

Setwise Coordinate Descent for Dual Asynchronous Decentralized Optimization

Marina Costantini, Nikolaos Liakopoulos, Panayotis Mertikopoulos, and Thrasyvoulos Spyropoulos

Abstract—In decentralized optimization over networks, synchronizing the updates of all nodes incurs significant communication overhead. For this reason, much of the recent literature has focused on the analysis and design of asynchronous optimization algorithms where nodes can activate anytime and contact a single neighbor to complete an iteration together. However, most works assume that the neighbor selection is done at random based on a fixed probability distribution (e.g., uniform), a choice that ignores the optimization landscape at the moment of activation. Instead, in this work we introduce an optimization-aware selection rule that chooses the neighbor providing the highest dual cost improvement (a quantity related to a dualization of the problem based on consensus). This scheme is related to the coordinate descent (CD) method with the Gauss-Southwell (GS) rule for coordinate updates; in our setting however, only a subset of coordinates is accessible at each iteration (because each node can communicate only with its neighbors), so the existing literature on GS methods does not apply. To overcome this difficulty, we develop a new analytical framework for smooth and strongly convex functions that covers our new class of *setwise CD algorithms*—a class that applies to both decentralized and parallel distributed computing scenarios— and we show that the proposed setwise GS rule can speed up the convergence in terms of iterations by a factor equal to the size of the largest coordinate set. We analyze extensions of these algorithms that exploit the knowledge of the smoothness constants when available and otherwise propose an algorithm to estimate these constants. Finally, we validate our theoretical results through extensive simulations.

Index Terms—convex optimization, coordinate descent, decentralized optimization, distributed machine learning, distributed optimization, multi-agent optimization, optimization over networks.

I. INTRODUCTION

Many timely applications require solving optimization problems over a network where nodes can only communicate with their direct neighbors. This may be due to the need of distributing storage and computation loads (e.g. training large machine learning models [1]), or to avoid transferring data that is naturally collected in a decentralized manner, either due to the communication costs or to privacy reasons (e.g. sensor networks [2], edge computing [3], and precision medicine [4]).

M. Costantini is with EURECOM, Sophia Antipolis, France (costanti@eurecom.fr).

N. Liakopoulos is with Amazon, Luxembourg City, Luxembourg (nliako@amazon.lu).

P. Mertikopoulos is with Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LIG, 38000 Grenoble, France (panayotis.mertikopoulos@imag.fr).

T. Spyropoulos is with EURECOM & the Technical University of Crete, Chania, Greece (spyropou@eurecom.fr).

Specifically, we consider a setting where the nodes want to solve the decentralized optimization problem

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

where each local function f_i is known only by node i and nodes can exchange optimization values (parameters, gradients) but *not* the local functions themselves. We represent the communication network as a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with $n = |\mathcal{V}| < \infty$ nodes (agents) and $E = |\mathcal{E}|$ edges, which are the links used by the nodes to communicate with their neighbors.

Existing methods to solve this problem usually assign to each node a local variable θ_i and execute updates that interleave a local gradient step at the nodes followed by an averaging step (usually carried out by a doubly-stochastic matrix) that aggregates the update of the node with those of its neighbors [5]–[10].

Alternatively, the consensus constraint can be stated explicitly between node pairs connected by an edge:

$$\underset{\theta_1, \dots, \theta_n \in \mathbb{R}^d}{\text{minimize}} \quad \sum_{i=1}^n f_i(\theta_i) \quad (2a)$$

$$\text{subject to} \quad \theta_i = \theta_j \quad \forall (i, j) \equiv \ell \in \mathcal{E}, \quad (2b)$$

where $\ell \equiv (i, j)$ indicates that edge ℓ links nodes i and j . This reformulation allows solving the decentralized optimization problem through its dual problem. Interestingly, dual decentralized algorithms are among the very few that have been shown to achieve optimal convergence rates [11]–[13] when using acceleration. However, these algorithms are *synchronous*, in the sense that they require coordinating the updates of all nodes at each iteration. In contrast, by exploiting connections with coordinate descent theory naturally derived from the dual formulation, here we propose an alternative way to speed up convergence in dual decentralized optimization that does *not* require network-wide synchronicity.

In our setup, nodes activate anytime at random and select one of their neighbors to make an update together. Methods with such minimal coordination requirements avoid incurring extra costs of synchronization that may also slow down convergence [14]–[18]. However, most of these works assume that when a node activates, it simply selects the neighbor to contact randomly, based on a predefined probability distribution. This approach overlooks the possibility of letting nodes *choose* the neighbor to contact taking into account the optimization landscape at the time of activation. Therefore, here we depart from

the probabilistic choice and ask: *can nodes pick the neighbor smartly to make the optimization converge faster?*

In this paper, we give an affirmative answer and propose algorithms that achieve this by solving the dual problem of (2). In the dual formulation, there is one dual variable $\lambda_\ell \in \mathbb{R}^d$ per constraint $\theta_i = \theta_j$, hence each dual variable can be associated with an edge ℓ in the graph. Our algorithms let an activated node i contact a neighbor j so that together they update their shared variable λ_ℓ , $\ell \equiv (i, j)$ with a gradient step. In particular, we propose to select the neighbor j such that the updated λ_ℓ is the one *whose directional gradient for the dual function is the largest*, and thus the one that provides the greatest cost improvement at that iteration. Beyond the preliminary version of this work that appeared in [19], such dual greedy choice has not yet been considered in the literature.

Interestingly, the above protocol where a node activates and selects a λ_ℓ to update can be seen as applying the coordinate descent (CD) method [20] to solve the dual problem of (2), with the following key difference: unlike standard CD methods, where *any* of the coordinates may be updated, now *only a small subset of coordinates are accessible at each step*, which are the coordinates associated with the edges connected to the node activated. Moreover, our proposal of updating the λ_ℓ with the largest gradient is similar to the Gauss-Southwell (GS) rule [21], but applied *only* to the parameters accessible by the activated node.

We name such protocols *setwise CD* algorithms, and we analyze a number of possibilities for the coordinate (or equivalently, neighbor) choice: random sampling, local GS, and two extensions that take into account the smoothness constants of the dual function. Compared to standard CD literature, three difficulties complicate the analysis and constitute the base of our contributions: (i) for arbitrary graphs, the dual problem of (2) has an objective function that is *not* strongly convex, even if the primal functions f_i are strongly convex, (ii) the fact that the GS rule is applied to a few coordinates prevents the use of standard norms to obtain the linear rate, as commonly done for CD methods [20]–[22], and (iii) the coordinate sets are overlapping (i.e. non-disjoint), which makes the problem even harder.

Our results also apply to the (primal) parallel distributed setting where multiple workers modify different sets of coordinates of the parameter vector, which is stored in a server accessible by all workers [23]–[25]. In particular, we show that for this setting the GS selection attains the maximum speedup promised by the theory (see Theorem 2).

Our contributions can be summarized as follows:

- We introduce the class of *setwise CD* algorithms for asynchronous optimization applicable to both the decentralized and the parallel distributed settings.
- We prove linear convergence rates for smooth and strongly convex f_i and four coordinate (equivalently, neighbor) choices: random uniform, GS, and their variants when the coordinate smoothness constants are known.
- For the cases when these constants are not known, we propose an algorithm based on backtracking [20] for estimating them online, and show that for certain problems this method achieves faster per-iteration convergence than

the exact knowledge of these constants.

- To obtain the rates of all considered algorithms, we prove strong convexity in unique norms derived from our dual coordinate updates. For the algorithms using the GS rule in particular, we deviate significantly from standard proofs and we resort to carefully-crafted support norms for each algorithm that allow us to bypass the difficulties introduced by the setwise choice.
- We show that the speedup in terms of number of iterations of GS selection with respect to random uniform can be up to N_{\max} (the size of the largest coordinate set).
- We prove that the versions accounting for coordinate smoothness are provably faster than those that do not account for these constants. In particular, we show that the algorithm exploiting both the GS rule *and* the smoothness knowledge is provably faster than all others.
- We support all our results with thorough simulations¹.

II. RELATED WORK

A number of algorithms have been proposed to solve (1) asynchronously, where we consider the edge-asynchronous model in Assumption 2 (other models have been considered in the literature, e.g. [26], [27]). These algorithms require transmitting a fixed number of vectors (either 1, 2 or 4 in the references below) over a single edge to complete an iteration. In [18], the activated node chooses a neighbor uniformly at random and both nodes average their primal local values. In [14] the authors adapted the ADMM algorithm to the decentralized setting, but it was the ADMM of [15] the first one shown to converge at the same rate as the centralized ADMM. The algorithm of [16] tracks the average gradients to converge to the exact optimum instead of just a neighborhood around it, as many algorithms back then. The algorithm of [28] can be used on top of directed graphs, which impose additional challenges. In [29], [30] a token is passed from node to node following a Markov chain and the node who receives the token completes an iteration. A key novelty of our scheme, compared to the works above, is that we consider the possibility of letting the nodes *choose smartly* the neighbor to contact in order to make convergence faster.

Work [31] is, to the best of our knowledge, the only work similarly considering smart neighbor selection. The authors propose Max-gossip, a version of the (primal) algorithm in [5] where the activated node averages its local parameter with that of the neighbor with whom the parameter difference is the largest. They show that the algorithm converges sublinearly to the optimum (for convex functions), and show in numerical simulations that it outperforms random neighbor selection. In contrast, here we propose dual algorithms for which we show linear convergence rates (for smooth and strongly convex functions), and most importantly, (i) we *prove analytically* that either applying the GS rule and/or using the Lipschitz information achieves faster convergence than random neighbor sampling, and (ii) we *quantify* the magnitude of the gains.

¹The code to reproduce the results is available at https://github.com/m-costantini/Set-wise_Coordinate_Descent/.

To achieve this, we apply coordinate descent in the dual, similar to a few related works. In [32] the authors propose a dual proximal algorithm for non-smooth functions whose edge-based variant, if applied to smooth functions, reduces to the SU-CD algorithm presented here. The ESDACD algorithm of [12] applies accelerated coordinate descent to the dual problem, achieving faster convergence at the expense of sacrificing complete asynchronicity (nodes need to agree on a sequence of node activations beforehand). Still, none of these works considered smart, non-randomized neighbor selection,

Finally, as mentioned earlier, our work relates to standard CD literature. In particular, our theorems extend the results in [21], where the GS rule was shown to be up to d times faster than uniform sampling for $f: \mathbb{R}^d \rightarrow \mathbb{R}$, to the case where this choice is constrained to a subset of the coordinates only, sets have different sizes, each coordinate belongs to exactly two sets, and sets activate uniformly at random. As we explain in Sections IV and V, these considerations bring important new challenges with respect to the standard single-machine CD algorithms. Furthermore, our algorithms are not only applicable to the decentralized case but also to parallel distributed settings such as [23]–[25]. For the latter, [33] also analyzed the GS applied to coordinate subsets, but their sets are disjoint, accessible by any worker, and they do not quantify the speedup of the method with respect to random uniform sampling.

III. DUAL FORMULATION

Here we define the notation, obtain the dual problem of (2), and analyze the properties of the dual objective function.

Assumption 1 (Smoothness and strong convexity). The f_i are M_i -smooth and μ_i -strongly convex, i.e. there exist finite constants $M_i \geq \mu_i > 0$, $i \in [n]$ such that:

$$\begin{aligned} f_i(y) &\leq f_i(x) + \langle \nabla f_i(x), y - x \rangle + (M_i/2) \|y - x\|_2^2 \\ f_i(y) &\geq f_i(x) + \langle \nabla f_i(x), y - x \rangle + (\mu_i/2) \|y - x\|_2^2. \end{aligned}$$

We define the concatenated primal and dual variables $\theta = [\theta_1^T, \dots, \theta_n^T]^T \in \mathbb{R}^{nd}$ and $\lambda = [\lambda_1^T, \dots, \lambda_E^T]^T \in \mathbb{R}^{Ed}$, respectively. The graph's incidence matrix $A \in \mathbb{R}^{n \times E}$ has exactly one 1 and one -1 per column ℓ , in the rows corresponding to nodes $i, j: \ell \equiv (i, j)$, and zeros elsewhere (the choice of sign for each node is irrelevant). We call $u_i \in \mathbb{R}^n$ the vector that has 1 in entry i and 0 elsewhere; we define $e_\ell \in \mathbb{R}^E$ analogously. We use $k \in [K]$ to indicate $k = 1, \dots, K$. Vectors $\mathbf{1}$ and $\mathbf{0}$ are respectively the all-one and all-zero vectors, and I_d is the $d \times d$ identity matrix. Finally, in order to use matrix operations in the equations below for some operations, we define the block arrays $\Lambda = A \otimes I_d \in \mathbb{R}^{nd \times Ed}$ and $U_i = u_i \otimes I_d \in \mathbb{R}^{nd \times d}$, where \otimes is the Kronecker product. This operation generates arrays analogous to A and u_i where the original entries 1, -1, and 0 have been replaced by I_d , $-I_d$, and the all-zero $d \times d$ matrix, respectively.

We can rewrite now (2b) as $\Lambda^T \theta = \mathbf{0}$, and the node variables

as $\theta_i = U_i^T \theta$. The minimum value of (2) satisfies:

$$\begin{aligned} \inf_{\theta: \Lambda^T \theta = \mathbf{0}} \sum_{i=1}^n f_i(U_i^T \theta) &\stackrel{(a)}{=} \inf_{\theta} \sup_{\lambda} \left[\sum_{i=1}^n f_i(U_i^T \theta) - \lambda^T \Lambda^T \theta \right] \\ &\stackrel{(b)}{=} \sup_{\lambda} \inf_{\theta} \left[\sum_{i=1}^n f_i(U_i^T \theta) - \lambda^T \Lambda^T \theta \right] \\ &= - \inf_{\lambda} \sup_{\theta} \sum_{i=1}^n \left[(U_i^T \Lambda \lambda)^T U_i^T \theta - f_i(U_i^T \theta) \right] \\ &= - \inf_{\lambda} \sum_{i=1}^n f_i^*(U_i^T \Lambda \lambda) \triangleq - \inf_{\lambda} F(\lambda), \end{aligned} \quad (3)$$

where (a) holds due to Lagrange duality and (b) holds by strong duality (see e.g. Section 5.4 in [34]). Functions f_i^* are the Fenchel conjugates of the f_i , and are defined as

$$f_i^*(y) = \sup_{x \in \mathbb{R}^d} (y^T x - f_i(x)).$$

Our setwise CD algorithms converge to the optimal solution of (2) by solving (3). In particular, they update a single dual variable λ_ℓ , $\ell \in [E]$ at each iteration and converge to some minimum value λ^* of $F(\lambda)$.

