

# AN ACCELERATED COMMUNICATION-EFFICIENT PRIMAL-DUAL OPTIMIZATION FRAMEWORK FOR STRUCTURED MACHINE LEARNING\*

CHENXIN MA<sup>†</sup>, MARTIN JAGGI<sup>‡</sup>, FRANK E. CURTIS<sup>†</sup>, NATHAN SREBRO<sup>§</sup>, AND  
MARTIN TAKÁČ<sup>†</sup>

**Abstract.** Distributed optimization algorithms are essential for training machine learning models on very large-scale datasets. However, they often suffer from communication bottlenecks. Confronting this issue, a communication-efficient primal-dual coordinate ascent framework (CoCoA) and its improved variant CoCoA+ have been proposed, achieving a convergence rate of  $\mathcal{O}(1/t)$  for solving empirical risk minimization problems with Lipschitz continuous losses. In this paper, an accelerated variant of CoCoA+ is proposed and shown to possess a convergence rate of  $\mathcal{O}(1/t^2)$  in terms of reducing suboptimality. The analysis of this rate is also notable in that the convergence rate bounds involve constants that, except in extreme cases, are significantly reduced compared to those previously provided for CoCoA+. The results of numerical experiments are provided to show that acceleration can lead to significant performance gains.

**1. Introduction.** The use of *distributed* optimization has become essential for training machine learning models on very large datasets. For example, distribution is a key ingredient when the dataset does not fit into the memory of a single machine, and rather must be stored in a distributed manner over many machines or computational agents, each having access only to their own (local) segment of the training data.

Efficiently training a machine learning model in such a network setting is challenging due to the cost of communicating information between machines, as compared to the relatively cheap cost of local computation on a single machine. Therefore, designing efficient distributed optimization algorithms, which are able to balance the amount of local computation with the amount of necessary communication between machines, is of crucial importance. This is especially the case since communication bottlenecks and the heterogeneity of large-scale computing systems have been increasing.

In this paper, we improve the existing CoCoA+ framework for communication-efficient distributed optimization [6, 10, 17] by equipping it with *acceleration* [12]. We do not simply apply a generic acceleration scheme, e.g., using the Universal Catalyst [7] (see §1.2 for more discussion), but rather develop a technique specific for our setting. We prove that our accelerated scheme possesses an improved convergence rate involving reduced constants compared to that of CoCoA+. Our theory holds for arbitrary local solvers used on the worker machines. We also provide the results of extensive numerical experiments, where we demonstrate that the new accelerated scheme leads to significantly improved performance in real application settings.

**1.1. Preliminaries.** We present our algorithmic framework in the context of solving the following pair of optimization problems, which are dual to each other:

$$(A) \quad \min_{\alpha \in \mathbb{R}^n} \left[ \mathcal{O}_A(\alpha) := f(A\alpha) + \sum_{i \in \mathcal{I}} g_i(\alpha_i) \right] \quad \text{and}$$

---

\*This material is based upon work supported by the U.S. National Science Foundation, Division of Computing and Communications Foundations, under Award Number CCF-1618717.

<sup>†</sup>Department of Industrial and Systems Engineering, Lehigh University, Bethlehem, PA, USA; [chm514@lehigh.edu](mailto:chm514@lehigh.edu), [frank.e.curtis@gmail.com](mailto:frank.e.curtis@gmail.com), [takac@lehigh.edu](mailto:takac@lehigh.edu)

<sup>‡</sup>School of Computer and Communication Sciences, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland; [martin.jaggi@epfl.ch](mailto:martin.jaggi@epfl.ch)

<sup>§</sup>Toyota Technological Institute at Chicago, Chicago, IL; [nati@ttic.edu](mailto:nati@ttic.edu)

$$(B) \quad \min_{\mathbf{w} \in \mathbb{R}^d} \left[ \mathcal{O}_B(\mathbf{w}) := f^*(\mathbf{w}) + \sum_{i \in \mathcal{I}} g_i^*(-A_i^\top \mathbf{w}) \right].$$

Here,  $h^*$  denotes the *convex conjugate* of a function  $h$ . The two problems are defined by the functions  $f : \mathbb{R}^d \rightarrow \mathbb{R}$  and  $g_i : \mathbb{R} \rightarrow \mathbb{R}$  for all  $i \in \mathcal{I}$  and a data matrix  $A \in \mathbb{R}^{d \times n}$ , about which we make the following assumption.

ASSUMPTION 1. *The functions  $f$  and  $g_i$  for all  $i \in \mathcal{I}$  are closed and convex. In addition, the gradient function  $\nabla f$  is Lipschitz continuous with constant  $\lambda$  and the function  $g_i^*$  for each  $i \in \mathcal{I}$  is Lipschitz continuous with constant  $L$ . Finally, without loss of generality, the columns of the data matrix satisfy  $\|A_i\| \leq 1$ .*

Observe that problem formulation (A) includes many important problems with nonsmooth regularizers, such as Lasso with  $\ell_1$ -norm regularization, i.e.,

$$(1) \quad \min_{\boldsymbol{\alpha} \in \mathbb{R}^n} \frac{1}{2} \|A\boldsymbol{\alpha} - \mathbf{b}\|_2^2 + \lambda_1 \|\boldsymbol{\alpha}\|_1.$$

More generally, many other smooth data-fitting functions can be used in place of  $f$  in our context of problem (A). Problem formulation (B) includes many regularized empirical loss minimization problems with strongly convex regularizers of the form

$$(2) \quad \min_{\mathbf{w} \in \mathbb{R}^d} \frac{1}{n} \sum_{i=1}^n \ell_i(A_i^\top \mathbf{w}) + \frac{\lambda}{2} \|\mathbf{w}\|^2,$$

such as support vector machines (SVMs) or logistic regression. Here, the  $i$ -th column  $A_i \in \mathbb{R}^d$  of  $A$  represents the  $i$ -th data example and  $\lambda > 0$  is a regularization parameter.

**1.2. Motivation and Literature Review.** The increased size of optimization problems of interest in machine learning, as well as the availability of parallel and distributed architectures, has led to various directions of research on the design of parallel algorithms for either solving (1) directly (e.g., see [2, 13, 3]) or indirectly by solving its dual (2) (e.g., see [18, 8, 19]). For example, some previous work has focused on shared memory systems in which the amount of parallelism possible is dependent upon the number of processing units (CPUs) on a single node, which is usually in the range of 16–64 cores per node in a contemporary high performance cluster. To avoid locking and, hence, improve speed, some *asynchronous* algorithms have been proposed and analyzed [8, 13]. However, despite the fact that such methods are efficient with shared memory, naïve extensions of these ideas to distributed environments can be terribly inefficient. This is due to the fact that such naïve extensions require constant communication between nodes, causing them to suffer large overhead.

This type of observation has led many to conclude that, in a distributed computing environment, one must focus not only on the number of data accesses, but also on the number of communication steps, which is usually tied to the number of iterations of the overall optimization algorithm [9]. For this reason, researchers have proposed the use of batch (i.e., full gradient) methods, potentially using (partial or approximate) second-order information as in Newton’s method [16, 21] or a quasi-Newton method such as L-BFGS [22]. The benefit of such methods is that they typically require fewer iterations to achieve high accuracy of the solution. On the other hand, a Newton-type method requires that a linear system be solved (approximately) in every iteration. A popular approach for this procedure is the linear conjugate gradient (CG) algorithm, such as proposed in DISCO [21]. Unfortunately, however, in many applications one might need a substantial number of CG steps in each Newton iteration. Since each CG step requires one pass over the data and one round of communication, this can also lead to substantial communication overhead.

To overcome these issues, the algorithms CoCoA [6], DisDCA [20] and CoCoA+ [10, 17] have been proposed to efficiently balance computation and communication in distributed optimization environments. The main idea of these methods, with the data partitioned across any number of nodes in a cluster, is to define meaningful auxiliary subproblems to be solved in each node using only locally stored data. These subproblems involve some inherited (partial) second-order information, which aids in yielding overall fast convergence for solving the original problems (A) and (B).

As is the case for many other optimization methods for machine learning, the existing CoCoA+ framework only yields a sublinear convergence rate of  $\mathcal{O}(1/t)$  for the general convex case. That said, we are motivated by the fact that several single machine solvers can be improved by incorporating Nesterov acceleration [12, 1], leading to an improved rate of  $\mathcal{O}(1/t^2)$ ; e.g., this was recently successfully done in [14]. In this paper, we are able to provably accelerate the distributed CoCoA+ framework, achieving a  $\mathcal{O}(1/t^2)$  rate for reducing suboptimality.

It is worthwhile to mention that acceleration could be achieved using the Universal Catalyst proposed in [7]. However, for our purposes, this approach is less appealing since it does not allow the local solver to produce randomized solutions which might only have sufficient quality in expectation. Perhaps the accelerated algorithms most relevant to the present work are those in [11, 5]. The subproblems in these papers are assumed to be strongly convex. In our case, this is not general enough; instead, we manage to exploit the general structure of the objective  $\mathcal{O}_A$ .

**1.3. Contributions.** Overall, the contributions represented by the work in this paper can be summarized as follows. Here, for any  $\alpha \in \mathbb{R}^n$  in the context of (A), we define a corresponding dual solution  $\mathbf{w} \in \mathbb{R}^d$  for problem (B) as

$$(3) \quad \mathbf{w} = \mathbf{w}(\alpha) := \nabla f(A\alpha).$$

- We propose and analyze an accelerated communication-efficient block descent algorithm for solving problems (A) and (B). Our analysis shows that our framework possesses a convergence rate of  $\mathcal{O}(1/t^2)$  for suboptimality, i.e., with  $\alpha_t$  representing the  $t$ -th algorithm iterate, we prove the accelerated reduction of  $\mathcal{O}_A(\alpha_t) - \mathcal{O}_A(\alpha_*)$  over time  $t$ , where  $\alpha_*$  represents an optimal solution of (A). Given the recent work in [4], our results for reducing suboptimality can then also be cast in terms of an accelerated rate of the duality gap  $G(\alpha) := \mathcal{O}_A(\alpha) + \mathcal{O}_B(\mathbf{w}(\alpha))$ , as a practically important accuracy certificate.
- The convergence analyses of CoCoA and CoCoA+ yielded bounds involving quantities which depend on the Lipschitz constants for the local subproblems. In this sense, the results might be no better than similar bounds for inexact block proximal gradient descent. By contrast, in this paper, we exploit the structure of the dual objective  $\mathcal{O}_B$  to obtain complexity bounds that do not depend on Lipschitz constants of the auxiliary subproblems. Instead, our bounds merely depend on quantities related to local curvature corresponding to the subproblems.
- We extend our accelerated framework to also cover general non-strongly convex regularizers. This is important, e.g., in the context of  $\ell_1$ -norm regularized Lasso, sparse logistic regression, and elastic net regularized problems.
- We have performed numerous numerical experiments to demonstrate that acceleration can lead to significant performance improvements. In particular, the gains are especially large for small values of the regularization parameter  $\lambda$ , which is important in very-large-scale settings where one often desires the regularization to be

inversely proportional to  $n$ . The C++ code for our implementation of our framework is available on github: <https://github.com/schemmy/CoCoA-Experiments>.

