

Stability of Decentralized Gradient Descent in Open Multi-Agent Systems

Julien M. Hendrickx and Michael G. Rabbat

Abstract—The aim of decentralized gradient descent (DGD) is to minimize a sum of n functions held by interconnected agents. We study the stability of DGD in open contexts where agents can join or leave the system, resulting each time in the addition or the removal of their function from the global objective. Assuming all functions are smooth, strongly convex, and their minimizers all lie in a given ball, we characterize the sensitivity of the global minimizer of the sum of these functions to the removal or addition of a new function and provide bounds in $O(\min(\kappa^{0.5}, \kappa/n^{0.5}, \kappa^{1.5}/n))$ where κ is the condition number. We also show that the states of all agents can be eventually bounded independently of the sequence of arrivals and departures. The magnitude of the bound scales with the importance of the interconnection, which also determines the accuracy of the final solution in the absence of arrival and departure, exposing thus a potential trade-off between accuracy and sensitivity. Our analysis relies on the formulation of DGD as gradient descent on an auxiliary function. The tightness of our results is analyzed using the PESTO Toolbox.

I. INTRODUCTION

Multi-agent consensus optimization methods aim to solve problems of the form

$$\begin{aligned} & \underset{x_1, \dots, x_n \in \mathbb{R}^d}{\text{minimize}} && \sum_{i=1}^n f_i(x_i) \\ & \text{subject to} && x_i = x_j \text{ for all } i, j \text{ such that } a_{i,j} > 0. \end{aligned} \quad (1)$$

Here there are n agents, agent $i \in [n]$ has access to a local objective function $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$ and a local decision vector $x_i \in \mathbb{R}^d$, and the agents seek to reach agreement on the minimizer of $f(x) := \sum_{i \in [n]} f_i(x)$. The constraints couple the agents' local decision vectors through a communication graph with (weighted) adjacency matrix $A \in \mathbb{R}_+^{n \times n}$, where entry $a_{i,j}$ is (strictly) positive if and only if agent i receives messages from agent j .

This class of methods has received much attention since the decentralized (sub)gradient descent method was introduced in [9]. Previous work has primarily focused on extending these methods, e.g., by developing more efficient methods, understanding the dependence of the rate of convergence on A , considering directed and/or time-varying communication graphs, as well as various different assumptions about the class of functions being optimized; see [8] for a review.

J.H. is with ICTEAM institute, UCLouvain (Belgium) julien.hendrickx@uclouvain.be. His work was supported by the ‘‘RevealFlight’’ Concerted Research Action (ARC) of the Federation Wallonie-Bruxelles, by the Incentive Grant for Scientific Research (MIS) ‘‘Learning from Pairwise Comparisons’’ of the F.R.S.-FNRS, and by a WBI World Excellence Fellowship. M.R. is with Facebook AI Research, Montreal, Canada, mikerabbat@fb.com.

Other works have also considered the case where the objective function may vary while the optimization method is being executed. Some studies formulate the problem as regret minimization [5], [11] — finding the single decision vector that minimizes the time-averaged objective function. Other studies aim to track the instantaneous minimizer in specific problem settings, such as linear regression [2], or when the local objectives may vary smoothly in time [12].

In this work we consider the setting in which one or more agent may change their objective function f_i in a discontinuous manner as long as the objective remains in the class of smooth, strongly convex functions with bounded minimizer, and we characterize behavior with respect to the instantaneous minimizer. To motivate this setting, consider the following three scenarios.

1. In open multi-agent systems [4], agents arrive and depart while the algorithm is executing. When an upper bound is available on the total number of agents that are active at any point in time, some such systems could be modeled by having agent objectives switch to $f_i(x) = 0$ when agent i is inactive.

2. When using clusters of servers to train large machine learning models, it is desirable to have *elastic* algorithms where the number of servers may vary over time [7]. Some servers may fail during execution and leave the system. In other cases, execution may start when some minimum number of servers is available, and additional servers may join during execution as they become available.

3. In secure distributed learning [3] and federated learning [6] systems, agents use techniques from secure multi-party computation to aggregate information without revealing the values of their individual gradients to any other agent in the network (including their neighbors). An outlier or adversarial agent may wish to exert influence by modifying their local objective f_i .

Our contributions are as follows. We characterize the extent to which a single agent can influence the minimizer when each function f_i is smooth, strongly convex, and its minimizer has bounded norm. Then we go on to study the extent to which one agent can move the minimizer by changing its objective function. Finally, we establish that decentralized gradient descent is stable: there exists a bounded set such that once the iterates of the decentralized gradient method enter the set they remain in the set even if agents change their objective at every iteration. Hence, agents cannot drive the iterates to become unbounded.

II. PROBLEM FORMULATION AND NOTATION

Throughout this work, we denote the Euclidean norm by $\|\cdot\|$; so $\|x\| = (x^\top x)^{1/2}$ for a vector x , and $\|A\|$ denotes the largest singular value of a matrix A . The set of the first n natural numbers is denoted by $[n]$, and \mathbb{R}_+ denotes the set of non-negative real numbers.

A. Objective functions

The following assumptions about the local objectives f_i hold throughout this paper.

Assumption 1. *Each function f_i , $i \in [n]$, is continuously differentiable, α -strongly convex ($f_i(x) - \frac{\alpha}{2}\|x\|^2$ is convex) and β -smooth ($\|\nabla f_i(x) - \nabla f_i(y)\| \leq \beta\|x - y\|$, $\forall x, y$).*

We denote by $\mathcal{F}_{\alpha,\beta}$ the set of functions satisfying assumption Assumption 1. The *condition number* of these functions is $\kappa = \beta/\alpha$, and plays an important role in the convergence analysis of algorithms optimizing such functions.

The strong convexity part of Assumption 1 implies that $f = \sum_{i \in [n]} f_i$ is strongly convex, and hence has a unique global minimizer which we denote by $x^* \in \mathbb{R}^d$. If a local function f_i is allowed to change arbitrarily, then clearly the minimizer of f can also be changed arbitrarily. To disallow this, we impose a form of normalization on the local objectives. Let $B(x, r) = \{y : \|x - y\| \leq r\}$ denote the ball of radius r centered at x .

