \leq 0$ with equality only at equilibrium. Then, by LaSalle's invariance principle, the dynamics (15) will converge asymptotically to its equilibrium point, and this completes the proof. \blacksquare

Theorem 1 shows that for the convergence of DMD, we do not have to appeal to a diminishing step-size (when discretized) to enforce consensus. We can remedy this using integral feedback without a loss in the convergence speed (as we see in the next section).

IV. NUMERICAL SIMULATION

In this section, we first derive a discretized version of our algorithm in (18) and then illustrate a numerical example that

shows the advantage of using integral feedback for speeding up the convergence of distributed mirror descent.

A. Discretization

Recall the continuous-time dynamics (9). We use Euler's method to derive a discrete version of the algorithm as follows. We first choose a time interval for discretization denoted by Δt . Let $t_k \triangleq k\Delta t$ and $\mathbf{x}^{(k)} \triangleq \mathbf{x}(t_k) = \mathbf{x}(k\Delta t)$. We can similarly define $\mathbf{y}^{(k)}$ and $\mathbf{z}^{(k)}$. We then have the following discrete updates

$$\begin{aligned} \frac{\mathbf{z}^{(k+1)} - \mathbf{z}^{(k)}}{\Delta t} &= -(\nabla f(\mathbf{x}^{(k)}) + \mathbf{L}\mathbf{x}^{(k)} + \mathbf{y}^{(k)}), \\ \frac{\mathbf{y}^{(k+1)} - \mathbf{y}^{(k)}}{\Delta t} &= \mathbf{L}\mathbf{x}^{(k)}. \end{aligned} \quad (17)$$

After re-arranging the terms, the fully distributed mirror descent algorithm with integral feedback takes the following (discrete) form

$$\begin{aligned} z_i^{(k+1)} &= z_i^{(k)} - \left(\nabla f_i(x_i^{(k)}) + y_i^{(k)} + \sum_{j \in \mathcal{N}_i} (x_i^{(k)} - x_j^{(k)}) \right) \Delta t, \\ y_i^{(k+1)} &= y_i^{(k)} + \sum_{j \in \mathcal{N}_i} (x_i^{(k)} - x_j^{(k)}) \Delta t, \\ x_i^{(k+1)} &= \nabla \phi^*(z_i^{(k+1)}). \end{aligned} \quad (18)$$

B. Numerical Example

We now provide a simulation for the update (18).

– *Network Structure*: We consider a 10-agent cycle network, where each agent is connected to its previous and next agent, and the last agent is connected to the first agent.

– *Generating Function for Mirror Descent*: To implement mirror descent, we employ the commonly used negative entropy as the generating function, where

$$\phi(x) = \sum_{j=1}^d [x]_j \log([x]_j) \implies [z]_i = [\nabla \phi(x)]_i = 1 + \log([x]_i),$$

and by convention $[x]_j \log([x]_j) = 0$ if $[x]_j = 0$. Here, x and z are both d -dimensional vectors. By simple calculations, it can be shown that ϕ^* , the convex conjugate of ϕ , takes the following form

$$\phi^*(z) = \sum_{j=1}^d e^{[z]_j - 1} \implies [x]_i = [\nabla \phi^*(z)]_i = e^{[z]_i - 1}.$$

Thus, we can now implement (18).

– *Global and Local Functions*: To construct the functions, we first generate a 100-dimensional vector u following a Gaussian distribution $\mathcal{N}(10 \times \mathbb{1}_d, I_d)$. We then perturb u to generate local optima $u_i = u + w_i$, where $w_i \sim \mathcal{N}(\mathbf{0}, I_d)$ for $i \in [n]$. We set the local functions $f_i(x) = \frac{1}{2} \|A_i x - b_i\|^2$, where $b_i = A_i u_i$ and $A_i \in \mathbb{R}^{20 \times 100}$ is a random matrix of rank 15. The global function becomes $F(x) = \frac{1}{2} \|\mathbf{A}x - \mathbf{b}\|^2$, where $\mathbf{A} \in \mathbb{R}^{200 \times 100}$ and $\mathbf{b} \in \mathbb{R}^{200}$ are stacked versions of their distributed counterparts. We can verify that F is strongly convex, and the closed-form solution for this problem is $x^* = \mathbf{A}^\dagger \mathbf{b}$.