**1.4. Organization.** The remainder of the paper is organized as follows. We start by introducing our new accelerated CoCoA+ (AccCoCoA+) algorithm, including the design of its subproblems and strategies for solving them. We then describe our main complexity result for the algorithm, showing its improvement over that for CoCoA+. Finally, we comment on the results of our numerical experiments.

**2. Accelerated CoCoA+.** In this section, we introduce our proposed accelerated CoCoA+ (AccCoCoA+) algorithm. We begin by defining notation related to the manner in which data is distributed across various machines, then define quantities related to the local subproblems to be solved in each iteration of the algorithm. Of central importance for these subproblems are a carefully defined regularization scheme and a loose assumption on the accuracy to which each subproblem must be solved in each step of AccCoCoA+.

**2.1. Data Partitioning.** Suppose that the  $n$  columns of the dataset are split across  $K$  machines (nodes). Let the index set of columns stored on node  $k \in [K] := \{1, \dots, k\}$  be denoted as  $\mathcal{P}_k$ . We assume that the dataset is partitioned in a disjoint manner such that  $\mathcal{P}_i \cap \mathcal{P}_j = \emptyset$  for any  $i \neq j$  while  $\cup_{k=1}^K \mathcal{P}_k = [n]$ . For notational convenience, we split the variable vector  $\alpha \in \mathbb{R}^n$  into the set of  $K$  vectors  $\{\alpha^{[k]}\}_{k=1}^K$  by employing, for each  $k \in [K]$ , the masking operator

$$(4) \quad (\alpha^{[k]})_i := \begin{cases} \alpha_i & \text{if } i \in \mathcal{P}_k, \\ 0 & \text{otherwise.} \end{cases}$$

As a consequence of separability of  $g$  in (A), we write

$$(5) \quad \psi_k(\alpha^{[k]}) := \sum_{i \in \mathcal{P}_k} g_i(\alpha_i),$$

so that  $g(\alpha) = \sum_{k=1}^K \psi_k(\alpha^{[k]})$ .

**2.2. Subproblem.** Iteration  $t$  of AccCoCoA+ involves the auxiliary vectors  $(\mathbf{y}_t, \mathbf{z}_t) \in \mathbb{R}^n \times \mathbb{R}^n$ , which one may split into  $\{\mathbf{y}_t^{[k]}\}_{k=1}^K$  and  $\{\mathbf{z}_t^{[k]}\}_{k=1}^K$ , respectively, in the same manner as  $\alpha \in \mathbb{R}^n$  in (4). Fundamentally, the goal in iteration  $t$  on each node  $k \in [K]$  is to (approximately) solve

$$(6) \quad \min_{\mathbf{z}_{t+1}^{[k]} \in \mathbb{R}^n} \mathcal{G}_k(\mathbf{z}_{t+1}^{[k]}; \mathbf{y}_t, \mathbf{z}_t),$$

where, for scalars  $\sigma' \geq 0$  and  $\theta_t \geq 0$  (see below), the local objective function is

$$(7) \quad \begin{aligned} \mathcal{G}_k(\mathbf{z}_{t+1}^{[k]}; \mathbf{y}_t, \mathbf{z}_t) &:= \psi_k(\mathbf{z}_{t+1}^{[k]}) + \frac{1}{K} f(A\mathbf{y}_t) \\ &\quad + \nabla f(A\mathbf{y}_t)^\top A(\mathbf{z}_{t+1}^{[k]} - \mathbf{y}_t^{[k]}) + \frac{\lambda\theta_t\sigma'}{2} \left\| A(\mathbf{z}_{t+1}^{[k]} - \mathbf{z}_t^{[k]}) \right\|^2. \end{aligned}$$

At first glance, it is not obvious that this subproblem can be solved only using local data on node  $k$  due to the presence of the term  $\nabla f(A\mathbf{y}_t)^\top A(\mathbf{z}_{t+1}^{[k]} - \mathbf{y}_t^{[k]})$ , which is dependent on the entire dataset. However, by simply making the single shared vector  $\mathbf{w}_t := \nabla f(A\mathbf{y}_t)$  available on each node, the local subproblem (6) only requires

knowledge of the pair  $(\mathbf{y}_t^{[k]}, \mathbf{z}_t^{[k]})$  and the local part of  $A$ , and not the full vectors  $(\mathbf{y}_t, \mathbf{z}_t)$ . Therefore, the storage of the variable vectors  $\mathbf{y}$  and  $\mathbf{z}$  can also be distributed.

The last term in (7) represents a regularization term in which the parameter  $\sigma' \geq 0$  plays a critical role. In short, it can be interpreted as a measure for the cross-dependency of the partitioning of the data. For our analysis, this term must be chosen, for some  $\gamma \in [\frac{1}{K}, 1]$ , to satisfy

$$(8) \quad \sigma' \geq \sigma'_{\min} := \gamma \max_{\boldsymbol{\alpha} \in R^n} \frac{\|A\boldsymbol{\alpha}\|^2}{\sum_{k=1}^K \|A\boldsymbol{\alpha}^{[k]}\|^2}.$$

It is easy to show that  $\sigma'_{\min}/\gamma = \max_{\boldsymbol{\alpha} \in R^n} \frac{\|A\boldsymbol{\alpha}\|^2}{\sum_{k=1}^K \|A\boldsymbol{\alpha}^{[k]}\|^2} \in [1, K]$ . If it is equal to 1, then any pair of samples from different elements of the partition must be orthogonal to each other. In such cases, the function  $\mathcal{O}_A$  is block-separable. On the other hand, if  $\sigma'_{\min}/\gamma$  is close to  $K$ , then the data across the partition are strongly correlated. Note that the choice  $\sigma' = \gamma K$  is “safe” in the sense that (8) holds; see [10, Lemma 4].

**2.3. Approximate Subproblem Solutions.** A strength of our framework is that each subproblem (6) need not be solved exactly. This is critical since, in the extreme, solving the subproblems exactly can be as difficult as solving the original problem. In ACCCoCoA+, we make the assumption that, in iteration  $t$ , the solver employed to solve the subproblem on node  $k$  yields an approximate solution with some additive error  $\epsilon_t \geq 0$ . To be precise, we make the following assumption.

**ASSUMPTION 2** ( $\epsilon_t$ -approximate solutions). *There exists a sequence  $\{\epsilon_t\}_{t=0}^\infty \geq 0$  such that, for each  $t \in \{0, 1, 2, \dots\}$  and  $k \in [K]$ , the local solver employed on node  $k$  in iteration  $t$  produces a (possibly random)  $\mathbf{z}_{t+1}^{[k]}$  satisfying*

$$(9) \quad \mathbf{E}_{t+1}[\mathcal{G}_k(\mathbf{z}_{t+1}^{[k]}; \mathbf{y}_t, \mathbf{z}_t) | t] \leq \mathcal{G}_k(\mathbf{z}_{t+1}^{\star [k]}; \mathbf{y}_t, \mathbf{z}_t) + \epsilon_t,$$

where  $\mathbf{z}_{t+1}^{\star [k]} := \arg \min_{\mathbf{z}_{t+1}^{[k]} \in \mathbb{R}^n} \mathcal{G}_k(\mathbf{z}_{t+1}^{[k]}; \mathbf{y}_t, \mathbf{z}_t)$  and  $\mathbf{E}_{t+1}[\cdot | t]$  indicates conditional expectation given the algorithm history up to time  $t$ .

**2.4. Algorithm.** ACCCoCoA+ is stated as Algorithm 1. Given an initial iterate vector  $\boldsymbol{\alpha}_0$  in the effective domain of  $g$  and the scalar  $\theta_0 = 1$ , each iteration involves a series of steps, only one of which involves communication between nodes. First, the auxiliary vectors  $\{\mathbf{y}_t^{[k]}\}_{k=1}^K$  are set on each node, each representing a convex combination of the variables  $\boldsymbol{\alpha}_t^{[k]}$  and  $\mathbf{z}_t^{[k]}$ . Then, for setting up the local objective (7) for each subproblem (6), the combined vector  $\mathbf{y}_t$  is used to compute  $\mathbf{w}_t$ , which must be communicated to all nodes. After this point in iteration  $t$ , all remaining steps involve local computation on each node: each subproblem is solved approximately to compute  $\{\mathbf{z}_{t+1}^{[k]}\}_{k=1}^K$ , after which the elements of  $\{\boldsymbol{\alpha}_{t+1}^{[k]}\}_{k=1}^K$  are set. Acceleration of the algorithm is due to the careful update for the sequence  $\{\theta_t\}$ , which, since it only involves a prescribed formula for a scalar quantity, can be performed identically on each node. Observe that the update sequence ensures that  $\theta_t \sim \mathcal{O}(1/t)$ .

**3. Convergence Analysis.** In this section, we study the convergence properties of the proposed ACCCoCoA+ algorithm. First, we prove general complexity results, then, respectively in Sections 3.1 and 3.2, we provide interpretations of this main theorem for cases when the subproblems are solved exactly or inexactly.

First, we prove the following lemma related to the sequence  $\{\theta_t\}_{t \geq 0}$ . The result is similar to that given as Lemma 1 in [5].

---

**Algorithm 1** Accelerated CoCoA+ (AccCoCoA+)

---

- 1: choose  $\alpha_0 \in \text{dom}(g) \subseteq \mathbf{R}^n$ ; set  $\mathbf{z}_0 := \alpha_0$ ,  $\gamma \in [\frac{1}{K}, 1]$  and  $\theta_0 := 1$
- 2: **for**  $t \in \{0, 1, 2, \dots\}$  **do**
- 3:   **for**  $k \in [K]$  in parallel, set

$$(10) \quad \mathbf{y}_t^{[k]} := (1 - \gamma\theta_t)\alpha_t^{[k]} + \gamma\theta_t\mathbf{z}_t^{[k]}$$

- 4:   set  $\mathbf{w}_t := \nabla f(A\mathbf{y}_t)$  and communicate to all nodes
- 5:   **for**  $k \in [K]$  in parallel, compute an  $\epsilon_t$ -approximate solution  $\mathbf{z}_{t+1}^{[k]}$  of subproblem (6)
- 6:   **for**  $k \in [K]$  in parallel, set

$$(11) \quad \alpha_{t+1}^{[k]} := \mathbf{y}_t^{[k]} + \gamma\theta_t(\mathbf{z}_{t+1}^{[k]} - \mathbf{z}_t^{[k]})$$

- 7:   set  $\theta_{t+1} := \frac{\sqrt{\gamma^2\theta_t^4 + 4\theta_t^2} - \gamma\theta_t^2}{2}$  (on each node)
  - 8: **end for**
- 

LEMMA 1. *The sequence  $\{\theta_t\}_{t \geq 0}$  is positive, monotonically decreasing, and has*

$$(12) \quad \frac{1 - \gamma\theta_{t+1}}{\theta_{t+1}^2} = \frac{1}{\theta_t^2}$$

and

$$(13) \quad \theta_t \leq \frac{2}{t\gamma + 2} \leq 1$$

for all  $t \geq 0$ .