Assumption 2. *Let $x_i^* := \arg \min_x f_i(x)$ denote the minimizer of f_i . For each $i \in [n]$, we assume that $f_i(x_i^*) = \min_x f_i(x) = 0$ and that $x_i^* \in B(\mathbf{0}_d, 1)$, where $\mathbf{0}_d$ denotes the vector of zeros in \mathbb{R}^d .*

The minimizer of a strongly convex function is finite, so the assumption that $x_i^* \in B(0, 1)$ for all i can be seen as holding generally by rescaling the space. Similarly, since we will mainly be concerned with understanding how x^* can change if one f_i changes and we do not use the actual values of the f_i , there is no loss of generality by assuming that $f_i(x_i^*) = 0$. In the sequel, we denote by $\mathcal{F}_{\alpha,\beta}$ the set of functions that simultaneously satisfy Assumptions 1 and 2.

B. Communication graph and mixing matrix

Recall that communication constraints, describing which agents communicate directly, are captured by an agent mixing matrix $A \in \mathbb{R}_+^{n \times n}$, where $a_{i,j} > 0$ if and only if agent i receives messages from agent j . The following assumptions about A hold throughout.

Assumption 3.

- i) *The matrix A is symmetric.*
- ii) *All diagonal entries $a_{i,i} > 0$ are positive.*
- iii) *The matrix A is primitive: there exists an integer $k > 0$ such that every entry of A^k is positive.*

It is common to view A as the (weighted) adjacency matrix of a graph $G = (V, E)$, with vertices $V = [n]$ corresponding to agents and edges $E = \{(i, j) : a_{i,j} > 0, i \neq j\}$ between agents that exchange messages. In this case, the third item

in Assumption 3 is equivalent to assuming that the graph G is connected.

We will make use of standard concepts from spectral graph theory. Let D denote a $n \times n$ diagonal matrix with entries $D_{i,i} = \sum_j a_{i,j}$ equal to the weighted degree of agent i . The graph Laplacian matrix is $L = D - A$. Since A is symmetric, by Assumption 3, L is too, and so it has an eigendecomposition. Let $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ denote the n eigenvalues of L sorted in ascending order. It follows from Assumption 3(iii) that $\lambda_1 = 0$ and $\lambda_2 > 0$, therefore L is positive semi-definite. In addition, the null space of L is spanned by $\mathbf{1}_d$, the d -dimensional vector of ones.

Let I_d denote the $d \times d$ identity matrix, and let $\mathbf{L} = L \otimes I_d$. Let $\mathbf{x} \in \mathbb{R}^{nd}$ denote the vector obtained by stacking x_1, \dots, x_n , where $x_i \in \mathbb{R}^d$ is the decision vector at agent i . One can verify that

$$\mathbf{x}^\top \mathbf{L} \mathbf{x} = \sum_{i \in [n]} \sum_{j > i} a_{i,j} \|x_i - x_j\|^2. \quad (2)$$

Similarly, if we let $\mathbf{y} = \mathbf{L} \mathbf{x} \in \mathbb{R}^{nd}$ denote a vector with blocks $y_1, \dots, y_n \in \mathbb{R}^d$, then one can verify that

$$y_i = - \sum_{j=1}^n a_{i,j} (x_i - x_j). \quad (3)$$

C. Decentralized Gradient Descent (DGD)

DGD is commonly defined by the iterations

$$x_i^{k+1} = \sum_{j=1}^n \tilde{a}_{i,j} x_j^k - \eta \nabla f_i(x_i^k), \quad (4)$$

with stochastic matrix \tilde{A} [9]. Here we consider a more general form of DGD given by

$$x_i^{k+1} = x_i^k - \eta \left(\nabla f_i(x_i^k) + \rho \sum_{j=1}^n a_{i,j} (x_i^k - x_j^k) \right). \quad (5)$$

Note that (5) and (4) are equivalent provided that $\eta\rho$ is not too large, as can be seen by taking $\tilde{a}_{i,j} = \eta\rho a_{i,j}$ for $i \neq j$ and $\tilde{a}_{i,i} = 1 - \eta\rho \sum_j a_{i,j}$.

We study (5) because it can be seen as applying gradient descent to a penalized objective function [14]. Recall that $\mathbf{x} \in \mathbb{R}^{nd}$ denotes the vector obtained by stacking x_1, \dots, x_n , and define $F(\mathbf{x}) := \sum_{i \in [n]} f_i(x_i)$. One way to incorporate the constraints in (1) is through quadratic penalties:

$$\begin{aligned} F_\rho(\mathbf{x}) &:= F(\mathbf{x}) + \frac{\rho}{2} \sum_{i \in [n]} \sum_{j > i} a_{i,j} \|x_i - x_j\|^2 \\ &= F(\mathbf{x}) + \frac{\rho}{2} \mathbf{x}^\top \mathbf{L} \mathbf{x}. \end{aligned} \quad (6)$$

where $\rho > 0$ controls the weight given to the penalty relative to the objective F and we have used (2) for the second equality. Minimizing F_ρ leads to a trade-off governed by ρ between the equality of the x_i and the minimization of the f_i .

Applying the gradient method with constant step-size $\eta > 0$ to minimize F_ρ leads to updates of the form

$$\mathbf{x}^{k+1} = \mathbf{x}^k - \eta (\nabla F(\mathbf{x}) + \rho \mathbf{L} \mathbf{x}),$$




Fig. 1. Representation of the level sets of f_1 and f_2 defined in (7) in (a), and of f_1 from (7) and f_b from (15) (b), together with the minimizer x^* of their sum, for $\kappa = 100$ and $\phi = \tan^{-1} \kappa^{-1/2} \simeq 0.0997$ (rad). One can see the minimizer can be at a distance $\simeq \sqrt{\kappa}/2$ of 0, and that a small rotation of a function can have an important effect on its position.

which is equivalent to (5). By observing that the updates (5) correspond to the gradient method applied to F_ρ , standard results (e.g., [10], [1]) directly give that the iterates (5) converge linearly to the minimizer of F_ρ when the step size η is chosen appropriately.

The following result establishes that f , F , and F_ρ are smooth and strongly convex; the proof involves standard algebraic manipulations and is omitted for space reasons.

Proposition 1. *Suppose that $f_1, \dots, f_n \in \mathcal{F}_{\alpha, \beta}$. Then:*

- i) F is α -strongly convex and β -smooth.
- ii) F_ρ is α -strongly convex and $(\beta + \rho\lambda_n)$ -smooth, where λ_n is the largest eigenvalue of L .
- iii) f is $n\alpha$ -strongly convex and $n\beta$ -smooth.

III. LOCATION OF THE MINIMIZERS

Since f , F , and F_ρ are strongly convex, they each have a unique minimizer. Moreover, since $F(x)$ is separable in x_1, \dots, x_n , its minimizer $x^* = \arg \min_x F(x)$ is trivially the vector obtained by stacking x_1^*, \dots, x_n^* , where $x_i^* = \arg \min_x f_i(x)$, and thus $x^* \in B(\mathbf{0}_d, 1)^n$.