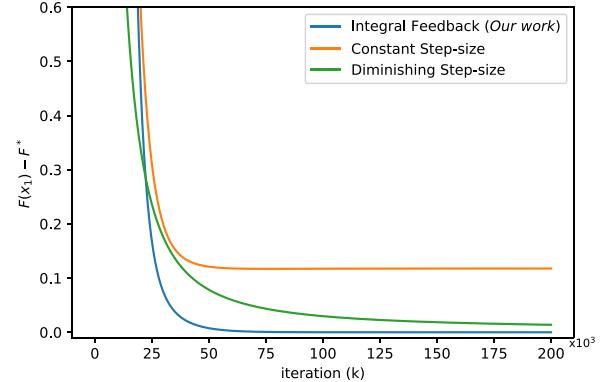


Fig. 1. The trajectory of the global objective evaluated at agent 1.

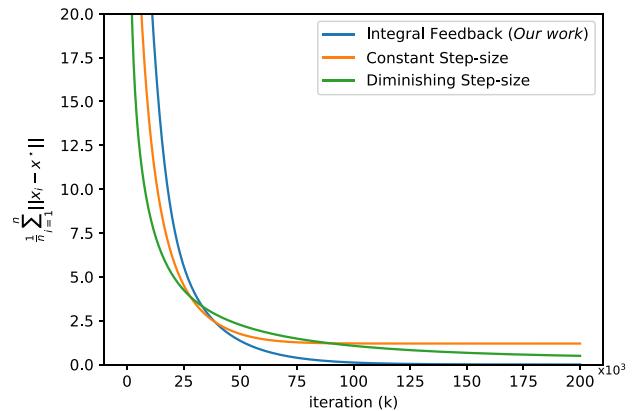


Fig. 2. The average distance of local variables from the optimum.

TABLE I
LIST OF RUN TIME AND ITERATION NEEDED FOR RESPECTIVE METHODS TO REACH $F(x_1) - F(x^*) \leq 0.01$

Method	Integral feedback	Diminishing	Constant
run time (sec)	31	130	∞
iter. number	51248	273044	∞

We run (18) with a feasible random initialization $x_i^{(0)}$, and let $y_i^{(0)} = \mathbf{0}$ and $z_i^{(0)} = \nabla \phi(x_i^{(0)})$ for every $i \in [n]$. Recall that $n = 10$ and $d = 100$, and we set $\Delta t = 10^{-2}$.

Note that the local objective functions are only convex, but the global objective function $F(x) = \sum_{i=1}^n f_i(x)$ is strongly convex in consistent with our theoretical assumptions.

– *Performance*: We compare our method with distributed mirror descent *without* integral feedback [7], [8]. These works were originally proposed for convex global functions with a suggested diminishing step-size $\frac{1}{\sqrt{k}}$. Beside that, we also include their performance with constant step-size.

For all three algorithms, we plot $F(x_1^{(k)}) - F^*$ and $\frac{1}{n} \sum_{i=1}^n \|x_i^{(k)} - x^*\|$ with respect to iteration k , in Fig. 1 and Fig. 2, respectively. We also provide a comparison on the run-time of the methods in Table I. The run-time is the required time to hit 0.01 precision on the objective value. We can see that our method converges faster than distributed mirror descent without integral feedback both in run time and in iteration number. In fact, without integral feedback, agents never converge to the global solution using a constant step-size, because the local objective functions have different

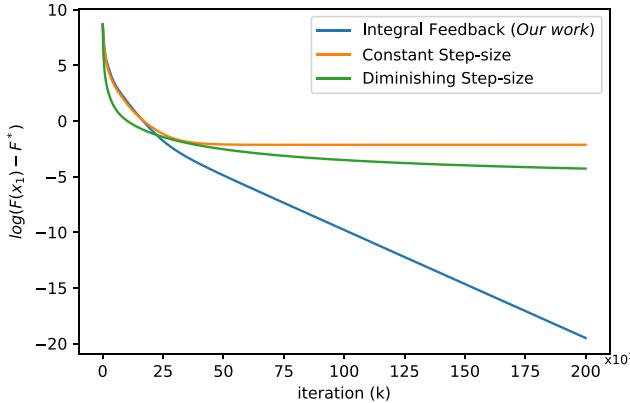


Fig. 3. The trajectory of the log-distance to global solution evaluated at agent 1.