*Proof.* For each  $t \geq 0$ , the value  $\theta_{t+1}$  can be seen from Step 7 of Algorithm 1 to be the positive root of the quadratic equation

$$(14) \quad \theta^2 + (\gamma\theta_t^2)\theta - \theta_t^2 = 0.$$

Since  $\theta_0 = 1$ , it follows from (14), the fact that a strongly convex quadratic univariate function with a negative vertical intercept has a positive real root, and a simple inductive argument that  $\theta_t > 0$  for all  $t \geq 0$ , as desired. Next, plugging in  $\theta_{t+1}$  for  $\theta$  in (14) and rearranging, we obtain (12). This can again be rearranged to yield

$$(15) \quad \frac{1}{\theta_{t+1}^2} = \frac{1}{\theta_t^2} + \frac{\gamma}{\theta_{t+1}},$$

from which it follows that

$$\frac{1}{\theta_t^2} = 1 + \sum_{i=1}^t \frac{\gamma}{\theta_i} \quad \text{for all } t \geq 0.$$

This shows that  $\{\theta_t\}_{t \geq 0}$  is monotonically decreasing, as desired. We now use mathematical induction to show (13). First, (13) clearly holds for  $t = 0$ , for which we have  $\theta_0 = \frac{2}{0+2} = 1$ . Assuming it holds up to  $t$ , we have from (15) that

$$(16) \quad \frac{1}{\theta_{t+1}^2} - \frac{\gamma}{\theta_{t+1}} = \frac{1}{\theta_t^2} \geq \frac{(\gamma t + 2)^2}{4}.$$

Now observe that the quadratic equation in the variable  $1/\theta$  given by

$$\frac{1}{\theta^2} - \frac{\gamma}{\theta} - \frac{(\gamma t + 2)^2}{4} = 0 \quad \text{has roots} \quad \frac{1}{\theta} = \frac{\gamma \pm \sqrt{\gamma^2 + (\gamma t + 2)^2}}{2}.$$

This shows that, by (16) and since  $\theta_{t+1} > 0$ , we have

$$\frac{1}{\theta_{t+1}} \geq \frac{\gamma + \sqrt{\gamma^2 + (\gamma t + 2)^2}}{2} \geq \frac{\gamma + \sqrt{(\gamma t + 2)^2}}{2} = \frac{\gamma(t+1) + 2}{2}. \quad \square$$

Therefore, we conclude  $\theta_{t+1} \leq \frac{2}{(\gamma t + 1)\gamma + 2} \leq 1$ , which concludes the proof.

Next, we prove the following lemma, which is a modification of Lemma 2 in [5].

LEMMA 2. *Let  $\{\alpha_t, \mathbf{z}_t\}_{t \geq 0}$  be generated by Algorithm 1. Then, for all  $t \geq 0$ ,*

$$(17) \quad \alpha_t = \sum_{l=0}^t \rho_t^l \mathbf{z}_l,$$

where the coefficients  $\{\rho_t^0, \rho_t^1, \dots, \rho_t^t\}$  are nonnegative and sum to 1; i.e.,  $\alpha_t$  is a convex combination of the vectors  $\{\mathbf{z}_0, \mathbf{z}_1, \dots, \mathbf{z}_t\}$ . More precisely, the coefficients are defined recursively in  $t$  as  $\rho_0^0 = 1, \rho_1^0 = 1 - \gamma\theta_0, \rho_1^1 = \gamma\theta_0$  and, for all  $t \geq 1$ ,

$$(18) \quad \rho_{t+1}^l = \begin{cases} (1 - \gamma\theta_t)\rho_t^l & \text{for } l \in \{0, \dots, t\}, \\ \gamma\theta_t & \text{for } l = t + 1. \end{cases}$$

*Proof.* We proceed by induction. First, notice that  $\alpha_0 = \mathbf{z}_0 = \rho_0^0 \mathbf{z}_0$  where  $\rho_0^0 := 1$ . By (10), this implies that  $\mathbf{y}_0 = \mathbf{z}_0$ , which together with  $\theta_0 = 1$  gives (see (11))

$$\alpha_1 = \mathbf{y}_0 + \gamma\theta_0(\mathbf{z}_1 - \mathbf{z}_0) = \mathbf{z}_0 + \gamma\theta_0(\mathbf{z}_1 - \mathbf{z}_0) = (1 - \gamma\theta_0)\mathbf{z}_0 + \gamma\theta_0\mathbf{z}_1,$$

which proves (17) for  $t = 1$ . Assuming now that (17) holds for some  $t \geq 1$ , we obtain

$$(19) \quad \begin{aligned} \alpha_{t+1} &\stackrel{(11)}{=} \mathbf{y}_t + \gamma\theta_t(\mathbf{z}_{t+1} - \mathbf{z}_t) \\ &\stackrel{(10)}{=} (1 - \gamma\theta_t)\alpha_t + \gamma\theta_t\mathbf{z}_t + \gamma\theta_t(\mathbf{z}_{t+1} - \mathbf{z}_t) \\ &\stackrel{(17)}{=} (1 - \gamma\theta_t) \sum_{l=0}^t \rho_t^l \mathbf{z}_l + \gamma\theta_t\mathbf{z}_t + \gamma\theta_t(\mathbf{z}_{t+1} - \mathbf{z}_t) \\ &\stackrel{(17)}{=} \sum_{l=0}^t \underbrace{(1 - \gamma\theta_t)\rho_t^l}_{\rho_{t+1}^l} \mathbf{z}_l + \underbrace{\gamma\theta_t}_{\rho_{t+1}^{t+1}} \mathbf{z}_{t+1}. \end{aligned}$$

From (13) and since  $\gamma \in [\frac{1}{K}, 1]$ , it follows that  $\rho_{t+1}^l \geq 0$  for all  $l \in \{0, \dots, t+1\}$ . It remains to show that the constants sum to 1. This is true since, for all  $t \geq 1$ , the weights in (19) (for  $\alpha_{t+1}$ ) are obtained by taking the corresponding weights for  $\alpha_t$ , multiplying them by  $\gamma\theta_t \in (0, 1]$ , then including the weight for  $\mathbf{z}_{t+1}$  as  $\gamma\theta_t$ .  $\square$

Our next result relates to the optimal solution of an instance of subproblem (6).

LEMMA 3. *Let  $\mathbf{u}^{*[k]} := \arg \min_{\mathbf{u}^{[k]}} \mathcal{G}_k(\mathbf{u}^{[k]}; \mathbf{y}, \mathbf{z})$ . Then, for all  $\mathbf{u}^{[k]}$ ,*

$$(20) \quad \mathcal{G}_k(\mathbf{u}^{[k]}; \mathbf{y}, \mathbf{z}) \geq \mathcal{G}_k(\mathbf{u}^{*[k]}; \mathbf{y}, \mathbf{z}) + \frac{\lambda\theta_t\sigma'}{2} \|A(\mathbf{u}^{[k]} - \mathbf{u}^{*[k]})\|^2.$$

*Proof.* Using convexity of  $f$  and  $g_i$  for all  $i \in \mathcal{I}$  (and, therefore, convexity of  $\psi_k$  defined in (5)), one finds that for  $\mathbf{v}^{[k]} \in \partial\psi_k(\mathbf{u}^{*[k]})$  one has

$$\begin{aligned}
& \mathcal{G}_k(\mathbf{u}^{[k]}; \mathbf{y}, \mathbf{z}) \\
& \stackrel{(7)}{=} \psi_k(\mathbf{u}^{[k]}) + \frac{1}{K}f(A\mathbf{y}) + \nabla f(A\mathbf{y})^\top A(\mathbf{u}^{[k]} - \mathbf{y}^{[k]}) + \frac{\lambda\theta_t\sigma'}{2} \left\| A(\mathbf{u}^{[k]} - \mathbf{z}^{[k]}) \right\|^2 \\
& \geq \psi_k(\mathbf{u}^{*[k]}) + \langle \mathbf{v}^{[k]}, \mathbf{u}^{[k]} - \mathbf{u}^{*[k]} \rangle \\
& \quad + \frac{1}{K}f(A\mathbf{y}) + \nabla f(A\mathbf{y})^\top A(\mathbf{u}^{[k]} - \mathbf{u}^{*[k]} + \mathbf{u}^{*[k]} - \mathbf{y}^{[k]}) \\
& \quad + \frac{\lambda\theta_t\sigma'}{2} \left\| A(\mathbf{u}^{[k]} - \mathbf{u}^{*[k]} + \mathbf{u}^{*[k]} - \mathbf{z}^{[k]}) \right\|^2 \\
& = \psi_k(\mathbf{u}^{*[k]}) + \langle \mathbf{v}^{[k]}, \mathbf{u}^{[k]} - \mathbf{u}^{*[k]} \rangle \\
& \quad + \frac{1}{K}f(A\mathbf{y}) + \nabla f(A\mathbf{y})^\top A(\mathbf{u}^{[k]} - \mathbf{u}^{*[k]}) + \nabla f(A\mathbf{y})^\top A(\mathbf{u}^{*[k]} - \mathbf{y}^{[k]}) \\
& \quad + \frac{\lambda\theta_t\sigma'}{2} \left\| A(\mathbf{u}^{[k]} - \mathbf{u}^{*[k]}) \right\|^2 + \frac{\lambda\theta_t\sigma'}{2} \left\| A(\mathbf{u}^{*[k]} - \mathbf{z}^{[k]}) \right\|^2 \\
& \quad + \lambda\theta_t\sigma'(\mathbf{u}^{[k]} - \mathbf{u}^{*[k]})^\top A^\top A(\mathbf{u}^{*[k]} - \mathbf{z}^{[k]}) \\
& = \underbrace{\psi_k(\mathbf{u}^{*[k]}) + \frac{1}{K}f(A\mathbf{y}) + \nabla f(A\mathbf{y})^\top A(\mathbf{u}^{*[k]} - \mathbf{y}^{[k]}) + \frac{\lambda\theta_t\sigma'}{2} \left\| A(\mathbf{u}^{*[k]} - \mathbf{z}^{[k]}) \right\|^2}_{=\mathcal{G}_k(\mathbf{u}^{*[k]}; \mathbf{y}, \mathbf{z})} \\
& \quad + \langle \mathbf{v}^{[k]} + A^\top \nabla f(A\mathbf{y}) + \lambda\theta_t\sigma' A^\top A(\mathbf{u}^{*[k]} - \mathbf{z}^{[k]}), \mathbf{u}^{[k]} - \mathbf{u}^{*[k]} \rangle \\
& \quad + \frac{\lambda\theta_t\sigma'}{2} \left\| A(\mathbf{u}^{[k]} - \mathbf{u}^{*[k]}) \right\|^2 \\
& \geq \mathcal{G}(\mathbf{u}^{*[k]}; \mathbf{y}, \mathbf{z}) + \frac{\lambda\theta_t\sigma'}{2} \left\| A(\mathbf{u}^{[k]} - \mathbf{u}^{*[k]}) \right\|^2,
\end{aligned}$$

