
Communication-Efficient Distributed Dual Coordinate Ascent

Martin Jaggi *
ETH Zurich

Virginia Smith *
UC Berkeley

Martin Takáč
Lehigh University

Jonathan Terhorst
UC Berkeley

Sanjay Krishnan
UC Berkeley

Thomas Hofmann
ETH Zurich

Michael I. Jordan
UC Berkeley

Abstract

Communication remains the most significant bottleneck in the performance of distributed optimization algorithms for large-scale machine learning. In this paper, we propose a communication-efficient framework, CoCoA, that uses local computation in a primal-dual setting to dramatically reduce the amount of necessary communication. We provide a strong convergence rate analysis for this class of algorithms, as well as experiments on real-world distributed datasets with implementations in Spark. In our experiments, we find that as compared to state-of-the-art mini-batch versions of SGD and SDCA algorithms, CoCoA converges to the same .001-accurate solution quality on average $25\times$ as quickly.

1 Introduction

With the immense growth of available data, developing distributed algorithms for machine learning is increasingly important, and yet remains a challenging topic both theoretically and in practice. On typical real-world systems, communicating data between machines is vastly more expensive than reading data from main memory, e.g. by a factor of several orders of magnitude when leveraging commodity hardware.¹ Yet, despite this reality, most existing distributed optimization methods for machine learning require significant communication between workers, often equalling the amount of local computation (or reading of local data). This includes for example popular mini-batch versions of online methods, such as stochastic subgradient (SGD) and coordinate descent (SDCA).

In this work, we target this bottleneck. We propose a distributed optimization framework that allows one to freely steer the trade-off between *communication* and *local computation*. In doing so, the framework can be easily adapted to the diverse spectrum of available large-scale computing systems, from high-latency commodity clusters to low-latency supercomputers or the multi-core setting.

Our new framework, CoCoA (**C**ommunication-efficient distributed dual **C**oordinate **A**scent), supports objectives for linear regularized loss minimization, encompassing a broad class of machine learning models. By leveraging the primal-dual structure of these optimization problems, CoCoA effectively combines partial results from local computation while avoiding conflict with updates simultaneously computed on other machines. In each round, CoCoA employs steps of an arbitrary dual optimization method on the local data on each machine, in parallel. A single update vector is then communicated to the master node. For example, when choosing to perform H iterations (usually order of the data size n) of an online optimization method locally per round, our scheme saves a factor of H in terms of communication compared to the corresponding naive distributed update

*Both authors contributed equally.

¹On typical computers, the latency for accessing data in main memory is in the order of 100 nanoseconds. In contrast, the latency for sending data over a standard network connection is around 250,000 nanoseconds.

scheme (i.e., updating a single point before communication). When processing the same number of datapoints, this is clearly a dramatic savings.

Our theoretical analysis (Section 4) shows that this significant reduction in communication cost comes with only a very moderate increase in the amount of total computation, in order to reach the same optimization accuracy. We show that, in general, the distributed CoCoA framework will inherit the convergence rate of the internally-used local optimization method. When using SDCA (randomized dual coordinate ascent) as the local optimizer and assuming smooth losses, this convergence rate is geometric.

In practice, our experiments with the method implemented on the fault-tolerant Spark platform [1] confirm both the clock time performance and huge communication savings of the proposed method on a variety distributed datasets. Our experiments consistently show order of magnitude gains over traditional mini-batch methods of both SGD and SDCA, and significant gains over the faster but theoretically less justified local SGD methods.

Related Work. As we discuss below (Section 5), our approach is distinguished from recent work on parallel and distributed optimization [2, 3, 4, 5, 6, 7, 8, 9] in that we provide a general framework for improving the communication efficiency of *any* dual optimization method. To the best of our knowledge, our work is the first to analyze the convergence rate for an algorithm with this level of communication efficiency, without making data-dependent assumptions. The presented analysis covers the case of smooth losses, but should also be extendable to the non-smooth case. Existing methods using mini-batches [4, 2, 10] are closely related, though our algorithm makes significant improvements by immediately applying all updates locally while they are processed, a scheme that is not considered in the classic mini-batch setting. This intuitive modification results in dramatically improved empirical results and also strengthens our theoretical convergence rate. More precisely, the convergence rate shown here only degrades with the number of workers K , instead of with the significantly larger mini-batch-size (typically order n) in the case of mini-batch methods.

Our method builds on a closely related recent line of work of [2, 3, 11, 12]. We generalize the algorithm of [2, 3] by allowing the use of arbitrary (dual) optimization methods as the local subroutine within our framework. In the special case of using coordinate ascent as the local optimizer, the resulting algorithm is very similar, though with a different computation of the coordinate updates. Moreover, we provide the first theoretical convergence rate analysis for such methods, without making strong assumptions on the data.