Since the minimizers x_i^* of the f_i are all in $B(0_d, 1)$ in view of Assumption 2, one may have hoped that $x^* = \arg \min f(x) = \sum_{i \in [n]} f_i(x)$ would be close to that ball. Unfortunately, this is far from true, as shown by the following example. Consider the following two quadratic functions

$$f_1(x, y) = \frac{\beta}{2}((x+1)\cos\phi - y\sin\phi)^2 + \frac{\alpha}{2}((x+1)\sin\phi + y\cos\phi)^2 \quad (7)$$

$$f_2(x, y) = \frac{\beta}{2}((x-1)\cos\phi + y\sin\phi)^2 + \frac{\alpha}{2}(-(x+1)\sin\phi + y\cos\phi)^2 \quad (8)$$

for scalar coordinates x and y and some angle ϕ , as illustrated in Figure 1. Their minimizers are clearly $(-1, 0)$ and $(1, 0)$, both in $B(0_2, 1)$, and one can verify that they are

α -strongly convex and β -smooth, as they can be obtained by applying a rotation $\pm\phi$ to $\frac{\alpha}{2}x^2 + \frac{\beta}{2}y^2$ followed by a translation. Simple derivations show that the minimum of $f_1 + f_2$ is reached at $\left(0, \frac{(\kappa-1)\tan\phi}{1+\kappa\tan^2\phi}\right)$, which, for $\tan\phi = 1/\sqrt{\kappa}$ becomes $\left(0, \frac{\kappa-1}{2\sqrt{\kappa}}\right)$. So even though the minimizer of each function is in $B(0_2, 1)$, the norm of the minimizer of their sum scales as $\mathcal{O}(\sqrt{\kappa})$. The next result shows that this is indeed the worst scaling possible up to a constant factor.

Theorem 1. *Let $\kappa = \beta/\alpha$ denote the condition number of f , and let $R_\kappa := 1 + \sqrt{\kappa}$. If $f_i \in \mathcal{F}_{\alpha, \beta}$ for all $i = 1, \dots, n$, then the following two inclusions, respectively in \mathbb{R}^d and \mathbb{R}^{nd} , hold:*

$$\arg \min_x f(x) \in B(\mathbf{0}_d, R_\kappa), \quad \text{and} \quad (9)$$

$$\arg \min_x F_\rho(x) \in B(\mathbf{0}_d, R_\kappa)^n. \quad (10)$$

Proof: By Assumption 2, we have $f_i(x_i^*) = 0$ and $x_i^* \in B(0, 1)$. Since f_i is β -smooth, this implies that

$$f_i(\mathbf{0}_d) \leq \frac{\beta}{2}. \quad (11)$$

For any $x \notin B(\mathbf{0}_d, R_\kappa)$ we have $\|x - x_i^*\| > R_\kappa - 1 = \sqrt{\kappa}$ since $\|x\| > R_\kappa$. Then α -strong convexity of f_i implies that

$$f_i(x) > \frac{\alpha\kappa}{2} = \frac{\beta}{2}, \quad \text{for all } x \notin B(\mathbf{0}_d, R_\kappa). \quad (12)$$

Using this bound along with (11), we obtain that for any $x \notin B(\mathbf{0}_d, R_\kappa)$,

$$f_i(x) > f_i(\mathbf{0}_d), \quad (13)$$

and thus

$$f(x) > \frac{n\beta}{2} \geq f(\mathbf{0}_d) \geq \min_{\bar{x} \in \mathbb{R}^d} f(\bar{x}).$$

Therefore, $\arg \min_x f(x) \in B(\mathbf{0}_d, R_\kappa)$, which completes the proof of (9).

Next, consider a point $\mathbf{x} = (x_1, x_2, \dots, x_n) \notin B(\mathbf{0}_d, R_\kappa)^n$, and let $\mathbf{x}' = (x'_1, \dots, x'_n)$ denote the projection of \mathbf{x} onto $B(\mathbf{0}_d, R_\kappa)^n$; i.e.,

$$x'_i := \begin{cases} x_i & \text{if } x_i \in B(\mathbf{0}_d, R_\kappa), \\ R_\kappa \frac{x_i}{\|x_i\|} & \text{otherwise.} \end{cases}$$

We will show that $F_\rho(\mathbf{x})$ is larger than $F_\rho(\mathbf{x}')$, which then implies the claim (10). Since a projection onto a convex set is non-expansive, we have $\|x'_i - x'_j\| \leq \|x_i - x_j\|$ for all $i, j \in [n]$, and thus $(\mathbf{x}')^\top \mathbf{L} \mathbf{x}' \leq \mathbf{x}^\top \mathbf{L} \mathbf{x}$. Moreover, there is at least one $i \in [n]$ such that $x'_i \neq x_i$ because $\mathbf{x} \notin B(\mathbf{0}_d, R_\kappa)^n$, and for every such i we know that x'_i is a convex combination of x_i and $\mathbf{0}_d$; i.e., $x'_i = \theta x_i + (1-\theta)\mathbf{0}_d$ for some $\theta \in (0, 1)$. For those indices i such that $x'_i \neq x_i$, we have $x_i \notin B(\mathbf{0}_d, R_\kappa)$, and so using convexity of f_i and (13), we get

$$f_i(x'_i) \leq \theta f_i(x_i) + (1-\theta)f_i(\mathbf{0}_d) < f_i(x_i).$$

Thus, for any $\mathbf{x} \notin B(\mathbf{0}_d, R_\kappa)^n$, we have $F(\mathbf{x}) > F(\mathbf{x}')$, and therefore $F_\rho(\mathbf{x}) > F_\rho(\mathbf{x}') \geq \min_{\bar{x}} F_\rho(\bar{x})$. Hence, $\arg \min_x F_\rho(x) \in B(\mathbf{0}_d, R_\kappa)^n$, which establishes (10). \square

The bound for f is tight up to a constant factor even for $n = 2$, as shown by the example (7). The bound for F_ρ should be tight for large ρ as its minimizer approaches $1_n \otimes x^*$. For small ρ , F_ρ is close to F and its minimizer should be closer to $B(0_d, 1)^n$. However, large values of ρ is the relevant regime, as the goal of the minimization of F_ρ is to find approximations of x^* .