where the fact that

$$\mathbf{v}^{[k]} + A^\top \nabla f(A\mathbf{y}) + \lambda\theta_t\sigma' A^\top A(\mathbf{u}^{*[k]} - \mathbf{z}^{[k]}) \in \partial\mathcal{G}_k(\mathbf{u}^{*[k]}; \mathbf{y}, \mathbf{z})$$

yields the final inequality.  $\square$

We now present our main convergence theorem. Our approach for proving the result is based on the use of randomized estimated sequences, as in [5]. However, we have included an important improvement to this approach that allows us to consider subproblems that are not strongly convex. This is only possible due to the special structure of our objective function  $\mathcal{O}_A$ .

**THEOREM 4.** *For any optimal solution  $\alpha_*$  of problem (A) and all  $t \geq 1$ ,*

$$\begin{aligned}
& \mathbf{E}[\mathcal{O}_A(\alpha_t) - \mathcal{O}_A(\alpha_*)] \\
(21) \quad & \leq \frac{4}{(t\gamma - \gamma + 2)^2} \left( (1 - \gamma)(\mathcal{O}_A(\alpha_0) - \mathcal{O}_A(\alpha_*)) + \frac{\gamma\lambda\sigma'}{2}C + K\epsilon_0\gamma + \sum_{j=1}^{t-1} E_j \right)
\end{aligned}$$

holds where

$$(22a) \quad C := \sum_{k=1}^K \|A(\alpha_*^{[k]} - \alpha_0^{[k]})\|^2,$$



$$(22b) \quad E_j := \frac{K\epsilon_j\gamma}{\theta_j} + \frac{\gamma}{\theta_{j-1}}\epsilon_{j-1} + \gamma\mathcal{R}\sqrt{\frac{2\lambda\sigma'}{\theta_{j-1}}}\epsilon_{j-1} \text{ for all } j \in \{1, \dots, t-1\},$$

$$(22c) \quad \text{and } \mathcal{R} := \max_{k,t} \|A(\boldsymbol{\alpha}_*^{[k]} - \mathbf{z}_t^{*[k]})\|.$$

*Proof.* By convexity of  $g$ , it follows that, for all  $t \geq 0$ , one has

$$(23) \quad g(\boldsymbol{\alpha}_t) \stackrel{(17)}{=} g\left(\sum_{l=0}^t \rho_t^l \mathbf{z}_l\right) \leq \sum_{l=0}^t \rho_t^l g(\mathbf{z}_l) =: \hat{g}^t.$$

Combining this definition and the result from Lemma 2, one gets that

$$(24) \quad \hat{g}^{t+1} \stackrel{(23),(18)}{=} \sum_{l=0}^{t+1} \rho_{t+1}^l g(\mathbf{z}_l) = \gamma\theta_t g(\mathbf{z}_{t+1}) + \sum_{l=0}^t \rho_{t+1}^l g(\mathbf{z}_l).$$

On the other hand, under Assumption 1, one finds that

$$\begin{aligned} & f(A\boldsymbol{\alpha}_{t+1}) \\ \stackrel{(11)}{=} & f(A\mathbf{y}_t + \gamma\theta_t A(\mathbf{z}_{t+1} - \mathbf{z}_t)) \\ \leq & f(A\mathbf{y}_t) + \langle \nabla f(A\mathbf{y}_t), \gamma\theta_t A(\mathbf{z}_{t+1} - \mathbf{z}_t) \rangle + \frac{\lambda\gamma^2\theta_t^2}{2} \|A(\mathbf{z}_{t+1} - \mathbf{z}_t)\|^2 \\ \stackrel{(8)}{\leq} & f(A\mathbf{y}_t) + \gamma\theta_t \left\langle A^\top \nabla f(A\mathbf{y}_t), \sum_{k=1}^K (\mathbf{z}_{t+1}^{[k]} - \mathbf{z}_t^{[k]}) \right\rangle + \frac{\lambda\gamma\sigma'\theta_t^2}{2} \sum_{k=1}^K \|A(\mathbf{z}_{t+1}^{[k]} - \mathbf{z}_t^{[k]})\|^2 \\ = & \sum_{k=1}^K \left\{ \frac{1}{K} (1 - \gamma\theta_t) f(A\mathbf{y}_t) + \gamma\theta_t \left\langle A^\top \nabla f(A\mathbf{y}_t), \mathbf{y}_t^{[k]} - \mathbf{z}_t^{[k]} \right\rangle \right. \\ & \left. + \gamma\theta_t \left( \frac{1}{K} f(A\mathbf{y}_t) + \left\langle A^\top \nabla f(A\mathbf{y}_t), \mathbf{z}_{t+1}^{[k]} - \mathbf{y}_t^{[k]} \right\rangle + \frac{\lambda\sigma'\theta_t}{2} \|A(\mathbf{z}_{t+1}^{[k]} - \mathbf{z}_t^{[k]})\|^2 \right) \right\}. \end{aligned}$$

Next, note that from the definition of  $\mathbf{y}_t$  in the algorithm one finds

$$\begin{aligned} & \mathbf{y}_t^{[k]} \stackrel{(10)}{=} (1 - \gamma\theta_t) \boldsymbol{\alpha}_t^{[k]} + \gamma\theta_t \mathbf{z}_t^{[k]}, \\ & \gamma\theta_t \mathbf{y}_t^{[k]} - \gamma\theta_t \mathbf{z}_t^{[k]} = (1 - \gamma\theta_t) \boldsymbol{\alpha}_t^{[k]} - (1 - \gamma\theta_t) \mathbf{y}_t^{[k]}, \\ (25) \quad & \gamma\theta_t (\mathbf{y}_t^{[k]} - \mathbf{z}_t^{[k]}) = (1 - \gamma\theta_t) (\boldsymbol{\alpha}_t^{[k]} - \mathbf{y}_t^{[k]}). \end{aligned}$$

Defining, for all  $t \geq 0$ , an upper-bound on  $\mathcal{O}_A(\boldsymbol{\alpha}_t)$  as

$$(26) \quad \hat{\mathcal{O}}_A^t := \hat{g}^t + f(A\boldsymbol{\alpha}_t) \stackrel{(23)}{\geq} \mathcal{O}_A(\boldsymbol{\alpha}_t),$$

it follows from above and convexity of  $f$  that one has

$$\begin{aligned} & \hat{\mathcal{O}}_A^{t+1} \\ \stackrel{(26)}{=} & \hat{g}^{t+1} + f(A\boldsymbol{\alpha}_{t+1}) \\ \stackrel{(24)}{=} & \gamma\theta_t g(\mathbf{z}_{t+1}) + \sum_{l=0}^t \rho_{t+1}^l g(\mathbf{z}_l) + f(A\boldsymbol{\alpha}_{t+1}) \end{aligned}$$

$$\begin{aligned}
&\leq \gamma\theta_t g(\mathbf{z}_{t+1}) + \sum_{l=0}^t \rho_{t+1}^l g(\mathbf{z}_l) \\
&\quad + \sum_{k=1}^K \left\{ \frac{1}{K} (1 - \gamma\theta_t) f(A\mathbf{y}_t) \right. \\
&\quad \quad + \gamma\theta_t \left\langle A^\top \nabla f(A\mathbf{y}_t), \mathbf{y}_t^{[k]} - \mathbf{z}_t^{[k]} \right\rangle \\
&\quad \quad + \gamma\theta_t \left( \frac{1}{K} f(A\mathbf{y}_t) + \left\langle A^\top \nabla f(A\mathbf{y}_t), \mathbf{z}_{t+1}^{[k]} - \mathbf{y}_t^{[k]} \right\rangle \right. \\
&\quad \quad \quad \left. \left. + \frac{\lambda\sigma'\theta_t}{2} \left\| A(\mathbf{z}_{t+1}^{[k]} - \mathbf{z}_t^{[k]}) \right\|^2 \right) \right\} \\
&\stackrel{(18),(25)}{=} \gamma\theta_t g(\mathbf{z}_{t+1}) + \underbrace{(1 - \gamma\theta_t) \sum_{l=0}^t \rho_{t+1}^l g(\mathbf{z}_l)}_{\hat{g}^t} \\
&\quad + \sum_{k=1}^K \left\{ \frac{1}{K} (1 - \gamma\theta_t) f(A\mathbf{y}_t) \right. \\
&\quad \quad + (1 - \gamma\theta_t) \left\langle A^\top \nabla f(A\mathbf{y}_t), \boldsymbol{\alpha}_t^{[k]} - \mathbf{y}_t^{[k]} \right\rangle \\
&\quad \quad + \gamma\theta_t \left( \frac{1}{K} f(A\mathbf{y}_t) + \left\langle A^\top \nabla f(A\mathbf{y}_t), \mathbf{z}_{t+1}^{[k]} - \mathbf{y}_t^{[k]} \right\rangle \right. \\
&\quad \quad \quad \left. \left. + \frac{\lambda\sigma'\theta_t}{2} \left\| A(\mathbf{z}_{t+1}^{[k]} - \mathbf{z}_t^{[k]}) \right\|^2 \right) \right\} \\
&\stackrel{(5)}{=} (1 - \gamma\theta_t) \hat{g}^t + (1 - \gamma\theta_t) (f(A\mathbf{y}_t) + \langle A^\top \nabla f(A\mathbf{y}_t), \boldsymbol{\alpha}_t - \mathbf{y}_t \rangle) \\
&\quad + \gamma\theta_t \sum_{k=1}^K \left( \psi_k(\mathbf{z}_{t+1}^{[k]}) + \frac{1}{K} f(A\mathbf{y}_t) + \left\langle A^\top \nabla f(A\mathbf{y}_t), \mathbf{z}_{t+1}^{[k]} - \mathbf{y}_t^{[k]} \right\rangle \right. \\
&\quad \quad \left. + \frac{\lambda\sigma'\theta_t}{2} \left\| A(\mathbf{z}_{t+1}^{[k]} - \mathbf{z}_t^{[k]}) \right\|^2 \right) \\
&\stackrel{(7)}{\leq} (1 - \gamma\theta_t) (\hat{g}^t + f(A\boldsymbol{\alpha}_t)) + \gamma\theta_t \sum_{k=1}^K \mathcal{G}_k(\mathbf{z}_{t+1}^{[k]}; \mathbf{y}_t, \mathbf{z}_t) \\
&\stackrel{(26)}{=} (1 - \gamma\theta_t) \hat{\mathcal{O}}_A^t + \gamma\theta_t \sum_{k=1}^K \mathcal{G}_k(\mathbf{z}_{t+1}^{[k]}; \mathbf{y}_t, \mathbf{z}_t).
\end{aligned}$$