The proposed CoCoA framework in its basic variant is entirely free of tuning parameters or learning rates, in contrast to SGD-based methods. The only choice to make is the selection of the internal local optimization procedure, steering the desired trade-off between communication and computation. When choosing a primal-dual optimizer as the internal procedure, the duality gap readily provides a fair stopping criterion and efficient accuracy certificates during optimization.

Paper Outline. The rest of the paper is organized as follows. In Section 2 we describe the problem setting of interest. Section 3 outlines the proposed framework, CoCoA, and the convergence analysis of this method is presented in Section 4. We discuss related work in Section 5, and compare against several other state-of-the-art methods empirically in Section 6.

2 Setup

A large class of methods in machine learning and signal processing can be posed as the minimization of a convex loss function of linear predictors with a convex regularization term:

$$\min_{\mathbf{w} \in \mathbb{R}^d} \left[P(\mathbf{w}) := \frac{\lambda}{2} \|\mathbf{w}\|^2 + \frac{1}{n} \sum_{i=1}^n \ell_i(\mathbf{w}^T \mathbf{x}_i) \right], \quad (1)$$

Here the data training examples are real-valued vectors $\mathbf{x}_i \in \mathbb{R}^d$; the loss functions $\ell_i, i = 1, \dots, n$ are convex and depend possibly on labels $y_i \in \mathbb{R}$; and $\lambda > 0$ is the regularization parameter. Using the setup of [13], we assume the regularizer is the ℓ_2 -norm for convenience. Examples of this class of problems include support vector machines, as well as regularized linear and logistic regression, ordinal regression, and others.

The most popular method to solve problems of the form (1) is the *stochastic subgradient method* (SGD) [14, 15, 16]. In this setting, SGD becomes an online method where every iteration only requires access to a single data example (\mathbf{x}_i, y_i) , and the convergence rate is well-understood.

The associated conjugate *dual* problem of (1) takes the following form, and is defined over one dual variable per each example in the training set.

$$\max_{\boldsymbol{\alpha} \in \mathbb{R}^n} \left[D(\boldsymbol{\alpha}) := -\frac{\lambda}{2} \|A\boldsymbol{\alpha}\|^2 - \frac{1}{n} \sum_{i=1}^n \ell_i^*(-\alpha_i) \right], \quad (2)$$

where ℓ_i^* is the conjugate (Fenchel dual) of the loss function ℓ_i , and the data matrix $A \in \mathbb{R}^{d \times n}$ collects the (normalized) data examples $A_i := \frac{1}{\lambda n} \mathbf{x}_i$ in its columns. The duality comes with the convenient mapping from dual to primal variables $\mathbf{w}(\boldsymbol{\alpha}) := A\boldsymbol{\alpha}$ as given by the optimality conditions [13]. For any configuration of the dual variables $\boldsymbol{\alpha}$, we have the duality gap defined as $P(\mathbf{w}(\boldsymbol{\alpha})) - D(\boldsymbol{\alpha})$. This gap is a computable certificate of the approximation quality to the unknown true optimum $P(\mathbf{w}^*) = D(\boldsymbol{\alpha}^*)$, and therefore serves as a useful stopping criteria for algorithms.

For problems of the form (2), *coordinate descent* methods have proven to be very efficient, and come with several benefits over primal methods. In randomized *dual coordinate ascent* (SDCA), updates are made to the dual objective (2) by solving for one coordinate completely while keeping all others fixed. This algorithm has been implemented in a number of software packages (e.g. LibLinear [17]), and has proven very suitable for use in large-scale problems, while giving stronger convergence results than the primal-only methods (such as SGD), at the same iteration cost [13]. In addition to superior performance, this method also benefits from requiring no stepsize, and having a well-defined stopping criterion given by the duality gap.

3 Method Description

The CoCoA framework, as presented in Algorithm 1, assumes that the data $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ for a regularized loss minimization problem of the form (1) is distributed over K worker machines. We associate with the datapoints their corresponding dual variables $\{\alpha_i\}_{i=1}^n$, being partitioned between the workers in the same way. The core idea is to use the dual variables to efficiently merge the parallel updates from the different workers without much conflict, by exploiting the fact that they all work on disjoint sets of dual variables.

Algorithm 1: CoCoA: Communication-Efficient Distributed Dual Coordinate Ascent

Input: $T \geq 1$, scaling parameter $1 \leq \beta_K \leq K$ (default: $\beta_K := 1$).