IV. IMPACT OF FUNCTION CHANGE

We now study to what extent the minimizer $x^* = \arg \min_x f(x)$ can change when the local objective f_i at one agent changes. Considering a single change is justified if changes are sufficiently infrequent to only have to consider one at the time. Suppose we have $n + 1$ functions $f_1, f_2, \dots, f_{n-1}, f_a, f_b \in \mathcal{F}_{\alpha, \beta}$. Let x^a and x^b be defined as the minimizers,

$$\begin{aligned} x^a &:= \arg \min_x (f_1(x) + f_2(x) + \dots + f_{n-1}(x) + f_a(x)) \\ x^b &:= \arg \min_x (f_1(x) + f_2(x) + \dots + f_{n-1}(x) + f_b(x)). \end{aligned} \quad (14)$$

We wish to analyze $\|x^a - x^b\|$. When $n = 2$, x^a and x^b can be very different even if f_a and f_b are very similar. If $f_a = f_2$ as defined in (7), then we have seen that the minimizer of $f_1 + f_a$ is $(0, \frac{\kappa-1}{2\sqrt{\kappa}})$. Consider now

$$\begin{aligned} f_b(x, y) &= \frac{\beta}{2}((x+1)\cos\phi - y\sin\phi)^2 \\ &+ \frac{\alpha}{2}((x+1)\sin\phi + y\sin\phi)^2, \end{aligned} \quad (15)$$

with $\tan\phi = \sqrt{\kappa}$. Note f_b can be obtained by applying a rotation of $2\phi = 2/\sqrt{\kappa}$ to f_a around its minimizer, which is small if the condition number is large, see Figure 1. Nevertheless, the minimum of $f_1 + f_b$ is $(0, 0)$, so changing f_a to f_b moves the minimizer by $\Omega(\sqrt{\kappa})$.

A first bound on $\|x^a - x^b\|$ in $O(\sqrt{\kappa})$ follows from Theorem 1 as both x^a and x^b belong to $B(0_d, R_\kappa)$. The next theorem presents stronger bounds exploiting n , as modifying a function clearly has a smaller potential impact when n is large. Its proof, omitted for space reasons, will appear in an extended version of this work.

Theorem 2. For x^a and x^b defined in (14),

$$\|x^a - x^b\| \leq \min\left(4 + 4\sqrt{\kappa}, \frac{4\sqrt{\kappa} + 2\kappa}{\sqrt{n-1}}, \frac{4\kappa + 2\kappa^{3/2}}{n}\right).$$

The tightness of Theorem 2 can be analyzed empirically using the PESTO toolbox [13], initially developed to compute the exact worst-case performances of optimization algorithms. This toolbox allows indeed solving exactly (up to numerical precision) optimization problems of the form $\max \|x^a - x^b\|$ over all functions $f_1, f_2, \dots, f_{n-1}, f_a, f_b \in \mathcal{F}_{\alpha, \beta}$ satisfying our assumptions, and x^a, x^b defined as in (14). It allows thus computing the worst possible value for $\|x^a - x^b\|$ appearing in Theorem 2, for given n and κ . Figure 2 clearly indicates that this worst-case value evolves as $\sqrt{\kappa}$ for a given n . Similar tests show a decrease of the worst-case distance with n for a given κ , but no clear expression of this




Fig. 2. Evolution with $\sqrt{\kappa}$ of the exact worst-case bound on the distance $\|x^a - x^b\|$ for x^a, x^b defined in (14). These bounds are computed numerically using the PESTO toolbox for performance estimation [13]. Results suggest a clear asymptotic scaling in $\sqrt{\kappa}$ for fixed n .

decay was identified so far. This suggests the tightness of our bound for fixed n and growing κ , but possible improvements in other regimes.

V. STABILITY OF DGD

We have seen in Section IV that modifications of certain f_i may result in large variations of the minimizer. If these happen during the execution of decentralized gradient descent (DGD), they can lead to sudden dramatic increases of the distance to the optimum. We now study the stability of the DGD iteration (5) in this context, taking advantage of its formulation as a gradient descent on the objective $F_\rho(x)$.

Recall the following standard result on gradient descent; see, e.g., [10], [1].

Theorem 3. Let $\bar{f} : \mathbb{R}^d \rightarrow \mathbb{R}$ be a $\bar{\beta}$ -smooth, $\bar{\alpha}$ -strongly convex function minimized at $x^* = \arg \min_x \bar{f}(x)$. Let η be a positive scalar which is at most $1/\bar{\beta}$. For any initial point $x^0 \in \mathbb{R}^d$, the sequence $\{x^k, k > 0\}$, produced by the gradient method $x^{k+1} = x^k - \eta \nabla \bar{f}(x^k)$ satisfies

$$\begin{aligned} \|x^{k+1} - x^*\|^2 &\leq (1 - \eta\bar{\alpha}) \|x^k - x^*\|^2 \\ &\leq (1 - \eta\bar{\alpha})^{k+1} \|x^0 - x^*\|^2. \end{aligned} \quad (16)$$

In addition to establishing linear convergence of the gradient method, this standard result ensures monotonic progress is made towards the minimizer at every step. Recall that applying the gradient method to the objective F_ρ in (6) leads to the decentralized algorithm (5). Thus, the iterates x^k of (5) converge monotonically to x^ρ , the minimizer of F_ρ , provided $\eta < 1/(\beta + \rho\lambda_n)$ in view of Proposition 1.

If the local objective function f_i at one or more agent may change during the execution of the optimization method, clearly convergence is no longer guaranteed, but we will show and quantify the stability of DGD under arbitrary changes of functions satisfying our assumptions. Specifically, we study the dynamics

$$x_i^{k+1} = x_i^k - \eta \left(\nabla f_i^k(x_i^k) + \rho \sum_{j=1}^n a_{i,j} (x_i - x_j) \right), \quad (17)$$

where the f_i^k can thus change at each iteration k . This can be seen as a gradient descent on the time-varying objective

$$F_\rho^k(\mathbf{x}) = F^k(\mathbf{x}) + \frac{\rho}{2} \mathbf{x}^\top \mathbf{L} \mathbf{x}, \quad (18)$$

with $F^k(\mathbf{x}) = \sum_{i=1}^n f_i^k(x_i)$. When $f_i^k \in \mathcal{F}_{\alpha,\beta}$ for all $i \in [n]$ and $k \geq 0$, we know from Theorem 1 that, for all $k \geq 0$, the minimizer of F_ρ^k lies in the ball $B(\mathbf{0}_d, R_\kappa)^n$, where $R_\kappa = 1 + \sqrt{\kappa}$. We will show that, even when every f_i^k may vary from step to step, the dynamics (17) cannot cause the sequence \mathbf{x}^k to grow unbounded if $f_i^k \in \mathcal{F}_{\alpha,\beta}$ for all $k \geq 0$.