Since $\sum_{i=1}^n f_i(U_i^T \theta)$ in (2a) is M_{\max} -smooth and μ_{\min} -strongly convex in θ , with $M_{\max} = \max_i M_i$ and $\mu_{\min} = \min_i \mu_i$, function F is L -smooth with $L = \frac{\gamma_{\max}}{\mu_{\min}}$, where γ_{\max} is the largest eigenvalue of $\Lambda^+ \Lambda$ (Section 4 in [13]). We call γ_{\min}^+ the smallest non-zero eigenvalue² of $\Lambda^+ \Lambda$.

However, as shown next, function F is *not* strongly convex in the standard L2 norm, which is the property that usually facilitates obtaining linear rates in optimization literature.

Lemma 1. F is not strongly convex in $\|\cdot\|_2$.

Proof. Since Λ does not have full column rank in the general case (i.e., unless the graph is a tree), there exist $w \in \mathbb{R}^{Ed}$ such that $w \neq \mathbf{0}$ and $F(\lambda) = F(\lambda + tw) \forall t \in \mathbb{R}$. \square

Nevertheless, we can still show linear rates for the setwise CD algorithms using the following result.

Lemma 2 (Appendix C of [12]). F is σ_A -strongly convex in the semi-norm $\|x\|_A \triangleq (x^T \Lambda^+ \Lambda x)^{\frac{1}{2}}$, with $\sigma_A = \frac{\gamma_{\min}^+}{M_{\max}}$.

Above, Λ^+ denotes the pseudo-inverse of Λ . A key fact for the proofs in the next section is that matrix $\Lambda^+ \Lambda$ is a projector onto $\text{range}(\Lambda^T)$, the column space of Λ^T .

To simplify the notation, in what follows we assume that $d = 1$, so that $\Lambda = A$, $U_i = u_i$, and the gradient $\nabla_{\ell} F(\lambda) = \frac{\partial F(\lambda)}{\partial \lambda_{\ell}}$ of $F(\lambda)$ in the direction of λ_{ℓ} is a scalar. In Section VI-C we discuss how to adapt our proofs to the case $d > 1$.

IV. SETWISE COORDINATE DESCENT ALGORITHMS

In this section, we present the *setwise CD* algorithms, which can solve generic convex problems (and (3) in particular) optimally and asynchronously. For solving (3) we rely on:

Assumption 2 (Asynchronous model). Nodes activate independently at random (potentially with different frequencies) and choose a neighbor with whom to complete an update.

²The “+” stresses that γ_{\min}^+ is the smallest *strictly positive* eigenvalue.

Multiple *pairs* of nodes may complete their updates in parallel, provided that each node updates with a single neighbor.

Here we analyze two possibilities for the coordinate choice within the accessible coordinate subset: (i) sampling uniformly at random (SU-CD), and (ii) applying the GS rule (SGS-CD). Their extensions when the coordinate-wise Lipschitz constants are known (or can be estimated) are analyzed in Section V.

If coordinate ℓ is updated at iteration k and assuming $d = 1$, the standard CD update applied to $F(\lambda)$ is [20]:

$$\lambda^{k+1} = \lambda^k - \eta^k \nabla_{\ell} F(\lambda^k) e_{\ell}, \quad (4)$$

where η^k is the stepsize. Since F is L -smooth, choosing $\eta^k = 1/L \forall k$ guarantees descent at each iteration [21]:

$$F(\lambda^{k+1}) \leq F(\lambda^k) - \frac{1}{2L} \left(\nabla_{\ell} F(\lambda^k) \right)^2. \quad (5)$$

Eq. (5) will be the departure point to prove the linear convergence rates of SU-CD and SGS-CD. We remark that since F is smooth and convex, the fixed stepsize choice above leads to exact convergence to an optimum of the function [20].

We now define formally the setwise CD algorithms.

Definition 1 (Setwise CD algorithm). In a setwise CD algorithm, every coordinate $\ell \in [E]$ is assigned to (potentially multiple) sets $\mathcal{S}_i, i \in [n]$, such that all coordinates belong to at least one set. At any time, a set \mathcal{S}_i may activate with uniform probability among the i ; a setwise CD algorithm then chooses a single coordinate $\ell \in \mathcal{S}_i$ to update using (4).

The next remark shows how the decentralized problem (2) can be solved asynchronously with setwise CD algorithms.

Remark 1. By letting (i) the E coordinates³ in Definition 1 correspond to the dual variables $\lambda_{\ell}, \ell \in [E]$, and (ii) the $\mathcal{S}_i, i \in [n]$ be the sets of dual variables corresponding to the edges that are connected to each node i , nodes can run a setwise CD algorithm to solve (3) (and thus, also (2)) asynchronously.

Furthermore, the setwise CD algorithms can also be used in the parallel distributed setting, as explained below.

Remark 2. In a parallel distributed setting, where the parameter vector $x \in \mathbb{R}^E$ is stored⁴ in a shared server and can be modified by n workers, the (primal) optimization can be done with setwise CD algorithms by letting each worker $i \in [n]$ modify a subset \mathcal{S}_i of coordinates such that when the worker activates it chooses one coordinate $\ell \in \mathcal{S}_i$ to modify, and all coordinates $\ell \in [E]$ belong to at least one set $\mathcal{S}_i, i \in [n]$.

We also note that, while in the parallel distributed setting each coordinate may belong to any number of sets in $[1, n]$, our results apply to the case where all coordinates belong to *exactly* two sets (i.e. they can be modified by two workers), since our analysis is driven by the decentralized setting.

In light of Remark 1, in the following we illustrate the steps that should be performed by the nodes to run the setwise CD

³If $d > 1$, the standard CD terminology calls each λ_{ℓ} a “block coordinate”, i.e. a vector of d coordinates out of the $E \cdot d$ of $F: \mathbb{R}^{E \cdot d} \rightarrow \mathbb{R}$.

⁴Also in this case we may actually have $x \in \mathbb{R}^{d \cdot E}$ and each worker may update a block coordinate $x_{\ell} \in \mathbb{R}^d$ at each iteration.

TABLE I: Set-related definitions

\mathcal{S}_i	Set of edges connected to node i
\mathcal{N}_i	Set of neighbors of node i
N_i	Degree of node i , i.e. $N_i = \mathcal{S}_i = \mathcal{N}_i $
N_{\max}	Maximum degree in the network, i.e. $\max_i N_i$
T_i	Selector matrix of set \mathcal{S}_i (see Definition 2)
\mathcal{S}'_i	Subset $\mathcal{S}'_i \subseteq \mathcal{S}_i$ such that $\mathcal{S}'_i \cap \mathcal{S}'_j = \emptyset$ if $i \neq j$
T'_i	Selector matrix of set \mathcal{S}'_i
$\overline{\mathcal{S}}'_i$	Complement set of \mathcal{S}'_i such that $\overline{\mathcal{S}}'_i = \mathcal{S}_i \setminus \mathcal{S}'_i$
\overline{T}'_i	Selector matrix of set $\overline{\mathcal{S}}'_i$

algorithms to find an optimal value λ^* . We first note that the gradient of F in the direction⁵ of λ_{ℓ} for $\ell \equiv (i, j)$ is

$$\nabla_{\ell} F(\lambda) = A_{i\ell} \nabla f_i^*(u_i^T A \lambda) + A_{j\ell} \nabla f_j^*(u_j^T A \lambda). \quad (6)$$

The nodes can use (4) and (6) to update the λ_{ℓ} of the edges they are connected to as follows: each node i keeps in memory the current values of $\lambda_{\ell}, \ell \in \mathcal{S}_i$, which are needed to compute $\nabla f_i^*(u_i^T A \lambda)$. Then, when edge $\ell \equiv (i, j)$ is updated (either because node i activated and contacted j , or vice versa), both i and j compute their respective terms in the right-hand side of (6) and exchange them through their link. Finally, both nodes compute (6) and update their copy of λ_{ℓ} applying (4).

Algorithms 1 and 2 below detail these steps for SU-CD and SGS-CD, respectively. We have used \mathcal{N}_i to indicate the set of neighbors of node i (note that $\mathcal{S}_i = \{\ell : \ell \equiv (i, j), j \in \mathcal{N}_i\}$). Table I shows this and other set-related notation we will use.

We now proceed to describe the SU-CD and SGS-CD algorithms in detail, and prove their linear convergence rates. We remark that the analyses that follow are given in terms of number of iterations. See Section VII-E for a discussion considering communication and computation complexity.

A. Setwise Uniform CD (SU-CD)

In SU-CD, the activated node chooses the neighbor uniformly at random, as shown in Algorithm 1. We can compute the per-iteration progress of SU-CD taking expectation in (5):

$$\begin{aligned} \mathbb{E}[F(\lambda^{k+1}) | \lambda^k] &\leq F(\lambda^k) - \frac{1}{2L} \mathbb{E} \left[\left(\nabla_{\ell} F(\lambda^k) \right)^2 | \lambda^k \right] \\ &= F(\lambda^k) - \frac{1}{2Ln} \sum_{i=1}^n \frac{1}{N_i} \sum_{\ell \in \mathcal{S}_i} \left(\nabla_{\ell} F(\lambda^k) \right)^2 \\ &\leq F(\lambda^k) - \frac{1}{LnN_{\max}} \|\nabla F(\lambda^k)\|_2^2 \end{aligned} \quad (7)$$

where $N_i = |\mathcal{S}_i|$, $N_{\max} = \max_i N_i$, and the factor 2 in the denominator disappears because each coordinate $\ell \equiv (i, j)$ is counted twice in the sums of \mathcal{S}_i and \mathcal{S}_j , respectively.

The standard procedure to show the linear convergence of CD in the single-machine case is to lower-bound $\|\nabla F(\lambda)\|_2^2$ using the strong convexity of the function [20], [21]. However, since F is *not* strongly convex (Lemma 1), we cannot apply this procedure to get the linear rate of SU-CD.

We can, however, use F 's strong convexity in $\|\cdot\|_A$ instead (Lemma 2). The next result gives the core of the proof.

⁵This is equivalent to saying “the ℓ -th (block) entry of the gradient ∇F ”.

Algorithm 1 Setwise Uniform CD (SU-CD)

- 1: **Input:** Functions f_i , step η , incidence matrix A , graph \mathcal{G}
- 2: Initialize $\theta_i^0, i = 1, \dots, n$ and $\lambda_\ell^0, \ell = 1, \dots, E$
- 3: **for** $k = 1, 2, \dots$ **do**
- 4: Sample activated node $i \in \{1, \dots, n\}$ uniformly
- 5: Node i picks neighbor $j \leftarrow \mathcal{U}\{h : h \in \mathcal{N}_i\}$
- 6: Node i computes $\nabla f_i^*(u_j^T A \lambda)$ and sends it to j
- 7: Node j computes $\nabla f_j^*(u_i^T A \lambda)$ and sends it to i
- 8: Nodes i, j : $(i, j) \equiv \ell$ use (6) to update their local copies of λ_ℓ by $\lambda_\ell^k \leftarrow \lambda_\ell^{k-1} - \eta \nabla_\ell F(\lambda)$
- 9: $\lambda_m^k \leftarrow \lambda_m^{k-1} \forall$ edges $m \neq \ell$

Lemma 3. It holds that

$$\|\nabla F(\lambda)\|_2 = \|\nabla F(\lambda)\|_A = \|\nabla F(\lambda)\|_A^*, \quad (8)$$

where $\|\cdot\|_A^*$ is the dual norm of $\|\cdot\|_A$, defined as (e.g. [34])

$$\|z\|_A^* = \sup_{x \in \mathbb{R}^d} \left\{ z^T x \mid \|x\|_A \leq 1 \right\}. \quad (9)$$

Proof. Note that $\forall w \neq \mathbf{0}$ such that $F(\lambda + tw) = F(\lambda) \forall t$, it holds that $w^T \nabla F(\lambda) = 0$ and thus $\nabla F(\lambda) \in \text{range}(A^T)$. This means that $A^+ A \nabla F(\lambda) = I_E \nabla F(\lambda)$, and therefore it holds that $\|\nabla F(\lambda)\|_A = \|\nabla F(\lambda)\|_2$. Finally, since the dual norm of the L2 norm is the L2 norm itself, we have that also $\|\nabla F(\lambda)\|_A^* = \|\nabla F(\lambda)\|_2$, which gives the result. \square

We now use Lemma 3 to prove the linear rate of SU-CD. Note that since λ^* is an optimum of F , convergence is exact. This is valid not only for SU-CD, but for all other algorithms presented in this paper.

Theorem 1 (Rate of SU-CD). SU-CD converges as

$$\mathbb{E}[F(\lambda^{k+1}) \mid \lambda^k] - F(\lambda^*) \leq \left(1 - \frac{2\sigma_A}{LnN_{\max}}\right) [F(\lambda^k) - F(\lambda^*)].$$

Proof. Since $F(\lambda)$ is strongly convex in $\|\cdot\|_A$ with strong convexity constant σ_A (Lemma 2), it holds

$$F(y) \geq F(x) + \langle \nabla F(x), y - x \rangle + \frac{\sigma_A}{2} \|y - x\|_A^2.$$

Minimizing both sides with respect to y as in [21] we get

$$F(x^*) \geq F(x) - \frac{1}{2\sigma_A} (\|\nabla F(x)\|_A^*)^2, \quad (10)$$

and rearranging terms we obtain the lower bound $(\|\nabla F(x)\|_A^*)^2 \geq 2\sigma_A(F(x) - F(x^*))$.