Data: $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ distributed over K machines

Initialize: $\boldsymbol{\alpha}_{[k]}^{(0)} \leftarrow \mathbf{0}$ for all machines k , and $\mathbf{w}^{(0)} \leftarrow \mathbf{0}$

for $t = 1, 2, \dots, T$

for all machines $k = 1, 2, \dots, K$ *in parallel*

$(\Delta\boldsymbol{\alpha}_{[k]}, \Delta\mathbf{w}_k) \leftarrow \text{LOCALDUALMETHOD}(\boldsymbol{\alpha}_{[k]}^{(t-1)}, \mathbf{w}^{(t-1)})$

$\boldsymbol{\alpha}_{[k]}^{(t)} \leftarrow \boldsymbol{\alpha}_{[k]}^{(t-1)} + \frac{\beta_K}{K} \Delta\boldsymbol{\alpha}_{[k]}$

end

reduce $\mathbf{w}^{(t)} \leftarrow \mathbf{w}^{(t-1)} + \frac{\beta_K}{K} \sum_{k=1}^K \Delta\mathbf{w}_k$

end

In each round, the K workers in parallel perform some steps of an arbitrary optimization method, applied to their local data. This internal procedure tries to maximize the dual formulation (2), only with respect to their own local dual variables. We call this local procedure LOCALDUALMETHOD, as specified in the template Procedure A. Our core observation is that the necessary information each worker requires about the state of the other dual variables can be very compactly represented by a single primal vector $\mathbf{w} \in \mathbb{R}^d$, without ever sending around data or dual variables between the machines.

Allowing the subroutine to process more than one local data example per round dramatically reduces the amount of communication between the workers. By definition, CoCoA in each outer iteration

Procedure A: LOCALDUALMETHOD: Dual algorithm for prob. (2) on a single coordinate block k

Input: Local $\alpha_{[k]} \in \mathbb{R}^{n_k}$, and $\mathbf{w} \in \mathbb{R}^d$ consistent with other coordinate blocks of α s.t. $\mathbf{w} = A\alpha$

Data: Local $\{(\mathbf{x}_i, y_i)\}_{i=1}^{n_k}$

Output: $\Delta\alpha_{[k]}$ and $\Delta\mathbf{w} := A_{[k]}\Delta\alpha_{[k]}$

Procedure B: LOCALSDCA: SDCA iterations for problem (2) on a single coordinate block k

Input: $H \geq 1$, $\alpha_{[k]} \in \mathbb{R}^{n_k}$, and $\mathbf{w} \in \mathbb{R}^d$ consistent with other coordinate blocks of α s.t. $\mathbf{w} = A\alpha$

Data: Local $\{(\mathbf{x}_i, y_i)\}_{i=1}^{n_k}$

Initialize: $\mathbf{w}^{(0)} \leftarrow \mathbf{w}$, $\Delta\alpha_{[k]} \leftarrow \mathbf{0} \in \mathbb{R}^{n_k}$

for $h = 1, 2, \dots, H$

choose $i \in \{1, 2, \dots, n_k\}$ uniformly at random

find $\Delta\alpha$ maximizing $-\frac{\lambda n}{2}\|\mathbf{w}^{(h-1)} + \frac{1}{\lambda n}\Delta\alpha \mathbf{x}_i\|^2 - \ell_i^*(-(\alpha_i^{(h-1)} + \Delta\alpha))$

$\alpha_i^{(h)} \leftarrow \alpha_i^{(h-1)} + \Delta\alpha$

$(\Delta\alpha_{[k]})_i \leftarrow (\Delta\alpha_{[k]})_i + \Delta\alpha$

$\mathbf{w}^{(h)} \leftarrow \mathbf{w}^{(h-1)} + \frac{1}{\lambda n}\Delta\alpha \mathbf{x}_i$

end

Output: $\Delta\alpha_{[k]}$ and $\Delta\mathbf{w} := A_{[k]}\Delta\alpha_{[k]}$

only requires communication of a single vector for each worker, that is $\Delta\mathbf{w}_k \in \mathbb{R}^d$. Further, as we will show in Section 4, COCOA inherits the convergence guarantee of any algorithm run locally on each node in the inner loop of Algorithm 1. We suggest to use randomized dual coordinate ascent (SDCA) [13] as the internal optimizer in practice, as implemented in Procedure B, and also used in our experiments.

Notation. In the same way the data is partitioned across the K worker machines, we write the dual variable vector as $\alpha = (\alpha_{[1]}, \dots, \alpha_{[K]}) \in \mathbb{R}^n$ with the corresponding coordinate blocks $\alpha_{[k]} \in \mathbb{R}^{n_k}$ such that $\sum_k n_k = n$. The submatrix $A_{[k]}$ collects the columns of A (i.e. rescaled data examples) which are available locally on the k -th worker. The parameter T determines the number of outer iterations of the algorithm, while when using an online internal method such as LOCALSDCA, then the number of inner iterations H determines the computation-communication trade-off factor.