Theorem 4. *Consider the sequence \mathbf{x}^k generated by updates (17) with possibly time-varying objective F_ρ^k (equiv., f_i^k). Assume that $f_i^k \in \mathcal{F}_{\alpha,\beta}$ for all $i \in [n]$ and $k \geq 0$, and $\eta \leq 1/(\beta + \rho\lambda_n)$. Let $\kappa_\rho = \frac{\beta + \rho\lambda_n}{\alpha}$ and let $\kappa = \beta/\alpha$. There exists a finite positive constant R ,*

$$R = \begin{cases} (1 + \sqrt{6})\sqrt{n}(1 + \sqrt{\kappa}) & \text{if } \kappa_\rho < 3 \\ (1 + \sqrt{2\kappa_\rho})\sqrt{n}(1 + \sqrt{\kappa}) & \text{otherwise,} \end{cases}$$

such that if $\|\mathbf{x}^{k_0}\| \leq R$ at some iteration k_0 , then $\|\mathbf{x}^k\| \leq R$ for all $k \geq k_0$.

A minor extension of our result not presented here shows the iterates \mathbf{x}^k approach $B(0, R)$ if $\|\mathbf{x}^{k_0}\| \geq R$. Moreover, since the proof is based on the analysis of a single iteration, Theorem 4 can directly be extended to time-varying network matrices with uniformly bounded λ_n . Finally, observe that the bound R scales with the square root of the condition number κ_ρ of F_ρ , which itself grows linearly with the penalty parameter ρ . This quantity can also be shown to determine the accuracy of the points to which each agent i converges. Hence our result reveals a potential trade-off between accuracy and stability.

A. Intermediate Results on Gradient Descent

The proof of Theorem 4 builds on an intermediate result in a more general setting: Let $\bar{f}: \mathbb{R}^d \rightarrow \mathbb{R}$ be an $\bar{\alpha}$ -strongly convex and $\bar{\beta}$ -smooth function, and let $\bar{\kappa} = \bar{\beta}/\bar{\alpha}$. Let $x^* = \arg \min_x \bar{f}(x)$ denote the minimizer of \bar{f} , and suppose that there exists a finite positive scalar b such that $\|x^*\| \leq b$. Let η be a non-negative scalar that is at most $1/\bar{\beta}$. For any initial point $x \in \mathbb{R}^d$, let $x^+ = x - \eta \nabla \bar{f}(x)$. We are interested in finding a constant R such that if $\|x\| \leq R$ then $\|x^+\| \leq R$, i.e. the radius of a ball centered on 0 stable under an iteration of gradient descent on \bar{f} .

The next lemma allows us to simplify our focus to the case where $\eta = 1/\bar{\beta}$.

Lemma 1. *Let R be a finite positive scalar. If $\|x\| \leq R$ and $\|x - (1/\bar{\beta})\nabla \bar{f}(x)\| \leq R$, then $\|x - \eta \nabla \bar{f}(x)\| \leq R$ for any η satisfying $0 \leq \eta \leq 1/\bar{\beta}$.*

Proof: By hypothesis, x and $x - (1/\bar{\beta})\nabla \bar{f}(x)$ are in the convex set $B(\mathbf{0}_d, R)$. Now, for $0 < \eta < 1/\bar{\beta}$, the point $x - \eta \nabla \bar{f}(x)$ is a convex combination of x and $x - (1/\bar{\beta})\nabla \bar{f}(x)$, and hence also belongs to $B(\mathbf{0}_d, R)$ by convexity. \square

The next proposition provides a stability result provided x is large enough.

Proposition 2. *Assume $\|x\| \geq b$ and $\bar{\kappa} \geq 3$, and let*

$$R = (1 + \sqrt{2\bar{\kappa}})b. \quad (19)$$

If $\|x\| \leq R$, then $\|x^+\| \leq R$.

Proof: By standard conditions for smoothness and strong convexity of \bar{f} , see e.g. [10], we have for any x ,

$$\nabla \bar{f}(x)^\top (x - x^*) \geq \frac{\bar{\beta}^{-1} \|\nabla \bar{f}(x)\|^2}{1 + \bar{\kappa}^{-1}} + \frac{\bar{\alpha} \|x - x^*\|^2}{1 + \bar{\kappa}^{-1}}.$$

Adding $\nabla \bar{f}(x)^\top x^*$ to both sides, applying the Cauchy-Schwartz inequality, and recalling that $\|x^*\| \leq b$ by assumption gives

$$\begin{aligned} \nabla \bar{f}(x)^\top x &\geq \nabla \bar{f}(x)^\top x^* + \frac{\bar{\beta}^{-1} \|\nabla \bar{f}(x)\|^2}{1 + \bar{\kappa}^{-1}} + \frac{\bar{\alpha} \|x - x^*\|^2}{1 + \bar{\kappa}^{-1}} \\ &\geq -b \|\nabla \bar{f}(x)\| + \frac{\bar{\beta}^{-1} \|\nabla \bar{f}(x)\|^2}{1 + \bar{\kappa}^{-1}} + \frac{\bar{\alpha} \|x - x^*\|^2}{1 + \bar{\kappa}^{-1}}. \end{aligned} \quad (20)$$

We focus on the case $\eta = 1/\bar{\beta}$. By definition of x^+ ,

$$\begin{aligned} \|x^+\|^2 &= \left\| x - \frac{1}{\bar{\beta}} \nabla \bar{f}(x) \right\|^2 \\ &= \|x\|^2 + \frac{1}{\bar{\beta}^2} \|\nabla \bar{f}(x)\|^2 - 2 \frac{1}{\bar{\beta}} \nabla \bar{f}(x)^\top x. \end{aligned}$$

Using (20), we can then bound $\|x^+\|^2$

$$\begin{aligned} \|x^+\|^2 &\leq \|x\|^2 + \frac{1}{\bar{\beta}^2} \|\nabla \bar{f}(x)\|^2 \\ &\quad + \frac{2}{\bar{\beta}} \left(b \|\nabla \bar{f}(x)\| - \frac{\bar{\beta}^{-1} \|\nabla \bar{f}(x)\|^2}{1 + \bar{\kappa}^{-1}} - \frac{\bar{\alpha} \|x - x^*\|^2}{1 + \bar{\kappa}^{-1}} \right) \\ &= \|x\|^2 + \frac{2}{\bar{\beta}} \left(b \|\nabla \bar{f}(x)\| - \frac{1}{2\bar{\beta}} \left(\frac{1 - \bar{\kappa}^{-1}}{1 + \bar{\kappa}^{-1}} \right) \|\nabla \bar{f}(x)\|^2 \right. \\ &\quad \left. - \frac{\bar{\alpha}}{1 + \bar{\kappa}^{-1}} \|x - x^*\|^2 \right). \end{aligned} \quad (21)$$