Finally, we can use Lemma 3 to replace $\|\nabla F(x)\|_2^2$ with $(\|\nabla F(x)\|_A^*)^2$ in (7), and use the lower bound on $(\|\nabla F(x)\|_A^*)^2$ to get the result. \square

Note that vector λ has $\frac{1}{2} \sum_i N_i = E \leq \frac{nN_{\max}}{2}$ coordinates, where the inequality holds with equality for regular graphs. We make the following remark.

Remark 3. If \mathcal{G} is regular, the linear convergence rate of SU-CD is $\frac{\sigma_A}{LE}$, which matches the rate of single-machine uniform CD for strongly convex functions [20], [21], with the only difference that now the strong convexity constant σ_A is defined over norm $\|\cdot\|_A$.

Algorithm 2 Setwise Gauss-Southwell CD (SGS-CD)

- 1: **Input:** Functions f_i , step η , incidence matrix A , graph \mathcal{G}
- 2: Initialize $\theta_i^0, i = 1, \dots, n$ and $\lambda_\ell^0, \ell = 1, \dots, E$
- 3: **for** $k = 1, 2, \dots$ **do**
- 4: Sample activated node $i \in \{1, \dots, n\}$ uniformly
- 5: All $h \in \mathcal{N}_i$ send their $\nabla f_h^*(u_h^T A \lambda)$ to node i
- 6: Node i computes $\nabla_\ell F(\lambda) \forall \ell \in \mathcal{S}_i$ (equivalently, $\forall h \in \mathcal{N}_i$ with (6) using the received $\nabla f_h^*(u_h^T A \lambda)$)
- 7: Node i selects $j \leftarrow \max_{h \in \mathcal{N}_i} |\nabla_\ell F(\lambda)|, \ell \equiv (i, h)$
- 8: Node i sends $\nabla f_i^*(u_j^T A \lambda)$ to j
- 9: Nodes i, j : $(i, j) \equiv \ell$ use (6) to update their local copies of λ_ℓ by $\lambda_\ell^k \leftarrow \lambda_\ell^{k-1} - \eta \nabla_\ell F(\lambda)$
- 10: Nodes i and j recompute their $\nabla f_h^*(u_h^T A \lambda), h \in \{i, j\}$
- 11: $\lambda_m^k \leftarrow \lambda_m^{k-1} \forall$ edges $m \neq \ell$

In the next section we analyze SGS-CD and show that its convergence rate can be up to N_{\max} times that of SU-CD.

B. Setwise Gauss-Southwell CD (SGS-CD)

In SGS-CD, as shown in Algorithm 2, the activated node i selects the neighbor j to contact applying the GS rule within the edges in \mathcal{S}_i :

$$\ell = \underset{m \in \mathcal{S}_i}{\text{argmax}} (\nabla_m F(\lambda))^2,$$

and then j is the neighbor that satisfies $\ell \equiv (i, j)$. In order to make this choice, all nodes $h \in \mathcal{N}_i$ must send their $\nabla f_h^*(u_h^T A \lambda)$ to node i (line 5 in Algorithm 2). Remark that the ∇f_h^* only need to be recomputed by the nodes that complete the update (line 10), and not by all neighbors originally pinged by the activated node. This keeps the computation complexity of SGS-CD equal to that of SU-CD, and only communication has increased. We discuss this fact in Section VIII.

To obtain the convergence rate of SGS-CD we will follow the steps taken for SU-CD in the proof of Theorem 1. As done for SU-CD, we start by computing the per-iteration progress taking expectation in (5) for SGS-CD:

$$\mathbb{E}[F(\lambda^{k+1}) \mid \lambda^k] \leq F(\lambda^k) - \frac{1}{2Ln} \sum_{i=1}^n \max_{\ell \in \mathcal{S}_i} (\nabla_\ell F(\lambda^k))^2. \quad (11)$$

Given this per-iteration progress, to proceed as we did for SU-CD we need to show that (i) the sum on the right-hand side of (11) defines a norm, and (ii) strong convexity holds in its dual norm. We start by defining the selector matrices T_i , which will significantly simplify notation.

Definition 2 (Selector matrices). The selector matrices $T_i \in \{0, 1\}^{N_i \times E}$, $i = 1, \dots, n$ select the coordinates of a vector in \mathbb{R}^E that belong to set \mathcal{S}_i . Note that any vertical stack of the unitary vectors $\{e_\ell^T\}_{\ell \in \mathcal{S}_i}$ gives a valid T_i .

We can now show that the sum in (11) is a (squared) norm. Since the operation involves applying $\max(\cdot)$ within each set \mathcal{S}_i , we will denote this norm $\|x\|_{\text{SM}}$, where the subscript SM stands for ‘‘Set-Max’’.

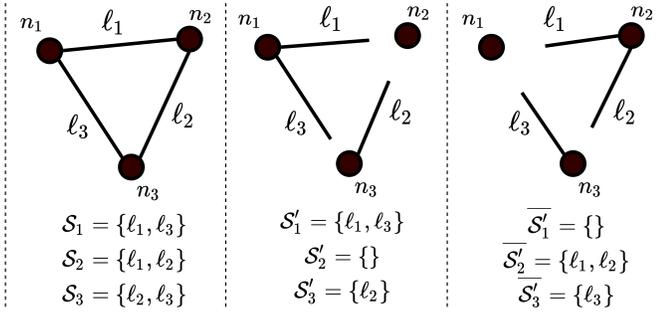


Fig. 1: Example of sets \mathcal{S}_i and one possibility for \mathcal{S}'_i and $\overline{\mathcal{S}'_i}$

Lemma 4. The function $\|x\|_{\text{SM}} \triangleq \sqrt{\sum_{i=1}^n \|T_i x\|_{\infty}^2} = \sqrt{\sum_{i=1}^n \max_{j \in \mathcal{S}_i} x_j^2}$ is a norm in \mathbb{R}^E .

Proof. Using $\max_{j \in \mathcal{S}_i} (x_j^2 + y_j^2) \leq \max_{j \in \mathcal{S}_i} x_j^2 + \max_{j \in \mathcal{S}_i} y_j^2$ and $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ we can show that $\|\cdot\|_{\text{SM}}$ satisfies the triangle inequality. It is straightforward to show that $\|\alpha x\|_{\text{SM}} = |\alpha| \|x\|_{\text{SM}}$ and $\|x\|_{\text{SM}} = 0$ if and only if $x = \mathbf{0}$. \square

Following the proof of Theorem 1, we would like to show that F is strongly convex in the dual norm $\|\cdot\|_{\text{SM}}^*$. Furthermore, we would like to compare the strong convexity constant σ_{SM} with σ_A to quantify the speedup of SGS-CD with respect to SU-CD. It turns out, though, that computing $\|\cdot\|_{\text{SM}}^*$ is not easy at all; the main difficulty stems from the fact that the sets \mathcal{S}_i are overlapping (or non-disjoint), since each coordinate $\ell \equiv (i, j)$ belongs to both \mathcal{S}_i and \mathcal{S}_j . The first scheme in Figure 1 illustrates this fact for the 3-node clique.

To circumvent this issue, we define a new norm $\|\cdot\|_{\text{SMNO}}^*$ (“Set-Max Non-Overlapping”) that we can directly relate to $\|\cdot\|_{\text{SM}}^*$ (Lemma 5) and whose value we can compute explicitly (Lemma 6), which will later allow us to relate the three strong convexity constants σ_{SM} , σ_{SMNO} , and σ_A (Theorem 2).

Definition 3 (Norm $\|\cdot\|_{\text{SMNO}}^*$). We assume that each coordinate $\ell \equiv (i, j)$ is assigned to only one of the sets $\mathcal{S}'_i \subseteq \mathcal{S}_i$ or $\mathcal{S}'_j \subseteq \mathcal{S}_j$, such that the new sets $\{\mathcal{S}'_i\}_{i=1}^n$ are non-overlapping (some sets can be empty), and all coordinates ℓ belong to *exactly one* set in $\{\mathcal{S}'_i\}$. We name the selector matrices of these new sets T'_i , so that each possible choice of $\{\mathcal{S}'_i\}$ defines a different set $\{T'_i\}$. Then, we define

$$\|z\|_{\text{SMNO}}^* = \sup_x \left\{ z^T x \mid \sqrt{\sum_{i=1}^n \|T'_i x\|_{\infty}^2} \leq 1 \right\}, \quad (12)$$

with the choice of non-overlapping sets

$$\{T_i^*\} = \arg \max_{\{T'_i\}} \sum_{i=1}^n \|T'_i x\|_{\infty}^2. \quad (13)$$

Note that the maximizations in (12) and (13) are coupled. We denote the value of x that attains (12) by x_{SMNO}^* .

The definition of sets \mathcal{S}'_i corresponds to assigning each edge ℓ to one of the two nodes at its endpoints, as illustrated in the second scheme of Figure 1. Therefore, for each possible pair $(\{\mathcal{S}'_h\}, \{T'_h\})$, $h \in [n]$ we can define a complementary pair $(\{\overline{\mathcal{S}'_h}\}, \{\overline{T'_h}\})$ such that if $\ell \equiv (i, j)$ was assigned to \mathcal{S}'_i in $\{\mathcal{S}'_h\}$,

then it is assigned to $\overline{\mathcal{S}'_j}$ in $\{\overline{\mathcal{S}'_h}\}$. This corresponds to assigning ℓ to the opposite endpoint (node) to the one originally chosen, as shown in the third scheme of Figure 1. With these definitions, it holds (potentially with some permutation of the rows):

$$T_i = \begin{bmatrix} T'_i \\ \overline{T'_i} \end{bmatrix} = \begin{bmatrix} T'_i \\ \mathbf{0} \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ \overline{T'_i} \end{bmatrix}, \quad i = 1, \dots, n.$$

We remark that the equality above holds for *any* $\{T'_i\}$ corresponding to a feasible assignment $\{\mathcal{S}'_i\}$, and in particular it holds for $(\{\mathcal{S}^*\}, \{T_i^*\})$. This fact is used in the proof of the following lemma, which relates norms $\|\cdot\|_{\text{SM}}^*$ and $\|\cdot\|_{\text{SMNO}}^*$. This will allow us to complete the analysis with $\|\cdot\|_{\text{SMNO}}^*$, which we can compute explicitly (Lemma 6).

Lemma 5. The dual norm of $\|\cdot\|_{\text{SM}}^*$, denoted $\|\cdot\|_{\text{SM}}^*$, satisfies $\frac{1}{2} (\|z\|_{\text{SMNO}}^*)^2 \leq (\|z\|_{\text{SM}}^*)^2 \leq (\|z\|_{\text{SMNO}}^*)^2$.

Proof. By definition

$$\|z\|_{\text{SM}}^* = \sup_x \left\{ z^T x \mid \sqrt{\sum_{i=1}^n \|T_i x\|_{\infty}^2} \leq 1 \right\}. \quad (14)$$

By inspection we can tell that the x that attains the supremum, denoted x_{SM}^* , will satisfy $\sum_{i=1}^n \|T_i x_{\text{SM}}^*\|_{\infty}^2 = 1$. Similarly, x_{SMNO}^* (defined under (13)) must satisfy $\sum_{i=1}^n \|T'_i x_{\text{SMNO}}^*\|_{\infty}^2 = 1$. Note that in these two equalities the $\{T_i\}$ are overlapping sets and the $\{T'_i\}$ are non-overlapping. Therefore, in order to satisfy both equalities it must hold that $\|x_{\text{SM}}^*\|_{\ell} \leq \|x_{\text{SMNO}}^*\|_{\ell}$, $\ell \in [E]$, i.e. the magnitude of the entries of x_{SMNO}^* are equal or larger than the magnitudes of the corresponding entries of x_{SM}^* . Referring to the definitions (12) and (14), this means that $\|z\|_{\text{SM}}^* \leq \|z\|_{\text{SMNO}}^*$.

We now proceed to show the first inequality in the lemma. We note that

$$\begin{aligned} \sum_{i=1}^n \|T_i x\|_{\infty}^2 &= \sum_{i=1}^n \left\| \begin{bmatrix} T'_i \\ \mathbf{0} \end{bmatrix} x + \begin{bmatrix} \mathbf{0} \\ \overline{T'_i} \end{bmatrix} x \right\|_{\infty}^2 \\ &\leq \sum_{i=1}^n \|T'_i x\|_{\infty}^2 + \sum_{i=1}^n \|\overline{T'_i} x\|_{\infty}^2 \leq 2 \sum_{i=1}^n \|\widehat{T'_i} x\|_{\infty}^2, \end{aligned} \quad (15)$$

with

$$\{\widehat{T'_i}\} = \arg \max_{\{T'_i, \overline{T'_i}\}} \left(\sum_{i=1}^n \|T'_i x\|_{\infty}^2 + \sum_{i=1}^n \|\overline{T'_i} x\|_{\infty}^2 \right). \quad (16)$$

We now evaluate (15) and (16) at x_{SMNO}^* . Due to (13) we have $\{\widehat{T'_i}\} = \{T_i^*\}$, and since $\sum_{i=1}^n \|T_i^* x_{\text{SMNO}}^*\|_{\infty}^2 = 1$, the rightmost member of (15) takes value 2. Then, dividing both sides of (15) by 2 we obtain

$$\frac{1}{2} \sum_{i=1}^n \|T_i x_{\text{SMNO}}^*\|_{\infty}^2 = \sum_{i=1}^n \left\| T_i \frac{x_{\text{SMNO}}^*}{\sqrt{2}} \right\|_{\infty}^2 \leq 1,$$

and since $\sum_{i=1}^n \|T_i x_{\text{SM}}^*\|_{\infty}^2 = 1$, we conclude that it must hold that $\frac{1}{\sqrt{2}} \|x_{\text{SMNO}}^*\|_{\ell} \leq \|x_{\text{SM}}^*\|_{\ell}$, $\ell \in [E]$, and thus $\frac{1}{\sqrt{2}} \|z\|_{\text{SMNO}}^* \leq \|z\|_{\text{SM}}^*$. \square

The next lemma gives the value of $\|x\|_{\text{SMNO}}^*$ explicitly, which will be needed to compare the strong convexity constant σ_{SMNO} with σ_A .