Conditioning on the history up to time  $t$ , it follows from above that

$$\begin{aligned}
\mathbf{E}_{t+1}[\hat{\mathcal{O}}_A^{t+1} | t] &\leq (1 - \gamma\theta_t) \hat{\mathcal{O}}_A^t + \gamma\theta_t \sum_{k=1}^K \mathbf{E}_{t+1}[\mathcal{G}_k(\mathbf{z}_{t+1}^{[k]}; \mathbf{y}_t, \mathbf{z}_t) | t] \\
&\stackrel{(9)}{\leq} (1 - \gamma\theta_t) \hat{\mathcal{O}}_A^t + \gamma\theta_t \sum_{k=1}^K \left( \mathcal{G}_k(\mathbf{z}_{t+1}^{*[k]}; \mathbf{y}_t, \mathbf{z}_t) + \epsilon_t \right),
\end{aligned}$$

meaning that, for any  $\mathbf{u}$ , one has

$$\begin{aligned} & \mathbf{E}_{t+1}[\hat{\mathcal{O}}_A^{t+1}|t] \\ & \stackrel{(20)}{\leq} (1 - \gamma\theta_t)\hat{\mathcal{O}}_A^t + \gamma\theta_t \sum_{k=1}^K \left( \mathcal{G}_k(\mathbf{u}^{[k]}; \mathbf{y}_t, \mathbf{z}_t) - \frac{\lambda\theta_t\sigma'}{2} \left\| A(\mathbf{u}^{[k]} - \mathbf{z}_{t+1}^{\star [k]}) \right\|^2 + \epsilon_t \right). \end{aligned}$$

In particular, choosing  $\mathbf{u} = \boldsymbol{\alpha}_\star$ , where  $\boldsymbol{\alpha}_\star$  is any optimal solution of problem (A), and taking the total expectation, one finds

$$\begin{aligned} & \mathbf{E}[\hat{\mathcal{O}}_A^{t+1}] \\ & \leq (1 - \gamma\theta_t) \mathbf{E}[\hat{\mathcal{O}}_A^t] + \gamma\theta_t \sum_{k=1}^K \mathbf{E} \left[ \mathcal{G}_k(\boldsymbol{\alpha}_\star^{[k]}; \mathbf{y}_t, \mathbf{z}_t) - \frac{\lambda\theta_t\sigma'}{2} \left\| A(\boldsymbol{\alpha}_\star^{[k]} - \mathbf{z}_{t+1}^{\star [k]}) \right\|^2 + \epsilon_t \right] \\ & = (1 - \gamma\theta_t) \mathbf{E}[\hat{\mathcal{O}}_A^t] \\ & \quad + \gamma\theta_t \sum_{k=1}^K \mathbf{E} \left[ \frac{\lambda\theta_t\sigma'}{2} \left\| A(\boldsymbol{\alpha}_\star^{[k]} - \mathbf{z}_t^{[k]}) \right\|^2 - \frac{\lambda\theta_t\sigma'}{2} \left\| A(\boldsymbol{\alpha}_\star^{[k]} - \mathbf{z}_{t+1}^{\star [k]}) \right\|^2 + \epsilon_t \right] \\ & \quad + \gamma\theta_t \sum_{k=1}^K \mathbf{E} \left[ \psi_k(\boldsymbol{\alpha}_\star^{[k]}) + \frac{1}{K} f(A\mathbf{y}_t) + \langle A^\top \nabla f(A\mathbf{y}_t), \boldsymbol{\alpha}_\star^{[k]} - \mathbf{y}_t^{[k]} \rangle \right] \\ & \leq (1 - \gamma\theta_t) \mathbf{E}[\hat{\mathcal{O}}_A^t] + \gamma\theta_t \mathcal{O}_A(\boldsymbol{\alpha}_\star) \\ & \quad + \gamma\theta_t \sum_{k=1}^K \mathbf{E} \left[ \frac{\lambda\theta_t\sigma'}{2} \left\| A(\boldsymbol{\alpha}_\star^{[k]} - \mathbf{z}_t^{[k]}) \right\|^2 - \frac{\lambda\theta_t\sigma'}{2} \left\| A(\boldsymbol{\alpha}_\star^{[k]} - \mathbf{z}_{t+1}^{\star [k]}) \right\|^2 + \epsilon_t \right], \end{aligned}$$

where the last inequality follows from convexity of  $f$ . Defining the scalar quantity  $r_{t+1}^2 = \sum_{k=1}^K \left\| A(\boldsymbol{\alpha}_\star^{[k]} - \mathbf{z}_{t+1}^{\star [k]}) \right\|^2$ , we may conclude from above that

$$\begin{aligned} & \mathbf{E} \left[ \hat{\mathcal{O}}_A^{t+1} - \mathcal{O}_A(\boldsymbol{\alpha}_\star) + \frac{\gamma\lambda\theta_t^2\sigma'}{2} r_{t+1}^2 \right] \\ (27) \quad & \leq (1 - \gamma\theta_t) \mathbf{E}[\hat{\mathcal{O}}_A^t - \mathcal{O}_A(\boldsymbol{\alpha}_\star)] + K\epsilon_t\gamma\theta_t + \frac{\gamma\lambda\theta_t^2\sigma'}{2} \mathbf{E} \left[ \sum_{k=1}^K \left\| A(\boldsymbol{\alpha}_\star^{[k]} - \mathbf{z}_t^{[k]}) \right\|^2 \right]. \end{aligned}$$

To bound the last term on the right-hand side, observe that, for all  $k \in \{1, \dots, K\}$ ,

$$\begin{aligned} \mathbf{E} \left[ \left\| A(\boldsymbol{\alpha}_\star^{[k]} - \mathbf{z}_t^{[k]}) \right\|^2 \right] &= \mathbf{E} \left[ \left\| A(\boldsymbol{\alpha}_\star^{[k]} - \mathbf{z}_t^{\star [k]} + \mathbf{z}_t^{\star [k]} - \mathbf{z}_t^{[k]}) \right\|^2 \right] \\ &\leq \mathbf{E} \left[ \left\| A(\boldsymbol{\alpha}_\star^{[k]} - \mathbf{z}_t^{\star [k]}) \right\|^2 \right] + \underbrace{\mathbf{E} \left[ \left\| A(\mathbf{z}_t^{\star [k]} - \mathbf{z}_t^{[k]}) \right\|^2 \right]}_{=: C_1} \\ &\quad + 2 \underbrace{\mathbf{E} \left[ \left\| A(\boldsymbol{\alpha}_\star^{[k]} - \mathbf{z}_t^{\star [k]}) \right\| \cdot \left\| A(\mathbf{z}_t^{\star [k]} - \mathbf{z}_t^{[k]}) \right\| \right]}_{=: C_2} \\ &\leq \mathbf{E}[r_t^2] + C_1 + C_2. \end{aligned}$$

It remains to bound  $C_1$  and  $C_2$ . From (20) and (9), one finds that

$$\begin{aligned}
(28) \quad & \mathcal{G}_k(\mathbf{z}_{t+1}^{*[k]}; \mathbf{y}_t, \mathbf{z}_t) + \frac{\lambda\theta_t\sigma'}{2} \mathbf{E}_{t+1}[\|A(\mathbf{z}_{t+1}^{*[k]} - \mathbf{z}_{t+1}^{[k]})\|^2 | t] \\
& \stackrel{(20)}{\leq} \mathbf{E}_{t+1}[\mathcal{G}_k(\mathbf{z}_{t+1}^{[k]}; \mathbf{y}_t, \mathbf{z}_t) | t] \\
& \stackrel{(9)}{\leq} \mathcal{G}_k(\mathbf{z}_{t+1}^{*[k]}; \mathbf{y}_t, \mathbf{z}_t) + \epsilon_t,
\end{aligned}$$

and hence we can conclude that

$$(29) \quad \frac{\lambda\theta_t\sigma'}{2} \mathbf{E}_{t+1}[\|A(\mathbf{z}_{t+1}^{*[k]} - \mathbf{z}_{t+1}^{[k]})\|^2 | t] \leq \epsilon_t.$$

Therefore,

$$\begin{aligned}
(30) \quad & \frac{\gamma\lambda\theta_t^2\sigma'}{2} C_1 = \frac{\gamma\lambda\theta_t^2\sigma'}{2} \mathbf{E} \left[ \|A(\mathbf{z}_t^{*[k]} - \mathbf{z}_t^{[k]})\|^2 \right] \\
& = \frac{\gamma\theta_t^2}{\theta_{t-1}} \frac{\lambda\theta_{t-1}\sigma'}{2} \mathbf{E}_{t-1} \left[ \mathbf{E}_t \left[ \|A(\mathbf{z}_t^{*[k]} - \mathbf{z}_t^{[k]})\|^2 | t-1 \right] \right] \stackrel{(29)}{\leq} \frac{\gamma\theta_t^2}{\theta_{t-1}} \epsilon_{t-1}.
\end{aligned}$$