Since $\bar{\kappa} \geq 3$ by assumption, we have $\frac{1 - \bar{\kappa}^{-1}}{1 + \bar{\kappa}^{-1}} > 0$. Thus

$$b \|\nabla \bar{f}(x)\| - \frac{1}{2\bar{\beta}} \left(\frac{1 - \bar{\kappa}^{-1}}{1 + \bar{\kappa}^{-1}} \right) \|\nabla \bar{f}(x)\|^2 - \frac{\bar{\alpha} \|x - x^*\|^2}{1 + \bar{\kappa}^{-1}}$$

is strictly concave in $\|\nabla \bar{f}(x)\|$, and hence can be bounded by its maximum (w.r.t. to $\|\nabla \bar{f}(x)\|$)

$$\frac{b^2 \bar{\beta}}{2} \left(\frac{1 + \bar{\kappa}^{-1}}{1 - \bar{\kappa}^{-1}} \right) - \frac{\bar{\alpha}}{1 + \bar{\kappa}^{-1}} \|x - x^*\|^2.$$

Using this in (21) gives that

$$\begin{aligned} \|x^+\|^2 &\leq \|x\|^2 + b^2 \frac{1 + \bar{\kappa}^{-1}}{1 - \bar{\kappa}^{-1}} - 2 \frac{\bar{\kappa}^{-1}}{1 + \bar{\kappa}^{-1}} \|x - x^*\|^2 \\ &\leq \|x\|^2 + b^2 \frac{1 + \bar{\kappa}^{-1}}{1 - \bar{\kappa}^{-1}} - 2 \frac{\bar{\kappa}^{-1}}{1 + \bar{\kappa}^{-1}} (\|x\| - \|x^*\|)^2. \end{aligned}$$

For $\bar{\kappa} \geq 3$, we have $\frac{1 + \bar{\kappa}^{-1}}{1 - \bar{\kappa}^{-1}} \leq 2$ and $-\frac{2\bar{\kappa}^{-1}}{1 + \bar{\kappa}^{-1}} \leq -\bar{\kappa}^{-1}$. Furthermore, since $\|x\| \geq b$, we have $\|x\| - b \geq 0$.

$$Q(\|x\|) = \|x\|^2 + 2b^2 - \bar{\kappa}^{-1} (\|x\| - b)^2.$$

Now observe that $Q(\|x\|)$ is strictly convex in $\|x\|$ since $1 - \bar{\kappa}^{-1} > 0$. Therefore, under the assumption that $0 \leq \|x\| \leq R$, $Q(\|x\|)$ achieves its maximum when either $\|x\| = 0$ or $\|x\| = R$. Observe that

$$Q\left((1 + \sqrt{2\bar{\kappa}})b\right) = (1 + \sqrt{2\bar{\kappa}})^2 b^2 = R^2$$

and, for $\bar{\kappa} \geq 3$,

$$Q(0) = (2 - \bar{\kappa}^{-1})b^2 < (1 + \sqrt{2\bar{\kappa}})^2 b^2 = R^2,$$

Therefore, if $\|x\|^2 \leq R^2$ and $\eta = 1/\bar{\beta}$, then $\|x^+\|^2 \leq R^2$ too. It follows from Lemma 1 that $\|x^+\|^2 \leq R^2$ for any $\eta \in [0, 1/\bar{\beta}]$, which completes the proof. \square

The previous proposition requires $\|x\| \geq b$ and a condition number $\bar{\kappa} = \bar{\beta}/\bar{\alpha} \geq 3$. We will see that smaller condition numbers can easily be treated. To complete our analysis, we just need thus to focus on x with norms smaller than the bound b on x^* .

Proposition 3. *If $\|x\| \leq b$ then $\|x^+\| \leq 3b$.*

Proof: We focus again first on the case $\eta = 1/\bar{\beta}$. Using Theorem 3, we have

$$\|x^+ - x^*\| \leq \sqrt{1 - \kappa^{-1}} \|x - x^*\|$$

Using $\|x\| \leq b$ and $\|x^*\| \leq b$, which implies $\|x - x^*\| \leq 2b$ we have then

$$\|x^+\| \leq \|x^+ - x^*\| + \|x^*\| \leq \|x - x^*\| + \|x^*\| \leq 3b.$$

The result for other values of $\eta \in [0, 1/\bar{\beta}]$ follows from Lemma 1. \square

We combine the results above in the following Proposition.

Proposition 4. *Let*

$$R := \begin{cases} (1 + \sqrt{2\bar{\kappa}})b & \text{if } \bar{\kappa} \geq 3, \\ (1 + \sqrt{6})b & \text{otherwise.} \end{cases}$$

If $\|x\| \leq R$ then $\|x^+\| \leq R$.

Proof: Suppose first that $\bar{\kappa} \geq 3$. If $\|x\| \geq b$, the result directly follows from Proposition 2. If $\|x\| < b$, Proposition 3 and $\bar{\kappa} \geq 3$ imply that

$$\|x^+\| \leq 3b \leq (1 + \sqrt{6})b \leq (1 + \sqrt{2\bar{\kappa}})b.$$

Smaller condition numbers $\bar{\kappa} = \bar{\beta}/\bar{\alpha}$ can be artificially increased to $\bar{\kappa}' = 3$ e.g. by considering a lower $\bar{\alpha}' = \bar{\beta}'/3 < \bar{\alpha}$. An $\bar{\alpha}$ -strongly convex function is indeed always $\bar{\alpha}'$ -strongly convex for any positive $\bar{\alpha}' < \bar{\alpha}$. The first part of the result can then be applied with $R = (1 + \sqrt{2\bar{\kappa}'})b = (1 + \sqrt{6})b$. \square

B. Proof of Theorem 4

Proof: The proof essentially follows by specializing the result of Proposition 4 to the setting of (17) and (18). By Theorem 1 we know that

$$\arg \min_x F_\rho^k(x) \in B(\mathbf{0}_d, 1 + \sqrt{\kappa})^n, \quad \text{for all } k \geq 0,$$

where $\kappa = \beta/\alpha$, and so we have $\|\arg \min_x F_\rho^k(x)\| \leq b := \sqrt{n}(1 + \sqrt{\kappa})$. In addition, it follows from Proposition 1 that the condition number of F_ρ^k is $\kappa_\rho = (\beta + \rho\lambda_n)/\alpha$ for all $k \geq 0$. The result follows then from Proposition 4. \square

VI. CONCLUSION

We have shown that DGD is stable: even if all agents change their objectives at every iteration, if the local per-agent objectives are constrained to be smooth, strongly convex, and their minimizer has bounded norm, then the iterates of DGD cannot be made to grow unbounded.