Lemma 6. It holds that $\|x\|_{\text{SMNO}}^* = \sqrt{\sum_{i=1}^n \|T_i^* x\|_1^2}$.

Proof. Since the sets $\{\mathcal{S}_i^*\}$ are non-overlapping and in (12) norm $\|\cdot\|_\infty$ is applied per-set, the entries x_ℓ of x_{SMNO}^* will have $|x_\ell| = x^{(i)} \geq 0 \forall \ell \in \mathcal{S}_i^*$ and the sign will match that of the entries of z , i.e. $\text{sign}(x_\ell) = \text{sign}(z_\ell)$. The maximization of (12) then becomes

$$\begin{aligned} & \underset{\{x^{(i)}\}}{\text{maximize}} && \sum_{i=1}^n \sum_{\ell \in \mathcal{S}_i^*} (|z_\ell| \cdot x^{(i)}) \\ & \text{subject to} && \sqrt{\sum_{i=1}^n (x^{(i)})^2} \leq 1. \end{aligned}$$

Factoring out $x^{(i)}$ in the objective and noting that $\sum_{\ell \in \mathcal{S}_i^*} |z_\ell| = \|T_i^* z\|_1$, we can define now both $w = [x^{(1)}, \dots, x^{(n)}]^T$ and $y = [\|T_1^* z\|_1, \dots, \|T_n^* z\|_1]^T$ so that (12) reads

$$\|z\|_{\text{SMNO}}^* = \sup_w \left\{ y^T w \mid \|w\|_2 \leq 1 \right\}.$$

The right-hand side is the definition of $\|\cdot\|_2^*$, the dual of the L2 norm, evaluated at y . Since $\|\cdot\|_2^* = \|\cdot\|_2$, we have that $\|z\|_{\text{SMNO}}^* = \|y\|_2 = \sqrt{\sum_{i=1}^n \|T_i^* z\|_1^2}$. \square

We can now prove the linear convergence rate of SGS-CD.

Theorem 2 (Rate of SGS-CD). SGS-CD converges as

$$\mathbb{E}[F(\lambda^{k+1}) \mid \lambda^k] - F(\lambda^*) \leq \left(1 - \frac{\sigma_{\text{SM}}}{Ln}\right) [F(\lambda^k) - F(\lambda^*)],$$

with

$$\frac{\sigma_A}{N_{\max}} \leq \sigma_{\text{SM}} \leq 2\sigma_A. \quad (17)$$

Proof. Similarly to what we did for SU-CD, we can depart from the strong convexity of F in the $\|\cdot\|_{\text{SM}}$ norm:

$$F(y) \geq F(x) + \langle \nabla F(x), y - x \rangle + \frac{\sigma_{\text{SM}}}{2} (\|y - x\|_{\text{SM}}^*)^2,$$

then minimize both sides with respect to y to obtain

$$F(x^*) \geq F(x) - \frac{1}{2\sigma_{\text{SM}}} (\|\nabla F(x)\|_{\text{SM}})^2, \quad (18)$$

which is analogous to (10), and then rearrange terms to obtain a lower bound on $\|\nabla F(x)\|_{\text{SM}}^2$. Using this lower bound in (11) gives the rate of SGS-CD.

Since this rate is given in terms of σ_{SM} and that of SU-CD in Theorem 1 is given in terms of σ_A , we need (17) to compare both rates. However, we cannot prove these inequalities directly because we cannot compare norms $\|\cdot\|_A$ and $\|\cdot\|_{\text{SM}}^*$ (due to the overlap of the coordinate sets, which prevents us from computing the latter). However, we can compare $\|\cdot\|_A$ with $\|\cdot\|_{\text{SMNO}}^*$ and $\|\cdot\|_{\text{SM}}^*$ with $\|\cdot\|_{\text{SMNO}}^*$ individually, from which we will obtain (17). In particular, we will show the inequalities

$$\frac{\sigma_A}{N_{\max}} \leq \sigma_{\text{SMNO}} \leq \sigma_A \quad (19)$$

and

$$\sigma_{\text{SMNO}} \leq \sigma_{\text{SM}} \leq 2\sigma_{\text{SMNO}}. \quad (20)$$

We start by proving (19). Below we assume $x \in \text{range}(A^T)$; the results here can then be directly applied to the proofs above because $\|\cdot\|_A, \|\cdot\|_{\text{SM}}, \|\cdot\|_{\text{SMNO}}$ and their duals are applied to ∇F , which is always in $\text{range}(A^T)$ (Lemma 3).

For $x \in \text{range}(A^T)$ it holds that (Lemmas 3 and 6):

$$\begin{aligned} \|x\|_A^2 &= \|x\|_2^2 = \sum_{i=1}^E x_i^2 = \sum_{i=1}^n \|T_i^* x\|_2^2 \\ (\|x\|_{\text{SMNO}}^*)^2 &= \sum_{i=1}^n \|T_i^* x\|_1^2. \end{aligned}$$

We also note that, using the Cauchy-Schwarz inequality and denoting $[v]_i$ the i^{th} entry of vector v , it holds both that

$$\begin{aligned} \sum_{i=1}^n \|T_i^* x\|_2^2 &\leq \sum_{i=1}^n \left(\sum_{j \in \mathcal{S}_i^*} |x_j| \right)^2 = \sum_{i=1}^n \|T_i^* x\|_1^2, \text{ and} \\ \sum_{i=1}^n \|T_i^* x\|_1^2 &= \sum_{i=1}^n \left(\mathbf{1}^T \left[|T_i^* x|_1, \dots, |T_i^* x|_{N_i^*} \right]^T \right)^2 \\ &\stackrel{\text{C.S.}}{\leq} \sum_{i=1}^n N_i^* \|T_i^* x\|_2^2 \leq N_{\max} \sum_{i=1}^n \|T_i^* x\|_2^2, \end{aligned}$$

where $N_i^* = |\mathcal{S}_i^*|$. We can summarize these relations as

$$\frac{1}{N_{\max}} (\|x\|_{\text{SMNO}}^*)^2 \leq \|x\|_A^2 \leq (\|x\|_{\text{SMNO}}^*)^2.$$

Using these inequalities in the strong convexity definitions, similarly to [21], we get both

$$\begin{aligned} F(y) &\geq F(x) + \langle \nabla F(x), y - x \rangle + \frac{\sigma_A}{2} (\|y - x\|_A)^2 \\ &\geq F(x) + \langle \nabla F(x), y - x \rangle + \frac{\sigma_A}{2N_{\max}} (\|y - x\|_{\text{SMNO}}^*)^2, \end{aligned} \quad (21)$$

and

$$\begin{aligned} F(y) &\geq F(x) + \langle \nabla F(x), y - x \rangle + \frac{\sigma_{\text{SMNO}}}{2} (\|y - x\|_{\text{SMNO}}^*)^2 \\ &\geq F(x) + \langle \nabla F(x), y - x \rangle + \frac{\sigma_{\text{SMNO}}}{2} (\|y - x\|_A)^2. \end{aligned} \quad (22)$$

Equation (21) says that F is at least $\frac{\sigma_A}{N_{\max}}$ -strongly convex in $\|\cdot\|_{\text{SMNO}}^*$, and eq. (22) says that F is at least σ_{SMNO} -strongly convex in $\|\cdot\|_A$. Together they imply (19).

We can show (20) by the same procedure applied in eqs. (21) and (22), but now using the strong convexity of F in norms $\|\cdot\|_{\text{SM}}^*$ and $\|\cdot\|_{\text{SMNO}}^*$ together with Lemma 5. From $\frac{1}{2}(\|z\|_{\text{SMNO}}^*)^2 \leq (\|z\|_{\text{SM}}^*)^2$ we get $\frac{1}{2}\sigma_{\text{SM}} \leq \sigma_{\text{SMNO}}$, and from $(\|z\|_{\text{SM}}^*)^2 \leq (\|z\|_{\text{SMNO}}^*)^2$ we get $\sigma_{\text{SMNO}} \leq \sigma_{\text{SM}}$.

Finally, putting (19) and (20) together gives (17). \square

Theorems 1 and 2 together allow us to compare the convergence rates of SU-CD and SGS-CD. We note that when σ_{SM} takes the upper value in (17), SGS-CD is (in expectation) N_{\max} times faster than SU-CD. The lower bound in (17), on the other hand, suggests that SGS-CD could be slower than SU-CD. We remark that (in expectation) this is not true and the lower bound is vacuous, since the following holds.

Remark 4. For the same sequence of node activations, the suboptimality reduction of SGS-CD at each iteration is equal to or larger than that of SU-CD.

Taking this fact into account, we have the following result.

Corollary 3. In expectation, SGS-CD converges at least as fast as SU-CD, and can be up to N_{\max} times faster.

Note that this result is analogous to that of [21] for single-machine CD, where they show that the GS rule can be up to d times faster than uniform sampling, d being the dimensionality of the problem.

We remark that achieving the upper bound of N_{\max} speedup may require designing a scenario particularly favorable to SGS-CD with respect to SU-CD. Similarly, having both algorithms converge at the same speed requires a particularly adversarial setting, but this is not hard to design: if all f_i are identical and the graph is complete, no particular benefits can be expected from SGS-CD over SU-CD.

In our simulations of Section VII for the decentralized setting, SGS-CD achieves a speedup approximately in the middle of the range between 1 and N_{\max} . We show that this speedup *increases linearly with N_{\max}* , achieving remarkable gains in terms of suboptimality reduction versus number of iterations (see Figure 2). Furthermore, in the same figure we show that for the parallel distributed setting *the maximum speedup of N_{\max} is attainable*. We explain this further in Section VI-A.

V. SETWISE LIPSCHITZ CD ALGORITHMS

In Section IV we stated that the dual function F is L -smooth and therefore a sufficient condition for the setwise algorithms to converge was using stepsize $\eta = 1/L$. However, the updates of some (and maybe *many*) coordinates could use larger stepsizes by exploiting the fact F has coordinate-wise smoothness $L_\ell \leq L$, i.e. for $\alpha \in \mathbb{R}$:

$$\|\nabla_\ell F(\lambda + \alpha e_\ell) - \nabla_\ell F(\lambda)\| \leq L_\ell \alpha. \quad (23)$$

Therefore, when the coordinate-wise Lipschitz constants L_ℓ are known or can be estimated (see Section V-C) we can apply the update (4) with stepsize $\eta^k = 1/L_\ell$, with ℓ being the coordinate updated at iteration k .

In the sections that follow we show that by using the knowledge (or estimation) of the coordinate-wise Lipschitz constants and per-coordinate stepsizes we can have:

- 1) An algorithm that has randomized but non-uniform neighbor selection that is provably faster than SU-CD. We call this algorithm Setwise Lipschitz CD (SL-CD).
- 2) An algorithm that applies locally the Gauss-Southwell Lipschitz rule [21] and that converges provably faster than both SL-CD and SGS-CD. We call this algorithm Setwise GS Lipschitz CD (SGSL-CD).

Once again, while the seminal work of [21] has analyzed both of these rules in the context of single-machine coordinate descent, their adaptation to setwise CD brings important new challenges. In this section, we prove that SL-CD is at least as fast as SU-CD, and that SGSL-CD is at least as fast as the fastest algorithm between SGS-CD and SL-CD.

In the proofs that follow we will use the following fact.

Fact 1. Denote $a \circ b$ the per-entry product of vectors a and b . Then, for any norm $\|\cdot\|$ and finite $a : a_i > 0 \forall i$, if we define $\|x\|_a := \|a \circ x\|$, then $\|x\|_a^* := \|a_{-1} \circ x\|^*$ with $a_{-1} = [\frac{1}{a_1}, \dots, \frac{1}{a_d}]$.

Proof. By definition

$$\|z\|_a^* = \sup_{\|x\|_a \leq 1} z^T x,$$

and defining $y := a \circ x$ we get

$$\|z\|_a^* = \sup_{\|y\| \leq 1} z^T (a_{-1} \circ y) = \|a_{-1} \circ z\|^*.$$

□

A. Setwise Lipschitz CD (SL-CD)

In SL-CD, an activated node i chooses the edge $\ell \in \mathcal{S}_i$ to update at random with probability

$$p_\ell = \frac{L_\ell}{\sum_{m \in \mathcal{S}_i} L_m} \quad (24)$$

and updates λ_ℓ applying (4) with stepsize $\eta^k = 1/L_\ell$.

For convenience, we define the quantities

$$L^{(i)} := \sum_{m \in \mathcal{S}_i} L_m$$

and

$$\mathcal{L}_\ell := \left(\frac{1}{L^{(i)}} + \frac{1}{L^{(j)}} \right) \quad \text{for } \ell \equiv (i, j).$$

With these definitions, and taking expectation in (5) for the Lipschitz-dependent sampling probabilities (24) gives

$$\begin{aligned} \mathbb{E}[F(\lambda^{k+1}) | \lambda^k] &\leq F(\lambda^k) - \frac{1}{2} \mathbb{E} \left[\frac{1}{L_{\ell_k}} \left[\nabla_{\ell_k} F(\lambda^k) \right]^2 \right] \\ &= F(\lambda^k) - \frac{1}{2n} \sum_{i=1}^n \frac{1}{L^{(i)}} \sum_{\ell \in \mathcal{S}_i} \left[\nabla_\ell F(\lambda^k) \right]^2 \\ &\stackrel{(a)}{=} F(\lambda^k) - \frac{1}{2n} \sum_{\ell=1}^E \mathcal{L}_\ell \left[\nabla_\ell F(\lambda^k) \right]^2 \end{aligned}$$

where in (a) we used that $\ell \equiv (i, j)$ implies $\ell \in \mathcal{S}_i, \mathcal{S}_j$.