Now, let us bound  $C_2$ . By defining  $\mathcal{R}$  as in (22c), Jensen's inequality gives

$$\begin{aligned}
\frac{\gamma\lambda\theta_t^2\sigma'}{2} C_2 &= \frac{\gamma\lambda\theta_t^2\sigma'}{2} \left( 2 \mathbf{E} \left[ \|A(\boldsymbol{\alpha}_*^{[k]} - \mathbf{z}_t^{*[k]})\| \cdot \|A(\mathbf{z}_t^{*[k]} - \mathbf{z}_t^{[k]})\| \right] \right) \\
& \stackrel{(22c)}{\leq} \mathcal{R} \gamma \lambda \theta_t^2 \sigma' \mathbf{E} \left[ \|A(\mathbf{z}_t^{*[k]} - \mathbf{z}_t^{[k]})\| \right] \\
& \leq \mathcal{R} \gamma \lambda \theta_t^2 \sigma' \sqrt{\frac{2}{\gamma \lambda \theta_t^2 \sigma'}} \sqrt{\frac{\gamma \lambda \theta_t^2 \sigma'}{2} \mathbf{E} \left[ \|A(\mathbf{z}_t^{*[k]} - \mathbf{z}_t^{[k]})\|^2 \right]} \\
& \stackrel{(30)}{\leq} \mathcal{R} \gamma \lambda \theta_t^2 \sigma' \sqrt{\frac{2}{\gamma \lambda \theta_t^2 \sigma'}} \sqrt{\frac{\gamma \theta_t^2}{\theta_{t-1}} \epsilon_{t-1}} = \gamma \theta_t^2 \mathcal{R} \sqrt{\frac{2\lambda\sigma'}{\theta_{t-1}}} \epsilon_{t-1}.
\end{aligned}$$

Putting everything together leads to

$$\begin{aligned}
\mathbf{E} \left[ \hat{\mathcal{O}}_A^{t+1} - \mathcal{O}_A(\boldsymbol{\alpha}_*) + \frac{\gamma\lambda\theta_t^2\sigma'}{2} r_{t+1}^2 \right] & \stackrel{(27)}{\leq} \mathbf{E} \left[ (1 - \gamma\theta_t)(\hat{\mathcal{O}}_A^t - \mathcal{O}_A(\boldsymbol{\alpha}_*)) + \frac{\gamma\lambda\theta_t^2\sigma'}{2} r_t^2 \right] \\
& \quad + K\epsilon_t\gamma\theta_t + \frac{\gamma\theta_t^2}{\theta_{t-1}} \epsilon_{t-1} + \gamma\theta_t^2 \mathcal{R} \sqrt{\frac{2\lambda\sigma'}{\theta_{t-1}}} \epsilon_{t-1}.
\end{aligned}$$

Dividing both sides by  $\theta_t^2$  and denoting  $\phi_t := \mathbf{E}[\hat{\mathcal{O}}_A^t - \mathcal{O}_A(\boldsymbol{\alpha}_*)]$  and  $\tilde{r}_t^2 := \mathbf{E}[r_t^2]$  gives

$$\begin{aligned}
(31) \quad & \frac{1}{\theta_t^2} \phi_{t+1} + \frac{\gamma\lambda\sigma'}{2} \tilde{r}_{t+1}^2 \stackrel{(27)}{\leq} \frac{1 - \gamma\theta_t}{\theta_t^2} \phi_t + \frac{\gamma\lambda\sigma'}{2} \tilde{r}_t^2 \\
& \quad + \underbrace{\frac{K\epsilon_t\gamma}{\theta_t} + \frac{\gamma}{\theta_{t-1}} \epsilon_{t-1} + \gamma \mathcal{R} \sqrt{\frac{2\lambda\sigma'}{\theta_{t-1}}} \epsilon_{t-1}}_{=: E_t}.
\end{aligned}$$

Now, by the property of  $\theta_t$  in (12), one finds

$$(32) \quad \frac{1 - \gamma\theta_{t+1}}{\theta_{t+1}^2} \phi_{t+1} + \frac{\gamma\lambda\sigma'}{2} \tilde{r}_{t+1}^2 \leq \frac{1 - \gamma\theta_t}{\theta_t^2} \phi_t + \frac{\gamma\lambda\sigma'}{2} \tilde{r}_t^2 + E_t.$$

Unrolling the recurrence, one obtains for  $t \geq 1$  that

$$(33) \quad \frac{1 - \gamma\theta_t}{\theta_t^2} \phi_t + \frac{\gamma\lambda\sigma'}{2} \tilde{r}_t^2 \stackrel{(32)}{\leq} \frac{1 - \gamma\theta_0}{\theta_0^2} \phi_0 + \frac{\gamma\lambda\sigma'}{2} \tilde{r}_0^2 + \sum_{i=0}^{t-1} E_i.$$

Hence, along with (12), one has for  $t \geq 1$  that

$$\begin{aligned} \phi_t &\leq \theta_{t-1}^2 \left( \frac{1 - \gamma\theta_0}{\theta_0^2} \phi_0 + \frac{\gamma\lambda\sigma'}{2} \tilde{r}_0^2 + \sum_{i=0}^{t-1} E_i \right) \\ &\stackrel{(13)}{\leq} \left( \frac{2}{t\gamma - \gamma + 2} \right)^2 \left( \frac{1 - \gamma\theta_0}{\theta_0^2} \phi_0 + \frac{\gamma\lambda\sigma'}{2} \sum_{k=1}^K \|A(\boldsymbol{\alpha}_\star^{[k]} - \mathbf{z}_0^{[k]})\|^2 + K\epsilon_0\gamma + \sum_{i=1}^{t-1} E_i \right) \end{aligned}$$

and (21) follows.  $\square$

Theorem 4 describes the behavior of suboptimality only. We can, however, use the following theorem from [4], which relates suboptimality with the duality gap.

THEOREM 5 ([4, Theorem 4]). *Suppose problem (A) is solved by a (possibly randomized) algorithm producing a sequence of iterates  $\{\boldsymbol{\alpha}_t\}_{t=0}^\infty$  such that, for all  $t \geq 1$ ,*

$$\mathbf{E}[\mathcal{O}_A(\boldsymbol{\alpha}_t) - \mathcal{O}_A(\boldsymbol{\alpha}_\star)] \leq \frac{F}{d(t)}$$

for some scalar  $F \geq 0$  and function  $d$ . If, for  $t \geq 1$ , it holds that

$$(34) \quad d(t) \geq \max \left\{ \frac{2F\lambda n}{\tilde{\sigma}^2 L^2}, \frac{2F\tilde{\sigma}^2 L^2}{\lambda n \epsilon^2} \right\},$$

where  $\tilde{\sigma}^2$  is the maximum eigenvalue of  $A^\top A$ , then the expected duality gap satisfies

$$\mathbf{E}[\mathcal{O}_A(\boldsymbol{\alpha}_t) + \mathcal{O}_B(\mathbf{w}(\boldsymbol{\alpha}_t))] \leq \epsilon.$$

Related to (34), henceforth, we assume that  $\epsilon$  is such that  $\frac{2F\lambda n}{\tilde{\sigma}^2 L^2} < \frac{2F\tilde{\sigma}^2 L^2}{\lambda n \epsilon^2}$ .

Before stating our key corollaries of Theorem 4 and comparisons with results for other methods in the literature, let us define a few important quantities on which these results depend (see Table 1). The first quantity is  $\tilde{\sigma}^2$ , already defined in Theorem 5. Due to the fact that for each data column  $A_i$  we have  $\|A_i\| \leq 1$ , it follows that this quantity is bounded by  $n$ . The second quantity is  $\sigma_k^2$ , the maximum eigenvalue of a Gram matrix for the local data on node  $k$ . This value will be large if the samples stored on node  $k$  are correlated. The next important quantity is  $\sigma^2$ , which depends on each  $\sigma_k^2$  and the size of each partition  $n_k$ . Let us now, for simplicity, assume that the data are split equally, i.e.,  $n_k = \frac{n}{K}$  for all  $k \in [K]$ . The last quantity is  $\mathcal{R}$ , which measures the maximum possible distance between the optimal solution and each iterate, i.e.,

$$\begin{aligned} \mathcal{R} &\stackrel{(22c)}{:=} \max_{k,t} \|A(\boldsymbol{\alpha}_\star^{[k]} - \mathbf{z}_t^{\star[k]})\| \leq \max_{k,t} \sigma_k \|(\boldsymbol{\alpha}_\star^{[k]} - \mathbf{z}_t^{\star[k]})\| \\ &\leq \max_k (\sigma_k 2L\sqrt{n_k}) \leq \max_k (2Ln_k). \end{aligned}$$

Here, we have used the fact that since  $g_i^*$  is  $L$ -Lipschitz, it follows that  $g_i$  has an effective domain bounded within a ball of radius  $L$  about the origin, which in turn means that for all  $(a, b) \in \text{dom}(g_i) \times \text{dom}(g_i)$  one has  $\|a - b\| \leq 2L$ .

TABLE 1  
Important quantities for our comparisons with related convergence results.

symbol	expression	worst-case upper bound	worst-case $\lim_{K \rightarrow n}$
$\tilde{\sigma}^2$	$\max_{\alpha} \frac{\ A\alpha\ ^2}{\ \alpha\ ^2}$	$n$	$n$
$\sigma_k^2$	$\max_{\alpha^{[k]}} \frac{\ A\alpha^{[k]}\ ^2}{\ \alpha^{[k]}\ ^2}$	$n_k = \frac{n}{K}$	1
$\sigma^2$	$\sum_{k=1}^K  \mathcal{P}_k  \sigma_k^2$	$\sum_k n_k^2 = K \frac{n^2}{K^2} = \frac{n^2}{K}$	$n$
$\mathcal{R}$	$\min_{\alpha_*} \max_{k,t} \ A(\alpha_*^{[k]} - \mathbf{z}_t^{*[k]})\ $	$\max_k 2Ln_k = \frac{2Ln}{K}$	$2L$

**3.1. Exact Subproblem Solves.** If the subproblems are solved exactly, i.e., if  $\epsilon_t = 0$  for all  $t \geq 0$ , then (21) has the form

$$(35) \quad \mathbf{E}[\mathcal{O}_A(\alpha_t) - \mathcal{O}_A(\alpha_*)] \leq \frac{4}{(t\gamma - \gamma + 2)^2} \left( (1 - \gamma)(\mathcal{O}_A(\alpha_0) - \mathcal{O}_A(\alpha_*)) + \frac{\gamma\lambda\sigma'}{2} \sum_{k=1}^K \|A(\alpha_*^{[k]} - \alpha_0^{[k]})\|^2 \right).$$

A nice property of this result is that the second term in the parentheses might be equal to zero even if  $\mathbf{z}_0 = \alpha_0 \neq \alpha_*$ . This is not the case for other results for accelerated algorithms as their subproblems are strongly convex [12, 1, 5, 7] and hence have the term  $\|\alpha_* - \alpha_0\|^2$  present in their complexity guarantees. Another nice property is that  $\sum_{k=1}^K \|A(\alpha_*^{[k]} - \alpha_0^{[k]})\|^2$  can be bounded above by  $\sum_{k=1}^K \sigma_k^2 \|\alpha_*^{[k]} - \alpha_0^{[k]}\|^2$ , though the former can be much smaller.