This paper contributes an initial exploration into the behavior of DGD, and consensus optimization methods in general, in the setting of open multi-agent systems. There are many interesting directions for future work, such as understanding the behavior of more advanced algorithms (e.g., those using gradient tracking), characterizing a stability region under stronger or different assumptions (e.g., if the frequency of function changes is constrained to be slower than the DGD update iterations, or alternative assumptions about local objective functions), and developing robust decentralized methods with tighter stability regions than DGD.

REFERENCES

- [1] S. Bubeck. Convex optimization: Algorithms and complexity. *Foundations and Trends in Machine Learning*, 8(3–4):231–357, 2015.
- [2] F. S. Cattivelli and A. H. Sayed. Diffusion LMS strategies for distributed estimation. *IEEE Trans. Signal Processing*, 58(3):1035–1048, 2010.
- [3] V. Chen, V. Pastro, and M. Raykova. Secure computation for machine learning with SPDZ. In *NeurIPS Workshop on Privacy Preserving Machine Learning*, Montréal, Canada, Dec. 2018.
- [4] J. M. Hendrickx and S. Martin. Open multi-agent systems: Gossiping with random arrivals and departures. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 763–768. IEEE, 2017.
- [5] S. Lee, A. Nedić, and M. Raginsky. Stochastic dual averaging for decentralized online optimization on time-varying communication graphs. *IEEE Trans. Automatic Control*, 62(12):6407–6414, 2017.
- [6] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 54. JMLR Workshops and Conference Proceedings, 2017.
- [7] S. Narayanamurthy, M. Weimer, D. Mahajan, T. Condie, S. Sellamany, and K. Selvaraj. Towards resource-elastic machine learning. In *NeurIPS BigLearn Workshop*, Lake Tahoe, NV, USA, 2013.
- [8] A. Nedić, A. Olshevsky, and M. G. Rabbat. Network topology and communication-computation tradeoffs in decentralized optimization. *Proceedings of the IEEE*, 106(5):953–976, 2018.
- [9] A. Nedić and A. Ozdaglar. Distributed subgradient methods for multi-agent optimization. *IEEE Trans. Automatic Control*, 54(1):48–61, 2009.
- [10] Y. Nesterov. *Introductory Lectures on Convex Optimization*. Springer, 2003.
- [11] S. Shahrampour and A. Jadbabaie. Distributed online optimization in dynamic environments using mirror descent. *IEEE Trans. Automatic Control*, 63(3):714–725, 2018.
- [12] A. Simonetto, A. Koppel, A. Mokhtari, G. Leus, and A. Ribeiro. Decentralized prediction-correction methods for networked time-varying convex optimization. *IEEE Trans. Automatic Control*, 62(11):5724–5738, 2017.
- [13] A. B. Taylor, J. M. Hendrickx, and F. Glineur. Performance estimation toolbox (pesto): automated worst-case analysis of first-order optimization methods. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 1278–1283. IEEE, 2017.
- [14] K. Yuan, Q. Ling, and W. Yin. On the convergence of decentralized gradient descent. *SIAM Journal on Optimization*, 26(3):1835–1854, 2016.

APPENDIX

In this Appendix we prove Theorem 2. Let $f^- := \sum_{i=1}^{n-1} f_i$ denote the sum of the common terms on the right-hand sides of the definitions of x^a and x^b in (14). Also let $x^- := \arg \min_x f^-(x)$ denote the minimizer of f^- . One approach to bound $\|x^a - x^b\|$ is to bound the distance between the minimizers of f^- and any function of the form $f^- + f_n$ where $f_n \in \mathcal{F}_{\alpha,\beta}$; i.e., to bound $\|x^* - x^-\|$ where $x^* = \arg \min_x (f^-(x) + f_n(x))$. Observe indeed that both x^a and x^b are particular cases of such x^* . Hence a bound on all $\|x^* - x^-\|$ together with the triangular inequality

$$\|x^a - x^b\| \leq \|x^a - x^-\| + \|x^b - x^-\| \quad (22)$$

gives us a bound on $\|x^a - x^b\|$. To this end, we have the following.

Theorem 5. *Let $f = f^- + f_n$, where f^- is as defined above and $f_n \in \mathcal{F}_{\alpha,\beta}$. Let $x^* = \arg \min_x f(x)$ and $x^- = \arg \min_x f^-(x)$. Then*

$$\|x^* - x^-\| \leq \min \left(2 + 2\sqrt{\kappa}, \frac{2\sqrt{\kappa} + \kappa}{\sqrt{n-1}}, \frac{2\kappa + \kappa^{3/2}}{n} \right). \quad (23)$$

The proof of Theorem 2 follows directly from the application of Theorem 5 to both terms in (22). Thus, all that remains is to prove Theorem 5. Before that, we derive the following intermediate result.

Proposition 5. *Let g_1 and g_2 be two functions such that g_1 is α_1 -strongly convex and β_1 -smooth and g_2 is α_2 -strongly convex and β_2 -smooth. Assume further that the minimizer x_1^* of g_1 lies in $B(0, R_1)$ and the minimizer x_2^* of g_2 lies in $B(0, R_2)$. Then $x^* = \arg \min_x (g_1(x) + g_2(x))$ satisfies*

$$\|x^* - x_1^*\| \leq \frac{\beta_2(R_1 + R_2)}{\alpha_1 + \alpha_2}.$$

Proof: It follows from the strong convexity and smoothness of g_1 and g_2 that

$$\begin{aligned} (\nabla g_1(x^*) - \nabla g_1(x_1^*))^\top (x^* - x_1^*) &\geq \frac{\alpha_1 \beta_1}{\alpha_1 + \beta_1} \|x^* - x_1^*\|^2 \\ &\quad + \frac{1}{\alpha_1 + \beta_1} \|\nabla g_1(x^*) - \nabla g_1(x_1^*)\|^2 \\ (\nabla g_2(x^*) - \nabla g_2(x_1^*))^\top (x^* - x_1^*) &\geq \frac{\alpha_2 \beta_2}{\alpha_2 + \beta_2} \|x^* - x_1^*\|^2 \\ &\quad + \frac{1}{\alpha_2 + \beta_2} \|\nabla g_2(x^*) - \nabla g_2(x_1^*)\|^2. \end{aligned}$$