local minima. We further plot $\log(F(x_1^{(k)}) - F^*)$ with respect to iteration k in Fig. 3. Interestingly, our method exhibits a linear convergence rate (i.e., exponentially fast), which is on par with the state-of-the-art distributed gradient descent methods (in the sense of achieving a linear rate). We reiterate that diminishing step-size is suitable for convex (and not strongly convex) global objective functions. The main purpose of the comparisons with other methods is to illustrate the power of integral feedback when the strong convexity assumption is satisfied.

V. CONCLUSION

In this letter, we considered a distributed optimization scenario where a network of agents aims at minimizing a strongly convex function, that can be written as a sum of local convex functions. The agents only have access to local gradients, but they are able to exchange information with one another. We proposed a fully decentralized mirror descent algorithm that enforces consensus among agents through a consensus term plus an additional integral feedback. We studied the continuous-time dynamics of the algorithm and provided asymptotic convergence using Lyapunov stability. Focusing on strongly convex problems, we presented empirical results verifying that distributed mirror descent with integral feedback enjoys a faster convergence rate, compared to its variants without integral feedback.

This letter provides technical analysis for the asymptotic convergence, but the simulations show that the algorithm (perhaps with smoothness assumption) can exhibit exponential convergence. Therefore, a potential future direction is the theoretical analysis of this behavior. Furthermore, studying the theoretical guarantees of (18) will shed more light on required technical assumptions in maintaining the same convergence rate when transitioning from the continuous-time update to the discrete-time update.