COROLLARY 6. Consider the extreme cases  $\gamma = \frac{1}{K}$  and  $\gamma = 1$ . As discussed in Section 2, for the first case, one can choose  $\sigma' = 1$  while for the second case one can choose  $\sigma' = K$ . To obtain  $\mathbf{E}[\mathcal{O}_A(\alpha_t) - \mathcal{O}_A(\alpha_*)] \leq \epsilon$ , one has to run Algorithm 1 for at least  $t > T$  iterations, where  $T$  is defined as follows for each case.

- Case  $\gamma = 1$ :

$$T = \sqrt{\frac{2\lambda\sigma'}{\epsilon} \sum_{k=1}^K \|A(\alpha_*^{[k]} - \alpha_0^{[k]})\|^2} \leq \sqrt{\frac{8L^2}{\lambda\epsilon}}$$

- Case  $\gamma = 1/K$ :

$$T \leq \sqrt{\frac{4K(K-1)}{\epsilon} (\mathcal{O}_A(\alpha_0) - \mathcal{O}_A(\alpha_*)) + \frac{8L^2}{\lambda\epsilon}}.$$

As found in [10], the complexity for  $\gamma = 1$  is better as one requires fewer iterations in order to have an  $\epsilon$ -approximate solution in expectation. Hence, let us focus only on the case of  $\gamma = 1$  and compare our rate with the results derived in [10]. From the proof of [10, Theorem 8], one obtains that for CoCoA+, if  $\alpha_0 = \mathbf{0}$ , it holds that  $\mathbf{E}[\mathcal{O}_A(\alpha_t) - \mathcal{O}_A(\alpha_*)] \leq \frac{4KL^2\lambda\sigma^2}{(1+\frac{1}{2}(t-t_0))}$ , where

$$t_0 \geq \max\{0, \log(2(\mathcal{O}_A(\alpha_0) - \mathcal{O}_A(\alpha_*)) / (4KL^2\lambda\sigma))\}.$$

To obtain suboptimality below  $\epsilon$ , CoCoA+ needs to be run for  $t \geq T = t_0 + \frac{8KL^2\lambda\sigma^2}{\epsilon} \leq t_0 + \frac{8L^2}{\lambda\epsilon}$  iterations. Neglecting the  $t_0$  term, CoCoA+ needs  $\frac{8L^2}{\lambda\epsilon}$  iterations, whereas AccCoCoA+ needs only  $\sqrt{\frac{8L^2}{\lambda\epsilon}}$ . This improvement is consistent with proximal

gradient and accelerated gradient descent [11], which is as expected since, in the worst-case, they will produce the same iterates as CoCoA+ and AccCoCoA+.

Let us also derive a complexity bound for the duality gap. Corollary 6 with Theorem 5 implies that whenever  $t \geq T$ , where  $T = \frac{2L}{\epsilon} \sqrt{\frac{\sigma' \tilde{\sigma}^2}{n} \sum_{k=1}^K \|A(\alpha_\star^{[k]} - \alpha_0^{[k]})\|^2} \leq \frac{4L^2}{\lambda \epsilon}$ , the expected duality gap satisfies  $\mathbf{E}[\mathcal{O}_A(\alpha_t) + \mathcal{O}_B(\mathbf{w}(\alpha_t))] \leq \epsilon$ . This is valid for the number of iterations of AccCoCoA+. Note also that for the worst-case complexity for non-accelerated batch-SDCA (or full gradient method) for hinge loss ( $L = 1$ ), the known results hold for the *average solution* and the number of iterations are  $\sim \mathcal{O}(\frac{1}{\lambda \epsilon})$  (see, e.g., [18, 19] with the batch size chosen as large as possible).

**3.2. Inexact Subproblem Solvers.** In order to get a better understanding of the case when the subproblems are solved approximately, let us define an auxiliary nonnegative sequence  $\{a_t\}_{t=0}^\infty$  such that  $\frac{\sum_{i=0}^\infty \sqrt{a_i}}{t^2} \rightarrow 0$ . In this section, we assume that the errors for the local solvers are set as  $\epsilon_t = a_t \theta_t$ .

Before analyzing this case in more detail, let us bound the total accumulated error up to iteration  $t$ , i.e.,  $\sum_{j=1}^{t-1} E_j$ . One finds that

$$\begin{aligned} \sum_{j=1}^{t-1} E_j &= K\gamma \sum_{j=1}^{t-1} a_j + \sum_{j=0}^{t-2} a_j + \mathcal{R} \sqrt{2\lambda\sigma'} \sum_{j=1}^{t-1} \sqrt{a_{j-1}} \\ &\leq (K+1) \sum_{j=0}^{t-1} a_j + \mathcal{R} \sqrt{2\lambda\sigma'} \sum_{j=0}^{t-1} \sqrt{a_j} \\ &\leq (K+1) \sum_{j=0}^{t-1} a_j + \sqrt{\frac{8L^2}{\lambda K}} \sum_{j=0}^{t-1} \sqrt{a_j} =: S_t. \end{aligned}$$

Let us now consider two cases:

- Suppose that  $a_t = r \in \mathbb{R}^+$ . In this case,  $\epsilon_t = r\theta_t = \mathcal{O}(1/t)$ . Moreover,  $S_t = \left((K+1)r + \sqrt{r \frac{8L^2}{\lambda K}}\right) t$ . This implies that

$$\begin{aligned} \mathbf{E}[\mathcal{O}_A(\alpha_t) - \mathcal{O}_A(\alpha_\star)] &\leq 4 \frac{\frac{\lambda\sigma'}{2} \sum_{k=1}^K \|A(\alpha_\star^{[k]} - \alpha_0^{[k]})\|^2 + K\epsilon_0 + S_t}{(t+1)^2} \\ &\sim \mathcal{O}\left(\frac{(K+1)r + \sqrt{r \frac{8L^2}{\lambda K}}}{t}\right). \end{aligned}$$

Note that, in this case,  $\epsilon \sim \mathcal{O}(1/t)$ . This might create the impression that the local solver has to do more work as  $t$  increases; however, note that the Lipschitz constant of the gradient of the smooth part of the subproblem solved by the local solver also scales as  $\theta_t \sim \mathcal{O}(1/t)$ .

- A second interesting case is when  $\lim_{t \rightarrow \infty} S_t =: S_\infty < \infty$ . For example, suppose  $a_t = \frac{r}{t^p}$  with  $p > 2$ . Then, indeed,  $\lim_{t \rightarrow \infty} S_t$  is finite. In this case, one obtains

$$\begin{aligned} \mathbf{E}[\mathcal{O}_A(\alpha_t) - \mathcal{O}_A(\alpha_\star)] &\leq 4 \frac{\frac{\lambda\sigma'}{2} \sum_{k=1}^K \|A(\alpha_\star^{[k]} - \alpha_0^{[k]})\|^2 + K\epsilon_0 + S_t}{(t+1)^2} \\ &\leq 4 \frac{\frac{\lambda\sigma'}{2} \sum_{k=1}^K \|A(\alpha_\star^{[k]} - \alpha_0^{[k]})\|^2 + K\epsilon_0 + S_\infty}{(t+1)^2} \\ &\leq \frac{\frac{8L^2}{\lambda} + 4K\epsilon_0 + 4(K+1) \sum_{i=0}^\infty a_i + 4\sqrt{\frac{8L^2}{\lambda K}} \sum_{i=0}^\infty \sqrt{a_i}}{(t+1)^2} \sim \mathcal{O}(1/t^2). \end{aligned}$$

**4. Numerical Experiments.** In this section, we report the results of numerical experiments. The purpose of providing the results of these experiments is twofold. For one thing, we use them to illustrate the benefits of acceleration by providing results that compare the performance of AccCoCoA+ versus CoCoA+ [10, 17]. In addition, we explore the communication/computation tradeoff and the scalability of AccCoCoA+. For a recent comparison of CoCoA+ to other distributed solvers, including Quasi-Newton methods and ADMM, we refer the reader to [10, 17].

Our implementations of CoCoA+ and AccCoCoA+ are written in C++ using MPI for communication, run on `m3.xlarge` Amazon EC2 instances. We run all the experiments across  $K = 4$  nodes using SDCA [15] as the local solver. The datasets we used are summarized in Table 2.<sup>1</sup>

TABLE 2  
*Datasets used in the numerical experiments.*

Dataset	$n$	$d$	size(GB)
url	2,396,130	3,231,961	2.21
rcv1.test	677,399	47,236	1.21
covtype	581,012	54	0.07
epsilon	400,000	2,000	3.6

We first compare CoCoA+ versus AccCoCoA+ for solving hinge-loss SVM problems of the form (2). For both algorithms, the local solver, SDCA, is run for  $H = 5 \times 10^5$  iterations (closed-form single coordinate solutions). In Figure 1, we compare the evolution of the duality gap with respect to the number of iterations and elapsed time. The results suggest a benefit of acceleration in terms of decreasing the duality gap, both when  $\gamma = 1/K$  and when  $\gamma = 1$ . In Figure 2, we show how the regularization parameter  $\lambda$  can affect the performance of both algorithms when solving the problem with the `url` dataset. In particular, the experiments suggest that as the value of  $\lambda$  becomes smaller, there will be a more significant benefit from employing the accelerated algorithm.

Figure 3 shows analogous results when the algorithms are employed to solve the Lasso problem in (1), for which increasing the regularization parameter  $\lambda_1$  typically leads to more sparsity of the solution vector. The results indicate that the accelerated algorithm offers faster convergence of the sub-optimality gap to zero, especially for small values of  $\lambda_1$ .

We also ran experiments to demonstrate how the performance of AccCoCoA+ depends on the number of iterations ( $H$ ) that SDCA runs for solving each subproblem. Figure 4 shows that for larger  $H$  the subproblems will be solved more accurately, and thus fewer outer iterations can be expected to reach a desired tolerance on the duality gap. However, in terms of running time, it is not always better to choose larger  $H$ . For example, for the `covtype` dataset, choosing  $H = 10^4$  results in less time to reach a tolerance of  $10^{-2}$  than is needed when  $H = 10^3$  or  $H = 10^5$ . However, for the `news20` dataset, the time required decreases with  $H$  for all values considered in our experiments. The reason that this occurs is that  $d$  is quite large in this dataset, which makes each round of communication quite time consuming. Therefore, by solving the subproblems more accurately (by running more iterations), one achieves a better balance between communication and computation.

<sup>1</sup>All data are available at <https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/>.