In order to prove the convergence rate of SL-CD, provided in Theorem 4, we define the norm

$$\|x\|_{\mathcal{L}} := \sqrt{\sum_{\ell=1}^E \mathcal{L}_\ell x_\ell^2},$$

so that we can write the per-iteration progress of SL-CD as

$$\mathbb{E}[F(\lambda^{k+1})] \leq F(\lambda^k) - \frac{1}{2n} \|\nabla F(\lambda^k)\|_{\mathcal{L}}^2. \quad (25)$$

Noting that $\|x\|_{\mathcal{L}} = \|x \circ [\sqrt{\mathcal{L}_1}, \dots, \sqrt{\mathcal{L}_E}]\|_2$ we can apply Fact 1 to get its dual norm:

$$\|x\|_{\mathcal{L}}^* = \sqrt{\sum_{\ell=1}^E \frac{1}{\mathcal{L}_\ell} x_\ell^2}.$$

We call $\sigma_{\mathcal{L}}$ the strong convexity constant of F in this norm:

$$F(y) \geq F(x) + \langle \nabla F(x), y - x \rangle + \frac{\sigma_{\mathcal{L}}}{2} (\|y - x\|_{\mathcal{L}}^*)^2. \quad (26)$$

We use the definitions of $\|\cdot\|_{\mathcal{L}}^*$ and $\sigma_{\mathcal{L}}$ in the proof of the linear rate of SL-CD, given in the theorem below.

Theorem 4 (Rate of SL-CD). SL-CD converges as

$$\mathbb{E}[F(\lambda^{k+1}) | \lambda^k] - F(\lambda^*) \leq \left(1 - \frac{\sigma_{\mathcal{L}}}{n}\right) [F(\lambda^k) - F(\lambda^*)]$$

and it holds that

$$\sigma_A \mathcal{L}_{\min} \leq \sigma_{\mathcal{L}} \leq \sigma_A \mathcal{L}_{\max} \quad (27)$$

with $\mathcal{L}_{\min} = \min_{\ell} \mathcal{L}_{\ell}$ and $\mathcal{L}_{\max} = \max_{\ell} \mathcal{L}_{\ell}$.

Proof. We start by proving the linear rate. Minimizing both sides of (26) with respect to y as done in (10) and (18) we get

$$F(x^*) \geq F(x) - \frac{1}{2\sigma_{\mathcal{L}}} \|\nabla F(x)\|_{\mathcal{L}}^2.$$

Rearranging terms gives a lower bound on $\|\nabla F(x)\|_{\mathcal{L}}^2$, and replacing in (25) gives the result.

We now move on to show (27). Once again, since the norms are evaluated at $\nabla F(\lambda)$ and (8) holds, to obtain the relation between $\sigma_{\mathcal{L}}$ and σ_A we will compare $\|\cdot\|_{\mathcal{L}}^*$ with $\|\cdot\|_2$ directly. We have that

$$c \|x\|_2^2 - (\|x\|_{\mathcal{L}}^*)^2 = c \sum_{\ell} x_{\ell}^2 - \sum_{\ell} \frac{1}{\mathcal{L}_{\ell}} x_{\ell}^2 = \sum_{\ell} \left(c - \frac{1}{\mathcal{L}_{\ell}} \right) x_{\ell}^2.$$

For $c \geq \max_{\ell} \frac{1}{\mathcal{L}_{\ell}} = \frac{1}{\mathcal{L}_{\min}}$ the expression is larger than zero, and thus

$$\frac{1}{\mathcal{L}_{\min}} \|x\|_2^2 \geq (\|x\|_{\mathcal{L}}^*)^2. \quad (28)$$

Similarly, we have that

$$c \|x\|_{\mathcal{L}}^2 - \|x\|_2^2 = \sum_{\ell} \left(\frac{c}{\mathcal{L}_{\ell}} - 1 \right) x_{\ell}^2$$

is larger than zero for $c \geq \mathcal{L}_{\max}$, and therefore

$$\mathcal{L}_{\max} (\|x\|_{\mathcal{L}}^*)^2 \geq \|x\|_2^2. \quad (29)$$

Using these inequalities (and Lemma 3) in the strong convexity definitions we have on the one hand:

$$\begin{aligned} f(y) &\geq f(x) + \langle \nabla f(x), y - x \rangle + \frac{\sigma_A}{2} \|y - x\|_A^2 \\ &\geq f(x) + \langle \nabla f(x), y - x \rangle + \frac{\sigma_A \mathcal{L}_{\min}}{2} (\|y - x\|_{\mathcal{L}}^*)^2, \end{aligned} \quad (30)$$

and on the other hand:

$$\begin{aligned} f(y) &\geq f(x) + \langle \nabla f(x), y - x \rangle + \frac{\sigma_{\mathcal{L}}}{2} (\|y - x\|_{\mathcal{L}}^*)^2 \\ &\geq f(x) + \langle \nabla f(x), y - x \rangle + \frac{\sigma_{\mathcal{L}}}{2\mathcal{L}_{\max}} \|y - x\|_A^2. \end{aligned} \quad (31)$$

Eqs. (30) and (31) indicate respectively that $\sigma_{\mathcal{L}} \geq \sigma_A \mathcal{L}_{\min}$ and that $\sigma_A \geq \frac{\sigma_{\mathcal{L}}}{\mathcal{L}_{\max}}$. Putting both together gives (27). \square

Having obtained the rate of SL-CD, we can compare it against that of SU-CD. We have the following result.

Corollary 5. In expectation, SL-CD converges as fast or faster than SU-CD.

Proof. The convergence rate of SU-CD is $\frac{2\sigma_A}{LnN_{\max}}$ (Theorem 1) and that of SL-CD is $\frac{\sigma_{\mathcal{L}}}{n}$ (Theorem 4). Since in the slowest case of SL-CD we have $\sigma_{\mathcal{L}} = \sigma_A \mathcal{L}_{\min}$, it suffices to show that $\mathcal{L}_{\min} \geq \frac{2}{LN_{\max}}$. Indeed, we have that

$$L^{(i)} = \sum_{\ell \in \mathcal{S}_i} L_{\ell} \leq L_{\max} |\mathcal{S}_i| \leq L_{\max} N_{\max}$$

and therefore

$$\mathcal{L}_{\min} = \min_{(i,j) \in \mathcal{E}} \left(\frac{1}{L^{(i)}} + \frac{1}{L^{(j)}} \right) \geq \frac{2}{L_{\max} N_{\max}}$$

The proof is complete by noting that it always holds that $L_{\max} \leq L$ [35]. \square

Since both SL-CD and SGS-CD can converge at the same speed as SU-CD in the worst case, we cannot claim that either of them is faster than the other. We can, however, exploit the knowledge of the Lipschitz constants to get an improved version of the GS rule, known as the Gauss-Southwell Lipschitz rule [21], that when combined with per-coordinate stepsizes allows for faster convergence than both SGS-CD and SL-CD. We call this algorithm SGSL-CD, and we analyze it next.

B. Setwise Gauss-Southwell Lipschitz CD (SGSL-CD)

If node i goes active, the Gauss-Southwell Lipschitz (GSL) rule chooses to update $\lambda_{\ell}, \ell \in \mathcal{S}_i$ according to

$$\ell = \operatorname{argmax}_{m \in \mathcal{S}_i} \frac{|\nabla_m f(x^k)|}{\sqrt{L_m}}.$$

If we now use the GSL rule with the per-coordinate stepsizes $\eta^k = 1/L_{\ell}$, the per-iteration progress given by (5) becomes:

$$\mathbb{E}[F(\lambda^{k+1})] \leq F(\lambda^k) - \frac{1}{2n} \sum_{i=1}^n \max_{\ell \in \mathcal{S}_i} \left(\frac{1}{L_{\ell}} [\nabla_{\ell} F(\lambda^k)]^2 \right). \quad (32)$$

Note the resemblance of this expression with the per-iteration progress of SGS-CD in (11). Similarly to the previous procedures, we define the ‘‘Set-Max Lipschitz’’ norm:

$$\|x\|_{\text{SML}} := \sqrt{\sum_{i=1}^n \max_{\ell \in \mathcal{S}_i} \left(\frac{1}{L_{\ell}} x_{\ell}^2 \right)}, \quad (33)$$

and call σ_{SML} the strong convexity constant of F in the dual norm $\|\cdot\|_{\text{SML}}^*$

$$F(y) \geq F(x) + \langle \nabla F(x), y - x \rangle + \frac{\sigma_{\text{SML}}}{2} (\|y - x\|_{\text{SML}}^*)^2. \quad (34)$$

We can now state the convergence rate of SGSL-CD.

Theorem 6 (Rate of SGSL-CD). SGSL-CD converges as

$$\mathbb{E}[F(\lambda^{k+1}) | \lambda^k] - F(\lambda^*) \leq \left(1 - \frac{\sigma_{\text{SML}}}{n} \right) [F(\lambda^k) - F(\lambda^*)], \quad (35)$$

and it holds that

$$\frac{\sigma_{\text{SM}}}{L} \leq \sigma_{\text{SML}}. \quad (36)$$

Proof. The expression of the rate is obtained with the procedure followed for the previous algorithms: minimizing both sides of (34) with respect to y and arranging terms we can find $\|\nabla F(x)\|_{\text{SML}} \geq 2\sigma_{\text{SML}}(F(x) - F(x^*))$, and using this in (32) gives (35).

We now show (36). By definition, the dual norm of $\|\cdot\|_{\text{SML}}$ is

$$\begin{aligned} \|z\|_{\text{SML}}^* &:= \sup_x \left\{ z^T x \mid \|x\|_{\text{SML}} \leq 1 \right\} = \\ &\sup_x \left\{ z^T x \mid \sqrt{\sum_{i=1}^n \max_{\ell \in \mathcal{S}_i} \left(\frac{1}{L_{\ell}} x_{\ell}^2 \right)} \leq 1 \right\}. \end{aligned}$$

Similarly, we can write the dual norm of $\|\cdot\|_{\text{SM}}$ provided in (14) also as

$$\|z\|_{\text{SM}}^* := \sup_x \left\{ z^T x \mid \|x\|_{\text{SM}} \leq 1 \right\} = \sup_x \left\{ z^T x \mid \sqrt{\sum_{\ell \in \mathcal{S}_i} \max_{\ell \in \mathcal{S}_i} x_\ell^2} \leq 1 \right\}.$$

We call the values that achieve the supremum x_{SML}^* and x_{SM}^* , respectively. To maximize $z^T x$, these values will satisfy the constraints of each dual norm with equality, i.e.

$$\|x_{\text{SML}}^*\|_{\text{SML}} = 1 \quad \text{and} \quad \|x_{\text{SM}}^*\|_{\text{SM}} = 1.$$

From these conditions and the definitions of the dual norms above we obtain

$$x_{\text{SML}}^* \circ \left[\frac{1}{\sqrt{L_1}}, \dots, \frac{1}{\sqrt{L_E}} \right] = x_{\text{SM}}^*.$$

Furthermore, using again $L_{\max} \leq L$, we have

$$x_{\text{SM}}^* = x_{\text{SML}}^* \circ \left[\frac{1}{\sqrt{L_1}}, \dots, \frac{1}{\sqrt{L_E}} \right] \geq \frac{1}{\sqrt{L_{\max}}} x_{\text{SML}}^* \geq \frac{1}{\sqrt{L}} x_{\text{SML}}^*,$$

where “ \geq ” indicates coordinate-wise inequality, and therefore

$$\|z\|_{\text{SM}}^* \geq \frac{1}{\sqrt{L}} \|z\|_{\text{SML}}^*.$$

Lastly, using this inequality in the strong convexity equation of F in $\|\cdot\|_{\text{SM}}^*$:

$$\begin{aligned} F(y) &\geq F(x) + \langle \nabla F(x), y - x \rangle + \frac{\sigma_{\text{SM}}}{2} (\|y - x\|_{\text{SM}}^*)^2 \\ &\geq F(x) + \langle \nabla F(x), y - x \rangle + \frac{\sigma_{\text{SM}}}{2L} (\|y - x\|_{\text{SML}}^*)^2, \end{aligned}$$

from where we obtain $\frac{\sigma_{\text{SM}}}{L} \leq \sigma_{\text{SML}}$ \square

Theorem 6 states that SGSL-CD converges (in expectation) at least as fast as SGS-CD. Algorithm SGSL-CD is also at least as fast as SL-CD by an argument analogous to Remark 4: for the same sequence of node activations, the setwise GSL rule achieves an equal or larger suboptimality reduction than the random coordinate sampling with the probabilities in (24). We thus have the following result.

Corollary 7. In expectation, SGSL-CD converges equally fast or faster than both SGS-CD and SL-CD.

We remark that we could have compared the convergence rates of SGSL-CD and SL-CD following a procedure similar to the one used to compare SGS-CD and SU-CD, where we would define a norm using non-overlapping sets (in this case, accounting also for the coordinate-wise Lipschitz constants) as an intermediate step to compare the strong convexity constants $\sigma_{\mathcal{L}}$ and σ_{SML} . We did this derivation and observed that just as it happened with σ_{SM} in eq. (17), the lower bound on σ_{SML} is not tight and suggests that SL-CD could be faster than SGSL-CD, which as we argued above, is not true.

Corollary 5 states that SL-CD is faster than SU-CD, and Corollary 7 states that SGSL-CD is the fastest of all algorithms analyzed here. However, these two methods depend on the knowledge of the coordinate-wise Lipschitz constants L_ℓ (see eq. (23)). These constants are the global upper bounds on the diagonal entries of the Hessian $H = \nabla^2 F$, given by

$$H_{\ell\ell}(\lambda) = \nabla^2 f_i^*(U_i^T \Lambda \lambda) + \nabla^2 f_j^*(U_j^T \Lambda \lambda), \quad \ell \equiv (i, j),$$

i.e. $H_{\ell\ell}(\lambda) \leq L_\ell \forall \lambda$. We next describe a decentralized algorithm to estimate these values when they are not known. In Section VII we show that the versions of SL-CD and SGSL-CD that use estimated constants, which we call SeL-CD and SGSeL-CD respectively, still perform remarkably well.

C. Smoothness constants estimation

In [20] the author proposed a method to estimate the value of the instantaneous Lipschitz constants $L_\ell(\lambda)$ when they are not known. By *instantaneous* we mean the value of the Lipschitz constants at the current point λ^k , and not global values valid for any value of λ .