REFERENCES

- [1] D. Yudin and A. Nemirovskii, *Problem Complexity and Method Efficiency in Optimization*. Hoboken, NJ, USA: Wiley, 1983.
- [2] A. Ben-Tal, T. Margalit, and A. Nemirovski, “The ordered subsets mirror descent optimization method with applications to tomography,” *SIAM J. Optim.*, vol. 12, no. 1, pp. 79–108, 2001.
- [3] A. Bloch, *Hamiltonian and Gradient Flows, Algorithms and Control*, vol. 3. Providence, RI, USA: Amer. Math. Soc., 1994.
- [4] A. A. Brown and M. C. Bartholomew-Biggs, “Some effective methods for unconstrained optimization based on the solution of systems of ordinary differential equations,” *J. Optim. Theory Appl.*, vol. 62, no. 2, pp. 211–224, 1989.
- [5] U. Helmke and J. B. Moore, *Optimization and Dynamical Systems*. Heidelberg, Germany: Springer, 2012.
- [6] W. Krichene, A. Bayen, and P. L. Bartlett, “Accelerated mirror descent in continuous and discrete time,” in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, 2015, pp. 2845–2853.
- [7] J. Li, G. Chen, Z. Dong, and Z. Wu, “Distributed mirror descent method for multi-agent optimization with delay,” *Neurocomputing*, vol. 177, pp. 643–650, Feb. 2016.
- [8] S. Shahrampour and A. Jadbabaie, “Distributed online optimization in dynamic environments using mirror descent,” *IEEE Trans. Autom. Control*, vol. 63, no. 3, pp. 714–725, Mar. 2018.
- [9] A. Nedic and A. Ozdaglar, “Distributed subgradient methods for multi-agent optimization,” *IEEE Trans. Autom. Control*, vol. 54, no. 1, pp. 48–61, Jan. 2009.
- [10] W. Shi, Q. Ling, G. Wu, and W. Yin, “Extra: An exact first-order algorithm for decentralized consensus optimization,” *SIAM J. Optim.*, vol. 25, no. 2, pp. 944–966, 2015.
- [11] Y. Sun, A. Daneshmand, and G. Scutari, “Convergence rate of distributed optimization algorithms based on gradient tracking,” 2019. [Online]. Available: arXiv:1905.02637.
- [12] P. Di Lorenzo and G. Scutari, “Next: In-network nonconvex optimization,” *IEEE Trans. Signal Inf. Process. Netw.*, vol. 2, no. 2, pp. 120–136, Jun. 2016.
- [13] S. Pu and A. Nedić, “Distributed stochastic gradient tracking methods,” *Math. Program.*, vol. 756, pp. 1–49, Mar. 2020.
- [14] P. Lin, W. Ren, C. Yang, and W. Gui, “Distributed continuous-time and discrete-time optimization with nonuniform unbounded convex constraint sets and nonuniform stepsizes,” *IEEE Trans. Autom. Control*, vol. 64, no. 12, pp. 5148–5155, Dec. 2019.
- [15] S. Liu, Z. Qiu, and L. Xie, “Continuous-time distributed convex optimization with set constraints,” *IFAC Proc. Vol.*, vol. 47, no. 3, pp. 9762–9767, 2014.
- [16] B. Gharesifard and J. Cortés, “Distributed continuous-time convex optimization on weight-balanced digraphs,” *IEEE Trans. Autom. Control*, vol. 59, no. 3, pp. 781–786, Mar. 2014.
- [17] X. Zeng, P. Yi, and Y. Hong, “Distributed continuous-time algorithm for constrained convex optimizations via nonsmooth analysis approach,” *IEEE Trans. Autom. Control*, vol. 62, no. 10, pp. 5227–5233, Oct. 2017.
- [18] S. S. Kia, J. Cortés, and S. Martínez, “Distributed convex optimization via continuous-time coordination algorithms with discrete-time communication,” *Automatica*, vol. 55, pp. 254–264, May 2015.
- [19] S. Yang, Q. Liu, and J. Wang, “A multi-agent system with a proportional-integral protocol for distributed constrained optimization,” *IEEE Trans. Autom. Control*, vol. 62, no. 7, pp. 3461–3467, Jul. 2017.
- [20] D. Yuan, Y. Hong, D. W. C. Ho, and S. Xu, “Distributed mirror descent for online composite optimization,” *IEEE Trans. Autom. Control*, early access, Apr. 17, 2020, doi: 10.1109/TAC.2020.2987379.
- [21] D. Yuan, Y. Hong, D. W. C. Ho, and G. Jiang, “Optimal distributed stochastic mirror descent for strongly convex optimization,” *Automatica*, vol. 90, pp. 196–203, Apr. 2018.
- [22] M. Rabbat, “Multi-agent mirror descent for decentralized stochastic optimization,” in *Proc. IEEE 6th Int. Workshop Comput. Adv. Multi-Sensor Adaptive Process. (CAMSAP)*, 2015, pp. 517–520.
- [23] S. Shahrampour, A. Rakhlin, and A. Jadbabaie, “Distributed detection: Finite-time analysis and impact of network topology,” *IEEE Trans. Autom. Control*, vol. 61, no. 11, pp. 3256–3268, Nov. 2016.
- [24] T. T. Doan, S. Bose, D. H. Nguyen, and C. L. Beck, “Convergence of the iterates in mirror descent methods,” *IEEE Control Syst. Lett.*, vol. 3, no. 1, pp. 114–119, Jan. 2019.
- [25] A. Borovykh, N. Kantas, P. Parpas, and G. A. Pavliotis, “To interact or not? The convergence properties of interacting stochastic mirror descent,” in *Proc. Int. Conf. Mach. Learn. (ICML) Workshop Beyond First Order Methods ML Syst.*, 2020.
- [26] M. Raginsky and J. Bouvrie, “Continuous-time stochastic mirror descent on a network: Variance reduction, consensus, convergence,” in *Proc. IEEE Conf. Decis. Control (CDC)*, 2012, pp. 6793–6800.
- [27] Y. Yu and B. Açıkgöz, “RLC circuits-based distributed mirror descent method,” *IEEE Control Syst. Lett.*, vol. 4, no. 3, pp. 548–553, Jul. 2020.
- [28] J.-B. Hiriart-Urruty and C. Lemaréchal, *Fundamentals of Convex Analysis*. Heidelberg, Germany: Springer, 2012.