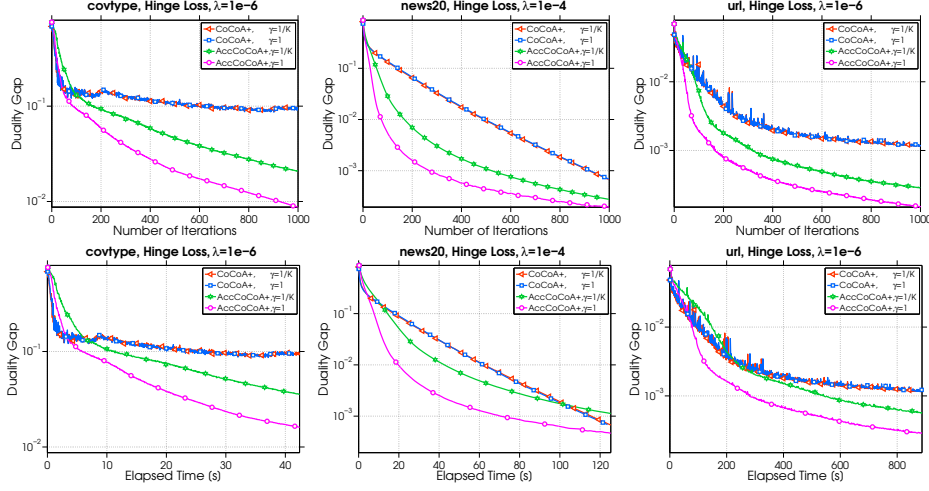


FIG. 1. Duality gap as a function of iterations (top row) and elapsed time (bottom row) when solving hinge-loss SVM problems.

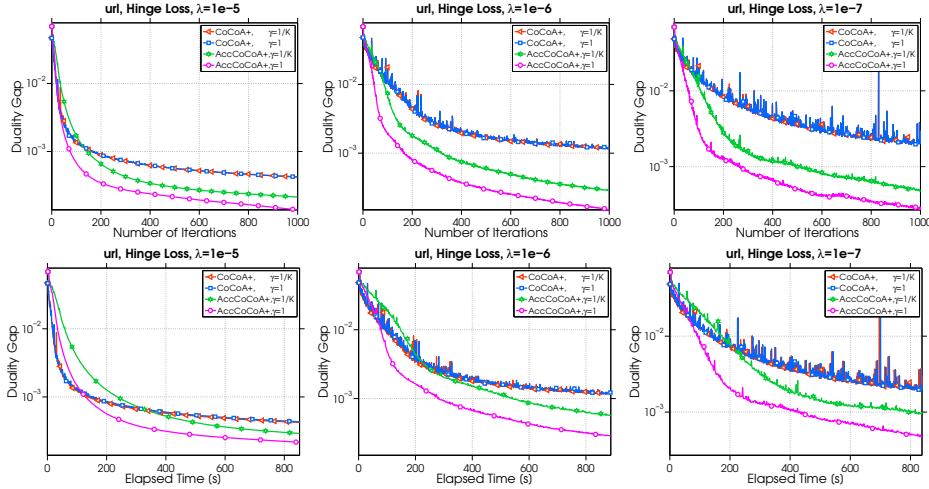


FIG. 2. Duality gap as a function of iterations (top row) and elapsed time (bottom row) when solving hinge-loss SVM problems with different regularization values ( $\lambda$ ) on the *url* dataset.

Such a tradeoff between communication and computation can also be observed in Figure 5, where we compare how the values of  $H$  can affect the convergence of the duality gap for the *covtype* and *rcv\_test* datasets. The left two plots illustrate that to reach the same tolerance on the duality gap, the number of iterations can always be reduced by reaching a more accurate solution for each subproblem. However, the curves in the other two plots indicate that there is always a best value of  $H$  that leads to fastest convergence with respect to running time.

Our last experiment shows how AccCoCoA+ scales with the number of machines ( $K$ ). The results are shown in Figure 6. We set  $H = \frac{n}{K}$  for every  $K$  to make sure that the same amount of data is utilized during each iteration of AccCoCoA+ across all  $K$  machines. By doing so, it is also expected that each subproblem has been




FIG. 3. Sub-optimality gap as a function of iterations (top row) and elapsed time (bottom row) when solving Lasso problems. The choice of the regularizer value ( $\lambda_1$ ) are such that for the url dataset the density of the optimal solution is 5.3%, while for the epsilon dataset the density is 13.56%.

solved to similar accuracy, according to the complexity result of SDCA in [15]. The results show that when  $\gamma = 1$ , ACCCoCoA+ takes almost the same amount of time regardless of how many machines are used, which demonstrates the better scaling properties it has than the  $\gamma = \frac{1}{K}$  case.

**5. Conclusion.** We proposed and analyzed ACCCoCoA+, an accelerated variant of CoCoA+ achieving the optimal  $\mathcal{O}(1/t^2)$  convergence rate. The method is robust to inaccurate subproblems both in theory and practice, and performs well in large-scale experiments. Our analysis provides constants in the convergence rate which are significantly tighter compared to those previously obtained for CoCoA+.

## REFERENCES

- [1] Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1):183–202, 2009.
- [2] Joseph K Bradley, Aapo Kyrola, Danny Bickson, and Carlos Guestrin. Parallel coordinate descent for l1-regularized loss minimization. *arXiv:1105.5379*, 2011.
- [3] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *The Journal of Machine Learning Research*, 12:2121–2159, 2011.
- [4] Celestine Dünner, Simone Forte, Martin Takáč, and Martin Jaggi. Primal-dual rates and certificates. *arXiv:1602.05205*, 2016.
- [5] Olivier Fercoq and Peter Richtárik. Accelerated, parallel and proximal coordinate descent. *SIAM Journal on Optimization*, 25(4):1997–2023, 2015.
- [6] Martin Jaggi, Virginia Smith, Martin Takáč, Jonathan Terhorst, Sanjay Krishnan, Thomas Hofmann, and Michael I Jordan. Communication-efficient distributed dual coordinate ascent. In *Advances in Neural Information Processing Systems*, pages 3068–3076, 2014.
- [7] Hongzhou Lin, Julien Mairal, and Zaid Harchaoui. A universal catalyst for first-order optimization. In *Advances in Neural Information Processing Systems*, pages 3366–3374, 2015.
- [8] Ji Liu and Stephen J Wright. Asynchronous stochastic coordinate descent: Parallelism and




FIG. 4. Comparison of performance for different inner iteration limits when  $\sigma' = \frac{1}{K}$  and  $\sigma' = 1$ .

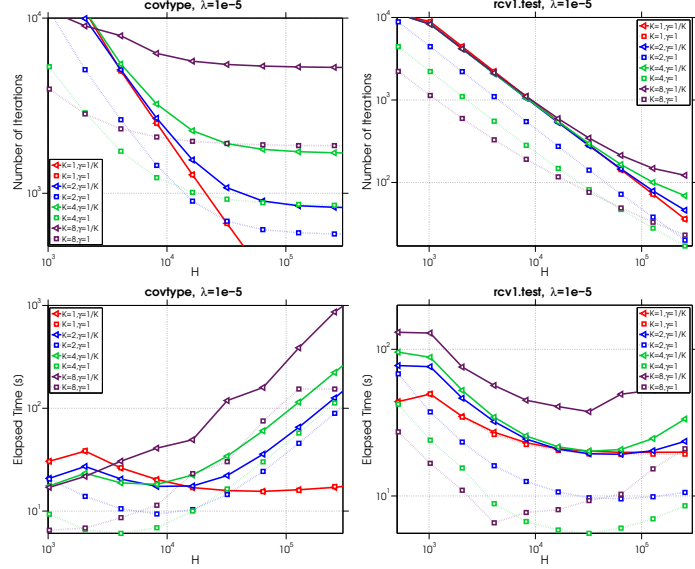


FIG. 5. Number of iterations and running time required to reach a tolerance of  $10^{-3}$  on the duality gap as the inner iteration limit ( $H$ ) is varied.

convergence properties. *SIAM Journal on Optimization*, 25(1):351–376, 2015.

- [9] Chenxin Ma, Jakub Konečný, Martin Jaggi, Virginia Smith, Michael I Jordan, Peter Richtárik, and Martin Takáč. Distributed optimization with arbitrary local solvers. *arXiv:1512.04039*, 2015.
- [10] Chenxin Ma, Virginia Smith, Martin Jaggi, Michael I Jordan, Peter Richtárik, and Martin Takáč. Adding vs. averaging in distributed primal-dual optimization. *In 32th International Conference on Machine Learning, ICML 2015*, 2015.
- [11] Yu Nesterov. Gradient methods for minimizing composite functions. *Mathematical Programming*,




FIG. 6. Running times on different numbers of machines  $K$ .

- 140(1):125–161, 2013.
- [12] Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2013.
  - [13] Benjamin Recht, Christopher Re, Stephen Wright, and Feng Niu. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *Advances in Neural Information Processing Systems*, pages 693–701, 2011.
  - [14] Shai Shalev-Shwartz and Tong Zhang. Accelerated mini-batch stochastic dual coordinate ascent. In *Advances in Neural Information Processing Systems*, pages 378–385, 2013.
  - [15] Shai Shalev-Shwartz and Tong Zhang. Stochastic dual coordinate ascent methods for regularized loss minimization. *The Journal of Machine Learning Research*, 14(1):567–599, 2013.
  - [16] Ohad Shamir, Nathan Srebro, and Tong Zhang. Communication efficient distributed optimization using an approximate newton-type method. *arXiv:1312.7853*, 2013.
  - [17] Virginia Smith, Simone Forte, Chenxin Ma, Martin Takáč, Michael I Jordan, and Martin Jaggi. CoCoA: A General Framework for Communication-Efficient Distributed Optimization. *arXiv*, 2016.
  - [18] Martin Takáč, Avleen Bijral, Peter Richtárik, and Nathan Srebro. Mini-batch primal and dual methods for svms. In *In 30th International Conference on Machine Learning, ICML 2013*, 2013.
  - [19] Martin Takáč, Peter Richtárik, and Nathan Srebro. Distributed mini-batch SDCA. *arXiv:1507.08322*, 2015.
  - [20] Tianbao Yang, Shenghuo Zhu, Rong Jin, and Yuanqing Lin. Analysis of distributed stochastic dual coordinate ascent. *arXiv:1312.1031*, 2013.
  - [21] Yuchen Zhang and Lin Xiao. Communication-efficient distributed optimization of self-concordant empirical loss. *arXiv:1501.00263*, 2015.
  - [22] Ciyong Zhu, Richard H Byrd, Peihuang Lu, and Jorge Nocedal. Algorithm 778: L-bfgs-b: Fortran subroutines for large-scale bound-constrained optimization. *ACM Transactions on Mathematical Software (TOMS)*, 23(4):550–560, 1997.