The procedure consists on finding, every time that variable λ_ℓ is going to be updated at iteration k , the *lowest* value $L_\ell(\lambda^k)$ such that after applying update (4) with stepsize $\eta = 1/L_\ell(\lambda^k)$ it holds that $\nabla_\ell F(\lambda^k) \cdot \nabla_\ell F(\lambda^{k+1}) > 0$. In other words, the procedure searches for a Lipschitz constant (or equivalently, a stepsize) for which the update (4) does not *overshoot*, making the gradient take a completely different direction.

The procedure to estimate $L_\ell(\lambda^k)$ is shown in Algorithm 3. In our numerical simulations, we denote SeL-CD and SGSeL-CD the versions of SL-CD and SGSL-CD that use estimated Lipschitz constants instead of the exact L values. SeL-CD is obtained by replacing line 8 in Algorithm 1 with Algorithm 3 and using the estimated values L_ℓ for the random sampling. SGSeL-CD is obtained by replacing line 9 in Algorithm 2 with Algorithm 3 and using the estimated values L_ℓ for the GSL neighbor choice.

The choice of the initial value \widehat{L}_ℓ^0 before entering the search loop is subject to a trade-off: if \widehat{L}_ℓ^0 is too big, the loop will be exited after only one iteration but we risk being too conservative and making a much smaller step than we could. Conversely, if \widehat{L}_ℓ^0 is too small, by repeated doubling we will eventually find the value \widehat{L}_ℓ that is closest to the true instantaneous smoothness $L_\ell(\lambda^k)$, but this may take many iterations inside the loop, which means many rounds of computation and communication for the nodes involved.

How are these estimated values expected to perform with respect to the analytical ones? This depends heavily on the problem at hand. We can easily construct a case where the exact constants perform better than the estimated: assume that we are in the parallel distributed setting, where we perform primal optimization, and the function to optimize is $F(x) = x^T \text{diag}(L_1, \dots, L_d)x$. If $x^0 \neq \mathbf{0}$, then the algorithm using the analytic constants can converge in d steps (one in each coordinate). This is actually what either SGS-CD or SGSL-CD would achieve. However, using estimated constants will most likely exit the search loop finding values $\widehat{L}_\ell \neq L_\ell$, and thus will need more iterations. Conversely, the analytical constants L_ℓ are *global* quantities, and therefore, although they are valid in the complete optimization space, they might be very different to the real instantaneous Lipschitz constants for many values of λ (or x in the example above). In that case, we may get a much better approximation to the instantaneous value using the estimations, and therefore a faster convergence due to using a larger stepsize. In Section VII-C we provide numerical tests where we observe both behaviors.

Algorithm 3 Online smoothness constant estimation

-
- 1: **Assumption:** Nodes i and j will update λ_ℓ , $\ell \equiv (i, j)$ and they have already exchanged their $\nabla f_x^*(u_x^T A \lambda)$, $x = i, j$.
Inputs: Instantaneous smoothness starting value \widehat{L}_ℓ^0
Each node $x = i, j$ **then runs:**
 - 2: Compute $\nabla_\ell F(\lambda^k)$ with (6) using $\nabla f_x^*(u_x^T A \lambda)$, $x = i, j$
 - 3: Set $\widehat{L}_\ell \leftarrow \widehat{L}_\ell^0$
 - 4: **do ...**
 - 5: Set $\widehat{L}_\ell \leftarrow 2 \cdot \widehat{L}_\ell$
 - 6: Compute $\widehat{\lambda}_\ell = \lambda_\ell^k - (1/\widehat{L}_\ell) \cdot \nabla_\ell F(\lambda^k)$
 - 7: Compute $\nabla f_x^*(u_x^T A \widehat{\lambda}_\ell)$ and send to neighbor
 - 8: Compute $\nabla_\ell F(\widehat{\lambda})$ with (6) using $\nabla f_x^*(u_x^T A \widehat{\lambda}_\ell)$, $x = i, j$
 - 9: **... while** $\nabla_\ell F(\lambda^k) \cdot \nabla_\ell F(\widehat{\lambda}) \leq 0$
 - 10: **end do-while**
 - 11: Set $L_\ell^{k+1} \leftarrow 0.5 \cdot \widehat{L}_\ell$ and $\lambda_\ell^{k+1} \leftarrow \widehat{\lambda}_\ell$
-

VI. ADDITIONAL CONSIDERATIONS

A. Application to parallel distributed optimization

In the parallel distributed setup, the parameter vector is stored in a server accessible by multiple workers, each of which modifies some or all of the coordinates of the parameter. We assume that coordinates are updated by a single worker at each iteration and workers always access the most recent value of the parameter.

In this setting, if there are E coordinates, n workers, and we let each worker i update a different set S_i of coordinates such that (i) the sets overlap, and (ii) each coordinate can be updated by exactly two workers, then all results presented previously (Theorems 1, 2, 4, and 6) hold also for this setting. We remark these two conditions are *not* necessary conditions to apply the setwise CD algorithms to the parallel distributed setting, but only to directly apply the results of the theorems, which were derived for the decentralized setting. In fact, the family of setwise CD algorithms *can always be applied* to the parallel distributed setting *independently* of the degree of overlapping of the sets and the number of the coordinates modified by each worker.

We can then also easily construct a setting where SGS-CD is N_{\max} times faster than SU-CD: let all sets have the same size $|S_i| = N_{\max} \forall i$, exactly $(N_{\max} - 1)$ coordinates in each set have $\nabla_m F(\lambda) = 0$, and only one ℓ have $\nabla_\ell F(\lambda) \neq 0$. In this case, on average *only* $\frac{1}{N_{\max}}$ *times will SU-CD choose the coordinate that gives some improvement*, while SGS-CD will do it at all iterations.

Note that achieving the maximum speedup for this carefully crafted scenario requires that the gradients of all coordinates are independent, which is not verified in the decentralized optimization setting: according to eq. (6), for a $\nabla_m F$ to be zero, it must hold that $\nabla f_i^* = \nabla f_j^*$ for $m \equiv (i, j)$. But unless this equality holds for *all* $(i, j) \in \mathcal{E}$ (i.e., unless the minimum has been attained), λ will continue to change, and the ∇f_i^* will differ. This prevents us from easily designing a scenario in the decentralized setting where SGS-CD attains the speedup upper bound with respect to SU-CD. Nevertheless, in Figure 2 we show examples where (i) *the speedup increases linearly with*

N_{\max} for the decentralized setting, and (ii) *the speedup matches* N_{\max} for the parallel distributed setting.

B. Dual-unfriendly functions and relation to Dual Ascent

The exposition that we have adopted up to this point may suggest that in order to run the setwise CD methods presented here, one should be able to compute the Fenchel conjugates f_i^* for $i \in [n]$. Computing these functions may be tedious, and in some cases simply impossible.

However, we remark that the dual coordinate algorithm presented here is equivalent to the dual decomposition method (Section 2.2 in [36]) and therefore the gradients ∇f_i^* can be directly computed by minimizing the per-node Lagrangian (see also Proposition 11.3 in [37])

$$\nabla f_i^*(u_i^T A \lambda) = \arg \min_{\theta_i} \left[f_i(\theta_i) + \sum_{\ell \in S_i} A_{i\ell} \lambda_\ell \theta_i \right]. \quad (37)$$

Therefore, to apply the algorithms presented here we do *not* need to be able to compute the Fenchel conjugates f_i^* , as long as we can solve (37) analytically or numerically to a high precision. The latter is what we do in Section VII for logistic regression. We remark that a formal treatment of solving (37) inexactly would require an error propagation analysis similar to that done in Section 6 of [13] for the synchronous setup.

C. Case $d > 1$

To extend the proofs above for $d > 1$, the block arrays Λ and U_i should be used instead of A and u_i , and the selector matrices T_i should be redefined in the same way (i.e., by making a Kronecker product with the identity). Then, all the operations that in the proofs above are applied *per entry* (scalar coordinate) of the vector λ , should now be applied to *the magnitude* of each vector coordinate $\lambda_\ell \in \mathbb{R}^d$ of $\lambda \in \mathbb{R}^{Ed}$. Also, since $\nabla_m F \in \mathbb{R}^d$, in this case the GS rule becomes $\operatorname{argmax}_{m \in S_i} \|\nabla_m F(\lambda)\|_2^2$ (and the GSL rule is modified analogously).

VII. NUMERICAL RESULTS

In this section, we test the algorithms proposed in numerical simulations and analyze their performance in a range of different scenarios. In all cases, we used (37) to compute the ∇f_i^* needed in (6). For quadratic and linear least squares problems (37) has a closed-form expression, while for logistic regression we used the SciPy module.

A. SU-CD vs SGS-CD: speedup increase with N_{\max}

Figure 2 shows an example in the decentralized setting where the speedup of SGS-CD compared to SU-CD increases linearly with N_{\max} (left plots), and an example in the parallel distributed setting where SGS-CD achieves the maximum speedup of N_{\max} (right plots).

For the decentralized setting, we created two regular graphs of $n = 32$ nodes and degrees $N_{\max} = 8$ and 12, respectively. The local functions were $f_i(\theta) = \theta^T c I_d \theta$ with $d = 5$, and the constant c being much larger for one node than all others. This choice gave a few edges with smoothness constants much smaller than

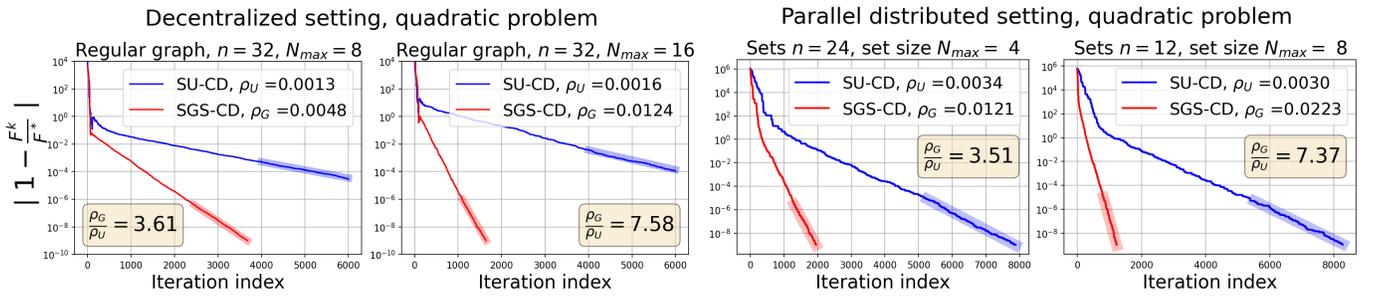


Fig. 2: Comparison of the convergence rates of SU-CD and SGS-CD for quadratic problems in two settings: decentralized optimization over a network (left plots), and parallel distributed computation with parameter server (right plots).

the rest, maximizing the chances to observe the advantages of SGS-CD versus SU-CD (see also the discussion in Section 4.1 of [21]).

For the parallel distributed setting, we created a problem that was separable per-coordinate, and we tried to recreate the conditions described in Section VI-A to approximate the N_{max} gain. We chose $F(x) = x^T \text{diag}(a_1, \dots, a_d)x$ with $d = 48$ and $a_i \sim \mathcal{N}(10, 3) \forall i$. We then created n sets of N_{max} coordinates such that each coordinate belonged to exactly two sets, and simulated two different distributions of the $d = 48$ coordinates: one with $n = 24$ sets of $N_{max} = 4$ coordinates, and another with $n = 12$ sets of $N_{max} = 8$ coordinates. Following the reasoning in Section VI-A, we set the initial value of $(N_{max} - 1)$ coordinates in each set to $x_m^0 = 1$ (close to the optimal value $x_m^* = 0$), and the one remaining to $x_\ell^0 = 100$ (far away from $x_\ell^* = 0$).

In all plots of Figure 2 we used the portion of the curves highlighted with thicker lines to estimate the suboptimality reduction factor $(1 - \rho)$, and called ρ_U and ρ_G the rates of SU-CD and SGS-CD, respectively. In all cases we see that $1 \leq \frac{\rho_G}{\rho_U} \leq N_{max}$, as predicted by Theorem 2. We additionally observe that this ratio increases approximately in the same proportion as N_{max} for the decentralized setting, and is approximately equal to N_{max} in the parallel distributed.

B. Scaling with the number of nodes

In addition to testing the effect of increasing N_{max} , we repeated the experiment of the previous section increasing the number of nodes n in the decentralized setting. Figure 3 shows the results of this experiment for $N_{max} = 8$ (we obtained analogous results for $N_{max} = 16$). The figures confirm the conclusions of the previous section and illustrate how, although both algorithms get slower as the number of nodes increases (as could be expected), the speedup of SGS-CD versus SU-CD remains constant.

C. Estimated vs exact Lipschitz constants

As mentioned in Section V-C, whether the estimated instantaneous Lipschitz constants \widehat{L}_i achieve a faster or slower per-iteration convergence of the exact global L_i than the Lipschitz-informed algorithms depends heavily on the problem at hand. In this section, we compare SL-CD against SeL-CD in two different polynomial functions and show that the fastest algorithm is different in each case (Figure 4).

We consider the parallel distributed setting of the last plot in Figure 2, where the parameter vector x has 48 coordinates and 12 workers update 8 coordinates each, such that each coordinate is updated by exactly two workers. We compare the performance of SL-CD and SeL-CD in two functions: a quadratic $F(x) = a_1 x_1^2 + \dots + a_d x_d^2 + 1$ and an order-four polynomial $F(x) = a_1 x_1^4 + \dots + a_d x_d^4 + 1$. For the former, the global (and also the instantaneous) coordinate-wise Lipschitz constants are $\{L_i\} = \{2a_i\}$, while for the latter they are $\{L_i\} = \{12a_i(\widehat{x}_i)^2\}$, where \widehat{x}_i is the maximum absolute value taken by the entry x_i throughout the optimization. The constants a_i were set to random integers sampled at uniform in the interval $[1, 100]$.