Summing these inequalities while taking into account that $\nabla g_1(x_1^*) = 0$ and $\nabla(g_1 + g_2)(x^*) = 0$, we obtain

$$\begin{aligned} -\nabla g_2(x_1^*)^\top (x^* - x_1^*) &\geq \frac{1}{\alpha_1 + \beta_2} \|\nabla g_1(x^*) - \nabla g_1(x_1^*)\|^2 \\ &\quad + \frac{1}{\alpha_2 + \beta_2} \|\nabla g_2(x^*) - \nabla g_2(x_1^*)\|^2 \\ &\quad + \|x^* - x_1^*\|^2 \left(\frac{\alpha_1 \beta_1}{\alpha_1 + \beta_1} + \frac{\alpha_2 \beta_2}{\alpha_2 + \beta_2} \right). \end{aligned}$$

Recall that for an α -strongly convex function f , for any x and y we have $\|\nabla f(x) - \nabla f(y)\| \geq \alpha \|x - y\|$. Therefore,

$$\begin{aligned} &-\nabla g_2(x_1^*)^\top (x^* - x_1^*) \\ &\geq \|x^* - x_1^*\|^2 \left(\frac{\alpha_1 \beta_1}{\alpha_1 + \beta_1} + \frac{\alpha_2 \beta_2}{\alpha_2 + \beta_2} + \frac{\alpha_1^2}{(\alpha_1 + \beta_1)} + \frac{\alpha_2^2}{\alpha_2 + \beta_2} \right) \\ &= \|x^* - x_1^*\|^2 (\alpha_1 + \alpha_2). \end{aligned} \quad (24)$$

By the Cauchy-Schwartz inequality, $-\nabla g_2(x_1^*)^\top (x^* - x_1^*) \leq \|\nabla g_2(x_1^*)\| \|x^* - x_1^*\|$. In addition, using the smoothness of g_2 along with the assumptions that $x_1^* \in B(0, R_1)$ and $x_2^* \in B(0, R_2)$, we obtain that

$$\begin{aligned} \|\nabla g_2(x_1^*)\| &= \|\nabla g_2(x_1^*) - \nabla g_2(x_2^*)\| \\ &\leq \beta_2 \|x_1^* - x_2^*\| \\ &\leq \beta_2 (\|x_1^*\| + \|x_2^*\|) \\ &\leq \beta_2 (R_1 + R_2). \end{aligned}$$

Combining these in (24) gives

$$\beta_2 (R_1 + R_2) \|x^* - x_1^*\| \geq (\alpha_1 + \alpha_2) \|x^* - x_1^*\|^2,$$

and rearranging terms completes the proof.

Proof: [Proof of Theorem 5] First observe that f^- is $(n-1)\alpha$ -strongly convex and $(n-1)\beta$ -smooth by similar reasoning as for f in Proposition 1. By Theorem 1 we have that both x^* and x^- lie in $B(0, 1 + \sqrt{\kappa})$, and thus

$$\|x^* - x^-\| \leq \|x^*\| + \|x^-\| \leq 2 + 2\sqrt{\kappa}. \quad (25)$$

This bound scales as $\mathcal{O}(\sqrt{\kappa})$ but does not involve n .

To obtain a bound involving n , we leverage properties of the functions f^- and f_n . Recall that x_n^* denotes the minimizer of f_n . Since f_n is β -smooth,

$$f(x^-) = f^-(x^-) + f_n(x^-) \leq f^-(x^-) + \frac{\beta}{2} \|x^- - x_n^*\|^2. \quad (26)$$

Also, since f^- is $(n-1)\alpha$ -strongly convex and $\min_x f_n(x) = 0$ by assumption, it follows that for any $x \in \mathbb{R}^d$,

$$f(x) = f^-(x) + f_n(x) \quad (27)$$

$$\geq f^-(x) \quad (28)$$

$$\geq f^-(x^-) + \frac{(n-1)\alpha}{2} \|x - x^-\|^2. \quad (29)$$

Combining equations (26) and (29) leads to

$$\begin{aligned} f(x) - f(x^-) &\geq f^-(x^-) + \frac{(n-1)\alpha}{2} \|x - x^-\|^2 \\ &\quad - f^-(x^-) - \frac{\beta}{2} \|x^- - x_n^*\|^2 \\ &= \frac{(n-1)\alpha}{2} \|x - x^-\|^2 - \frac{\beta}{2} \|x^- - x_n^*\|^2. \end{aligned}$$

By Assumption 2, we know that $\|x_n^*\| \leq 1$, and applying Theorem 1 to f^- gives $\|x^-\| \leq 1 + \sqrt{\kappa}$. Therefore $\|x^- - x_n^*\| \leq 2 + \sqrt{\kappa}$, and so

$$\begin{aligned} f(x) - f(x^-) &\geq \frac{(n-1)\alpha}{2} \|x - x^-\|^2 - \frac{\beta}{2} (2 + \sqrt{\kappa})^2 \\ &= \frac{(n-1)\alpha}{2} \left(\|x - x^-\|^2 - \frac{\kappa(2 + \sqrt{\kappa})^2}{n-1} \right). \end{aligned}$$

A point x cannot be the minimizer of f if there is another point x^- for which $f(x) - f(x^-) > 0$, or equivalently, if

$$\|x - x^-\|^2 > \frac{\kappa(2 + \sqrt{\kappa})^2}{n-1},$$

and therefore

$$\|x^* - x^-\| \leq \frac{\kappa + 2\sqrt{\kappa}}{\sqrt{n-1}}. \quad (30)$$

Finally, applying Proposition 5 with $g_1 = f^-$ (hence $\alpha_1 = (n-1)\alpha$ and $R_1 = 1 + \sqrt{\kappa}$) and $g_2 = f_n$ (hence $\alpha_2 = \alpha$, $\beta_2 = \beta$, and $R_2 = 1$), we also obtain that

$$\|x^* - x^-\| \leq \frac{\beta(2 + \sqrt{\kappa})}{n\alpha} = \frac{2\kappa + \kappa^{3/2}}{n}. \quad (31)$$

Combining this with (25) and (30) gives the desired result.