Figure 4 shows the performance of SL-CD and SeL-CD in the quadratic (left plot) and the order-four (right plot) problems. As expected, SL-CD converges faster than SeL-CD in the quadratic problem, where once a coordinate is selected for the first time it is set to its optimal value in that single iteration. Note also that due to this behavior and the fact that the coordinate sampling probabilities of SL-CD are fixed (cf. eq. (24)), sampling a not-yet-optimized coordinate becomes more and more difficult as the iterations progress, which is why the convergence of SL-CD in the left plot of Figure 4 shows a stairwise pattern where the length of the steps becomes larger with the iterations. The estimated constants \widehat{L}_ℓ approximate the optimal values $2a_i$ as well as possible, but cannot match them exactly and thus SeL-CD converges more slowly.

Conversely, in the order-four function SeL-CD converges faster than SL-CD, since in this case the estimated constants will always be closer to the true instantaneous coordinate-wise Lipschitz values than the global smoothness constants L_ℓ . Note, however, that if the exact instantaneous values could be known at each iteration (which is the case of the quadratic function in the left plot of Figure 4), SL-CD using these values would

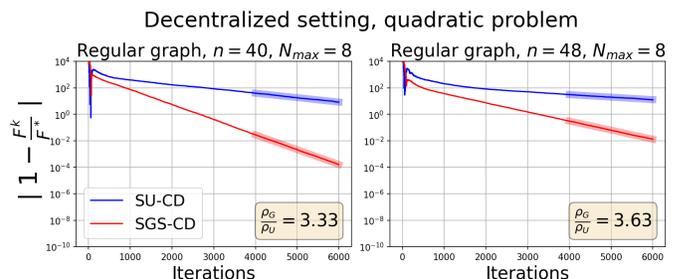


Fig. 3: Comparison of SU-CD and SGS-CD for increasing n .

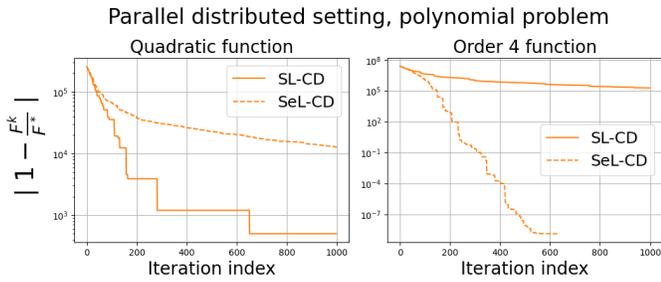


Fig. 4: Comparison of the performance of SL-CD (using global exact L_i values) and SeL-CD in the parallel distributed setting. The experiment was designed to show that which of the two methods is faster depends on the problem considered.

always be faster than SeL-CD.

D. Number of iterations vs communication complexity

Figure 5 shows the performance of all algorithms proposed for the linear least squares (LLS) problem

$$f_i(\theta) = \frac{1}{M} \|X_i \theta - Y_i\|_2^2, \quad X_i \in \mathbb{R}^{M \times d}, \quad Y_i \in \mathbb{R}^M,$$

in two random graphs of $n = 32$ nodes and link probabilities of 0.1 (left plots) and 0.5 (right plots), respectively. The data was generated with the model of [11], $d = 5, M = 30$, and the Y values were additionally multiplied by the index of the corresponding node to have non-iid data between the nodes. Here we show convergence both against number of iterations (top) and against number of vectors in \mathbb{R}^d transmitted through the network (bottom). Table II shows the communication complexity of each algorithm in these terms.

In terms of the number of iterations, we confirm the conclusions of all our corollaries, namely (i) SGS-CD converges faster than SU-CD (Corollary 3), (ii) SL-CD converges faster than SU-CD (Corollary 5), and (iii) SGSL-CD converges faster than both SL-CD and SGS-CD (Corollary 7). Whether the versions with estimated Lipschitz constants SeL-CD and SGSeL-CD are faster than their counterparts with exact Lipschitz knowledge SL-CD and SGSL-CD depends on the problem instance, as discussed in the previous section. Once again, the speedup of applying either the GS or the GSL rule increases radically as the graph becomes more connected.

Although the plots against number of iterations allow us to confirm the predictions of our theory, they mask the important fact that the algorithms with highest rates have also larger communication complexity. The plots in terms of the number of vectors transmitted illustrate better this trade-off. In the bottom plots we see that the suboptimality reduction per number of vectors transmitted can be comparable between the algorithms

TABLE II: Communication complexity of each algorithm: number of vectors in \mathbb{R}^d transmitted in one iteration for an arbitrary activated node i . Variable I indicates the number of iterations inside the do-while loop in Algorithm 3.

SU-CD	2	SGS-CD	$N_i + 1$
SL-CD	2	SGSL-CD	$N_i + 1$
SeL-CD	$2 + 2I$	SGSeL-CD	$N_i + 1 + 2I$

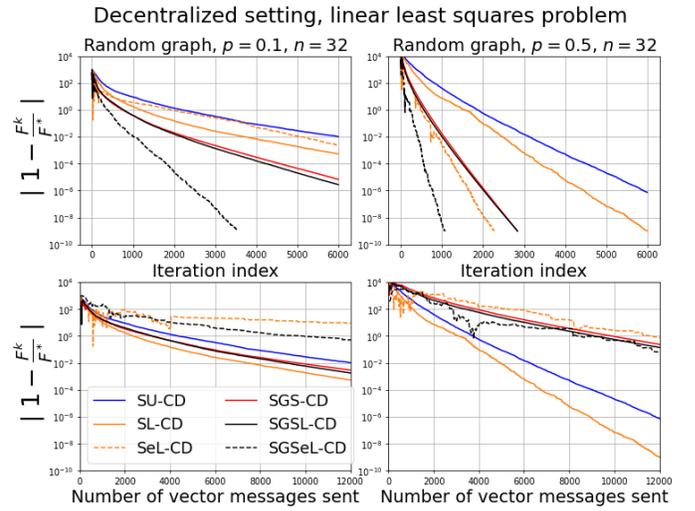


Fig. 5: Performance of the algorithms presented in an LLS problem and two random graphs with different numbers of edges (left and right columns). Top plots: convergence in terms of the number of iterations. Bottom plots: convergence in terms of the number of vectors in \mathbb{R}^d transmitted.

with randomized neighbor choice and those applying the GS(L) rule (left), and the former can even be more efficient as the network becomes denser (right). The greater overhead of the algorithms using estimated constants also becomes apparent in these plots. However, while the plots in Figure 5 support our theory and illustrate the communication overhead of the different algorithms, they do not reflect what we would expect to observe in a real system, where updates can happen in parallel. We analyze this in the next section.

E. Asynchronism and communication overhead

In this section, we analyze how the characteristics of a realistic setting, with asynchronous node activations and communication overhead, can affect the relative performance of the algorithms. For clarity, we limited our analysis to SU-CD and SGS-CD; Table II and Figure 5 allow to extrapolate the conclusions here to the rest of the algorithms presented.

In Figure 6 we repeat the simulations of Figure 2 for the decentralized setting and graph degree $N_{\max} = 16$, now plotting against time (where units are arbitrary). We consider that nodes activate following an exponential distribution with average activation rate κ and that communication takes time τ , during which the transmitting nodes are blocked and cannot either start or participate in other updates. We consider $\kappa \in \{10, 5\}$ and $\tau \in \{0, 1\}$. We also consider that the activation rates can be equal to κ for all nodes (top two plots), or can have a skewed distribution with expected value κ (bottom plot). This skewed distribution (we chose a Zipf with parameter 2) intends to model the presence of slow nodes.

The top plot in Figure 6 shows that if communication time is negligible ($\tau = 0$), we recover the convergence curves of Figure 2, independently of the parameter κ . There is only a shrinkage of the time axis as we go from $\kappa = 10$ to 5 due to updates occurring faster.

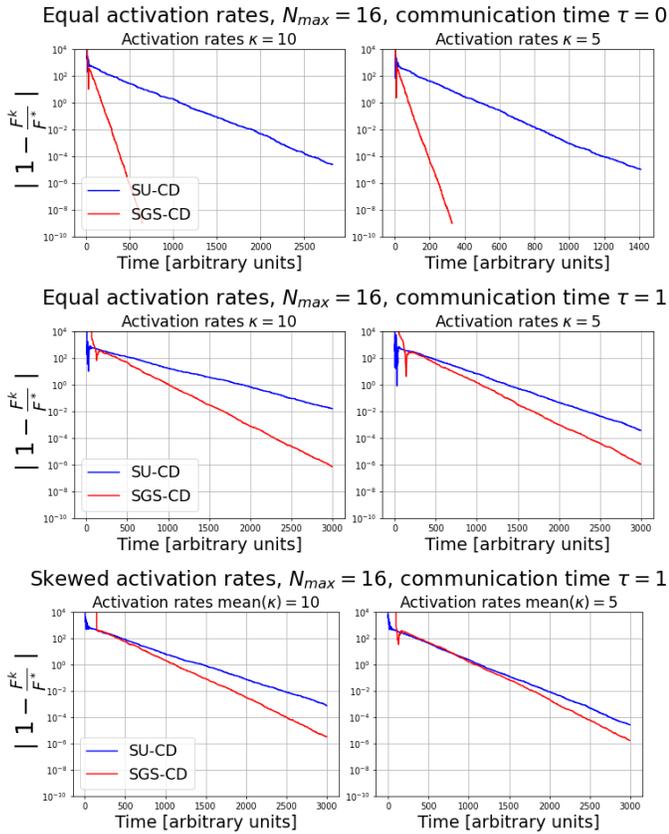


Fig. 6: Comparison of SU-CD vs. SGS-CD in a realistic setting where (i) nodes activation intervals are exponentially distributed with average parameter κ , and (ii) communication takes time τ . We consider equal and skewed activation rates.

The differences show up when communication is not instant anymore ($\tau = 1$, middle plot). Here the gains of SGS-CD versus SU-CD are smaller than in the top plot, since now in SGS-CD nodes have to wait for *all* their neighbors to be available. When we increase the activation rates from $\kappa = 10$ to 5, the gap between the algorithms reduces further, since SU-CD can complete updates more quickly but SGS-CD suffers of greater overhead.

The bottom plot shows that having skewed node activations exacerbates this effect: most nodes have faster rates than κ and this speeds up the convergence of SU-CD, but SGS-CD is constrained by its overhead. We remark, however, that both algorithms converge faster than a synchronous algorithm would, since in the synchronous setting the updates would be completed at the rate of the slowest node.

F. Comparison against token algorithms

In this section we compare the performance of SU-CD and SL-CD against that of token algorithms [29], [30], which have the same edge asynchronicity model of the algorithms considered in this paper and share similar computation and communication complexities.

Figure 7 shows the performance of the algorithms in terms of number of iterations (left) and in the realistic setting described in Section VII-E (right) with parameters $\kappa = 10$ and $\tau = 1$ for the quadratic problem in a regular graph of $n = 32$ and

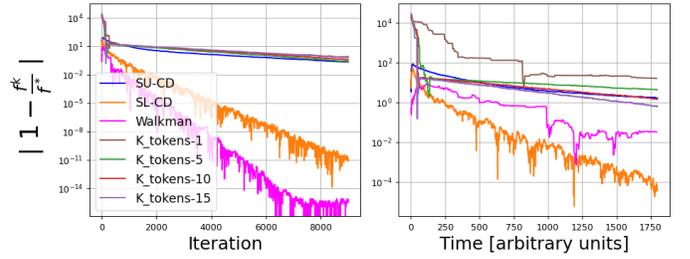


Fig. 7: SU-CD and SL-CD against token algorithms.

degree $N_{\max} = 8$. For Walkman [29], we set their stepsize β to the maximum allowed by their theory. For the multi-token algorithms we simulate $K \in \{1, 5, 10, 15\}$ tokens and perform one computation followed by one communication (among the many variants that the authors propose in their paper). Remark that since Walkman is a primal algorithm, here we plot primal suboptimality.

Walkman is the algorithm that makes the largest progress per iteration (left). Note that in terms of iterations the multi-token algorithms perform equally, since it does not matter what token makes the update. The advantages of the asynchronicity and multiple updates of our algorithms and the multi-token are seen in the right plot, for the realistic setting, where many updates can happen in parallel. Here SL-CD is the fastest of all thanks to allowing many updates in parallel, in contrast to Walkman whose updates are sequential. While SU-CD also allows parallel updates, its convergence is slower due to the use of the same stepsize for all dual coordinates. Still, it performs similarly to the multi-token algorithm with $K = 10$. We remark that the authors of [30] provided variants of their algorithms with local updates and acceleration that would make the convergence of multi-token faster, but would not be a fair comparison in terms of communication and computation.

G. A dual-unfriendly problem with no L knowledge

Figure 8 shows the convergence of SeL-CD and SGSeL-CD in the logistic regression problem

$$f_i(\theta) = \frac{1}{M} \sum_{j=1}^M \log(1 + \exp(-[Y_i]_j \cdot ([X_i]_j)^T \theta_i)) + c \|\theta_i\|_2^2$$

where $[Y_i]_j$ and $([X_i]_j)^T$ are the j -th component and the j -th row of arrays $Y_i \in \mathbb{R}^M$ and $X_i \in \mathbb{R}^{M \times d}$, respectively. We ran the simulation for the same graphs and parameters used for the experiments in Section VII-D. In this case, we cannot compute analytically the optimal value of (37), so we did the optimization in (37) using the SciPy module. For the same reason, we do not know the true coordinate Lipschitz values L_ℓ , so we test only the algorithms using estimated constants.

As in the previous examples, both algorithms converge linearly, and SGSeL-CD is faster than SeL-CD. We may remark, however, that the gap between the two algorithms does not increase with the graph connectivity, as observed between SL-CD and SGSL-CD in the top plots of Figure 5. Indeed, we observed that when the L_ℓ are estimated, the gap between SGSeL-CD and SeL-CD may or may not increase with the connectivity of the graph. We attribute this effect to the fact that the performance of

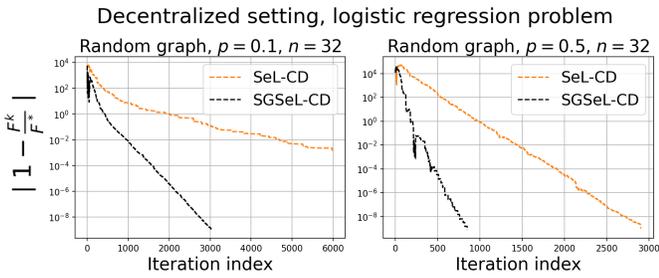


Fig. 8: Convergence of SeL-CD and SGSeL-CD in a logistic regression problem.

the algorithms depends very much on how close to optimal the estimated Lipschitz constants are, and therefore, some instances that allow for a better fit of the true constants using Algorithm 3 have advantage over others whose true constants cannot be well approximated. In the next section we discuss how to improve the estimation of Algorithm 3, and the associated costs of this improvement.

VIII. DISCUSSION AND CONCLUSION

We have presented the class of *setwise CD* optimization algorithms, where in a multi-agent system workers are allowed to modify only a subset of the total number of coordinates at each iteration. These algorithms are suitable for (dual) asynchronous decentralized optimization and (primal) distributed parallel optimization.

We studied the convergence of a number of setwise CD variants: random uniform and Gauss-Southwell setwise coordinate selection (SU-CD and SGS-CD), and their Lipschitz-informed versions (SL-CD and SGSL-CD). We showed linear convergence for all variants for smooth and strongly convex functions f_i , which required developing a new methodology that extends previous results on CD methods.

In particular, we proved that in expectation, when convergence is measured in terms of number of iterations, both SGS-CD and SL-CD are faster than SU-CD, and SGSL-CD is the fastest of them all. However, the algorithms applying the GS(L) rule entail higher communication, which can potentially lower the rate at which iterations are completed and give milder speedups in a real system. The same holds for the algorithms using estimated Lipschitz constants, which require transmitting even more vectors per iteration. Nevertheless, our simulations for a realistic setting with asynchronous activations show that the gains of the GS rule can still be significant with respect to the uniform neighbor choice.

An interesting question is whether the speedup derived for SGS-CD with respect to SU-CD can be obtained with more lenient assumptions than smoothness and strong convexity of the local objectives. Since we solve the problem in the dual, the smoothness assumption may be removed (leading to a just-convex dual) using the methodology of [38]. Removing the strong convexity assumption would lead to a non-smooth dual, and we are not aware of any results that provide the speedup of the GS rule without the smoothness assumption. Filling this gap in the literature would constitute an important contribution to the current knowledge of CD algorithms and may be explored in future work.

ACKNOWLEDGMENTS

M.C. acknowledges funding from the European Union's H2020 MonB5G (grant no. 871780) project. P.M. is also a member of Archimedes/Athena RC—NKUA, and was partially supported by the PEPR IA FOUNDRY project (ANR-23-PEIA-0003) of the French National Research Agency (ANR), and project MIS 5154714 of the National Recovery and Resilience Plan Greece 2.0 funded by the European Union under the NextGenerationEU Program. T.S. acknowledges funding from the Hellenic Foundation for Research & Innovation (HFRI) project 8017: “AI4RecNets: Artificial Intelligence (AI) Driven Co-design of Recommendation and Networking Algorithms”.

REFERENCES

- [1] X. Lian, C. Zhang, H. Zhang, C.-J. Hsieh, W. Zhang, and J. Liu, “Can decentralized algorithms outperform centralized algorithms? a case study for decentralized parallel stochastic gradient descent,” *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [2] P. Wan and M. D. Lemmon, “Event-triggered distributed optimization in sensor networks,” in *2009 International Conference on Information Processing in Sensor Networks*, pp. 49–60, IEEE, 2009.
- [3] M. Alrowaily and Z. Lu, “Secure edge computing in IoT systems: review and case studies,” in *2018 IEEE/ACM Symposium on Edge Computing (SEC)*, pp. 440–444, IEEE, 2018.
- [4] S. Warnat-Herresthal, H. Schultze, K. L. Shastry, S. Manamohan, S. Mukherjee, V. Garg, R. Sarveswara, K. Händler, P. Pickkers, N. A. Aziz, et al., “Swarm learning for decentralized and confidential clinical machine learning,” *Nature*, vol. 594, no. 7862, pp. 265–270, 2021.
- [5] A. Nedic and A. Ozdaglar, “On the rate of convergence of distributed subgradient methods for multi-agent optimization,” in *2007 46th IEEE Conference on Decision and Control*, pp. 4711–4716, IEEE, 2007.
- [6] G. Neglia, C. Xu, D. Towsley, and G. Calbi, “Decentralized gradient methods: does topology matter?,” in *International Conference on Artificial Intelligence and Statistics*, pp. 2348–2358, PMLR, 2020.
- [7] D. Jakovetic, D. Bajovic, A. K. Sahu, and S. Kar, “Convergence rates for distributed stochastic optimization over random networks,” in *2018 IEEE Conference on Decision and Control (CDC)*, pp. 4238–4245, IEEE, 2018.
- [8] X. Lian, W. Zhang, C. Zhang, and J. Liu, “Asynchronous decentralized parallel stochastic gradient descent,” in *International Conference on Machine Learning*, pp. 3043–3052, PMLR, 2018.
- [9] K. Yuan, Q. Ling, and W. Yin, “On the convergence of decentralized gradient descent,” *SIAM Journal on Optimization*, vol. 26, no. 3, pp. 1835–1854, 2016.
- [10] W. Shi, Q. Ling, G. Wu, and W. Yin, “EXTRA: An exact first-order algorithm for decentralized consensus optimization,” *SIAM Journal on Optimization*, vol. 25, no. 2, pp. 944–966, 2015.
- [11] K. Scaman, F. Bach, S. Bubeck, Y. T. Lee, and L. Massoulié, “Optimal algorithms for smooth and strongly convex distributed optimization in networks,” in *international conference on machine learning*, pp. 3027–3036, PMLR, 2017.
- [12] H. Hendrikx, F. Bach, and L. Massoulié, “Accelerated decentralized optimization with local updates for smooth and strongly convex objectives,” in *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 897–906, PMLR, 2019.
- [13] C. A. Uribe, S. Lee, A. Gasnikov, and A. Nedić, “A dual approach for optimal algorithms in distributed optimization over networks,” in *2020 Information Theory and Applications Workshop (ITA)*, pp. 1–37, IEEE, 2020.
- [14] F. Iutzeler, P. Bianchi, P. Ciblat, and W. Hachem, “Asynchronous distributed optimization using a randomized alternating direction method of multipliers,” in *52nd IEEE Conference on Decision and Control (CDC)*, pp. 3671–3676, IEEE, 2013.
- [15] E. Wei and A. Ozdaglar, “On the $O(1/k)$ convergence of asynchronous distributed alternating direction method of multipliers,” in *2013 IEEE Global Conference on Signal and Information Processing*, pp. 551–554, IEEE, 2013.
- [16] J. Xu, S. Zhu, Y. C. Soh, and L. Xie, “Convergence of asynchronous distributed gradient methods over stochastic networks,” *IEEE Transactions on Automatic Control*, vol. 63, no. 2, pp. 434–448, 2017.
- [17] K. Srivastava and A. Nedic, “Distributed asynchronous constrained stochastic optimization,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 5, no. 4, pp. 772–790, 2011.

- [18] S. S. Ram, A. Nedić, and V. V. Veeravalli, "Asynchronous gossip algorithms for stochastic optimization," in *Proceedings of the 48th IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference*, pp. 3581–3586, IEEE, 2009.
- [19] M. Costantini, N. Liakopoulos, P. Mertikopoulos, and T. Spyropoulos, "Pick your neighbor: Local gauss-southwell rule for fast asynchronous decentralized optimization," in *2022 IEEE 61st Conference on Decision and Control (CDC)*, pp. 1602–1609, IEEE, 2022.
- [20] Y. Nesterov, "Efficiency of coordinate descent methods on huge-scale optimization problems," *SIAM Journal on Optimization*, vol. 22, no. 2, pp. 341–362, 2012.
- [21] J. Nutini, M. Schmidt, I. Laradji, M. Friedlander, and H. Koepke, "Coordinate descent converges faster with the Gauss-Southwell rule than random selection," in *International Conference on Machine Learning*, pp. 1632–1641, PMLR, 2015.
- [22] J. Nutini, I. Laradji, and M. Schmidt, "Let's make block coordinate descent go fast: Faster greedy rules, message-passing, active-set complexity, and superlinear convergence," *arXiv preprint arXiv:1712.08859*, 2017.
- [23] J. Tsitsiklis, D. Bertsekas, and M. Athans, "Distributed asynchronous deterministic and stochastic gradient optimization algorithms," *IEEE Transactions on Automatic Control*, vol. 31, no. 9, pp. 803–812, 1986.
- [24] Z. Peng, Y. Xu, M. Yan, and W. Yin, "Arock: An algorithmic framework for asynchronous parallel coordinate updates," *SIAM Journal on Scientific Computing*, vol. 38, no. 5, pp. A2851–A2879, 2016.
- [25] L. Xiao, A. W. Yu, Q. Lin, and W. Chen, "DSCOVER: Randomized primal-dual block coordinate algorithms for asynchronous distributed optimization," *The Journal of Machine Learning Research*, vol. 20, no. 1, pp. 1634–1691, 2019.
- [26] N. Bastianello, R. Carli, L. Schenato, and M. Todescato, "Asynchronous distributed optimization over lossy networks via relaxed ADMM: Stability and linear convergence," *IEEE Transactions on Automatic Control*, vol. 66, no. 6, pp. 2620–2635, 2020.
- [27] Y. Tian, Y. Sun, and G. Scutari, "Achieving linear convergence in distributed asynchronous multiagent optimization," *IEEE Transactions on Automatic Control*, vol. 65, no. 12, pp. 5264–5279, 2020.
- [28] S. Pu, W. Shi, J. Xu, and A. Nedić, "Push-pull gradient methods for distributed optimization in networks," *IEEE Transactions on Automatic Control*, vol. 66, no. 1, pp. 1–16, 2020.
- [29] X. Mao, K. Yuan, Y. Hu, Y. Gu, A. H. Sayed, and W. Yin, "Walkman: A communication-efficient random-walk algorithm for decentralized optimization," *IEEE Transactions on Signal Processing*, vol. 68, pp. 2513–2528, 2020.
- [30] H. Hendrikx, "A principled framework for the design and analysis of token algorithms," in *International Conference on Artificial Intelligence and Statistics*, pp. 470–489, PMLR, 2023.
- [31] A. Verma, M. M. Vasconcelos, U. Mitra, and B. Touri, "Maximal dissent: a state-dependent way to agree in distributed convex optimization," *IEEE Transactions on Control of Network Systems*, 2023.
- [32] I. Notarnicola and G. Notarstefano, "Asynchronous distributed optimization via randomized dual proximal gradient," *IEEE Transactions on Automatic Control*, vol. 62, no. 5, pp. 2095–2106, 2016.
- [33] Y. You, X. Lian, J. Liu, H.-F. Yu, I. S. Dhillon, J. Demmel, and C.-J. Hsieh, "Asynchronous parallel greedy coordinate descent," *Advances in Neural Information Processing Systems*, vol. 29, 2016.
- [34] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [35] S. J. Wright, "Coordinate descent algorithms," *Mathematical Programming*, vol. 151, no. 1, pp. 3–34, 2015.
- [36] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein, *et al.*, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends® in Machine learning*, vol. 3, no. 1, pp. 1–122, 2011.
- [37] R. T. Rockafellar and R. J.-B. Wets, *Variational analysis*, vol. 317. Springer Science & Business Media, 2009.
- [38] S. U. Stich, A. Raj, and M. Jaggi, "Approximate steepest coordinate descent," in *International Conference on Machine Learning*, pp. 3251–3259, PMLR, 2017.



Marina Costantini received her Electronic Engineering degree (with honors) from the National University of La Plata (Argentina) in 2015, her M.Sc. in communication systems from EURECOM (Sophia Antipolis, France) in 2019, and her Ph.D. degree in engineering from EURECOM and Sorbonne University in 2023. Her research interests span different aspects of mathematical optimization, and in her Ph.D. dissertation she proposed new algorithms for multi-agent optimization in distributed computing environments.



Nikolaos Liakopoulos received the Diploma degree in physics and the M.Sc. degree in control and computing from the National and Kapodistrian University of Athens, and the Ph.D. degree in communication systems from Sorbonne University in 2019. He was with the Mathematical and Algorithmic Sciences Lab, Paris Research Center, Huawei Technologies SASU, from 2016 until 2021. He is currently with Amazon Luxembourg, where he is Senior Research Scientist in Amazon Transportation Services.



Panayotis Mertikopoulos graduated valedictorian from the Physics Department of the University of Athens in 2003. He received the M.Sc. and M.Phil. degrees in mathematics from Brown University in 2005 and 2006, and the Ph.D. degree from the University of Athens in 2010. During 2010–2011, he was a post-doctoral researcher at École Polytechnique, Paris, France. Since 2011, he has been affiliated with the French National Center for Scientific Research (CNRS), as a research director since 2023. He currently serves on the editorial board of

several journals, as a Senior Area Chair at NeurIPS, and as managing co-editor of the Open Journal of Mathematical Optimization. He was the recipient of the INFORMS Best Paper Award in network science in 2022 and IEEE NETGCOOP 2012, and was nominated for the *médaille de bronze* of the CNRS in computer science in 2020. His research interests span the interface of game theory, optimization, and learning, with a view toward their applications to machine learning and operations research.



Thrasyvoulos Spyropoulos received the Diploma in Electrical and Computer Engineering from the National Technical University of Athens, Greece, and a Ph.D degree in Electrical Engineering from the University of Southern California. He was a post-doctoral researcher at INRIA and then, a senior researcher with the Swiss Federal Institute of Technology (ETH) Zurich. He has served as an Associate Professor (2010-2020) and Professor (2020-2022) in EURECOM, Sophia-Antipolis, France. He is currently a Full Professor at the

Technical University of Crete, Greece. He is the co-recipient of the best paper award in IEEE SECON 2008, and IEEE WoWMoM 2012, and best paper award runner-up for ACM Mobihoc 2011, IEEE WoWMoM 2015, 2021.