4 Convergence Analysis

Considering the dual problem (2), we define the local suboptimality on each coordinate block as:

$$\epsilon_{D,k}(\alpha) := \max_{\hat{\alpha}_{[k]} \in \mathbb{R}^{n_k}} D((\alpha_{[1]}, \dots, \hat{\alpha}_{[k]}, \dots, \alpha_{[K]})) - D((\alpha_{[1]}, \dots, \alpha_{[k]}, \dots, \alpha_{[K]})), \quad (3)$$

that is how far we are from the optimum on block k with all other blocks fixed. Note that this differs from the global suboptimality $\max_{\hat{\alpha}} D(\hat{\alpha}) - D((\alpha_{[1]}, \dots, \alpha_{[K]}))$.

Assumption 1 (Local Geometric Improvement of LOCALDUALMETHOD). *We assume that there exists $\Theta \in [0, 1)$ such that for any given α , LOCALDUALMETHOD when run on block k alone returns a (possibly random) update $\Delta\alpha_{[k]}$ such that*

$$\mathbf{E}[\epsilon_{D,k}((\alpha_{[1]}, \dots, \alpha_{[k-1]}, \alpha_{[k]} + \Delta\alpha_{[k]}, \alpha_{[k+1]}, \dots, \alpha_{[K]}))] \leq \Theta \cdot \epsilon_{D,k}(\alpha). \quad (4)$$

Note that this assumption is satisfied for several available implementations of the inner procedure LOCALDUALMETHOD, in particular for LOCALSDCA, as shown in the following Proposition.

From here on, we assume that the input data is scaled such that $\|\mathbf{x}_i\| \leq 1$ for all datapoints. Proofs of all statements are provided in the appendix.

Proposition 1. *Assume the loss functions ℓ_i are $(1/\gamma)$ -smooth. Then for using LOCALSDCA, Assumption 1 holds with*

$$\Theta = \left(1 - \frac{\lambda n \gamma}{1 + \lambda n \gamma \tilde{n}}\right)^H. \quad (5)$$

where $\tilde{n} := \max_k n_k$ is the size of the largest block of coordinates.

Theorem 2. Assume that Algorithm 1 is run for T outer iterations on K worker machines, with the procedure LOCALDUALMETHOD having local geometric improvement Θ , and let $\beta_K := 1$. Further, assume the loss functions ℓ_i are $(1/\gamma)$ -smooth. Then the following geometric convergence rate holds for the global (dual) objective:

$$\mathbf{E}[D(\boldsymbol{\alpha}^*) - D(\boldsymbol{\alpha}^{(T)})] \leq \left(1 - (1 - \Theta) \frac{1}{K} \frac{\lambda n \gamma}{\sigma + \lambda n \gamma}\right)^T \left(D(\boldsymbol{\alpha}^*) - D(\boldsymbol{\alpha}^{(0)})\right). \quad (6)$$

Here σ is any real number satisfying

$$\sigma \geq \sigma_{\min} := \max_{\boldsymbol{\alpha} \in \mathbb{R}^n} \lambda^2 n^2 \frac{\sum_{k=1}^K \|A_{[k]} \boldsymbol{\alpha}_{[k]}\|^2 - \|A \boldsymbol{\alpha}\|^2}{\|\boldsymbol{\alpha}\|^2} \geq 0. \quad (7)$$

Lemma 3. If $K = 1$ then $\sigma_{\min} = 0$. For any $K \geq 1$, when assuming $\|\mathbf{x}_i\| \leq 1 \ \forall i$, we have

$$0 \leq \sigma_{\min} \leq \tilde{n}.$$

Moreover, if datapoints between different workers are orthogonal, i.e. $(A^T A)_{i,j} = 0 \ \forall i, j$ such that i and j do not belong to the same part, then $\sigma_{\min} = 0$.

If we choose $K = 1$ then, Theorem 2 together with Lemma 3 implies that

$$\mathbf{E}[D(\boldsymbol{\alpha}^*) - D(\boldsymbol{\alpha}^{(T)})] \leq \Theta^T \left(D(\boldsymbol{\alpha}^*) - D(\boldsymbol{\alpha}^{(0)})\right),$$

as expected, showing that the analysis is tight in the special case $K = 1$. More interestingly, we observe that for any K , in the extreme case when the subproblems are solved to optimality (i.e. letting $H \rightarrow \infty$ in LOCALSDCA), then the algorithm as well as the convergence rate match that of serial/parallel block-coordinate descent [18, 19].

Note: If choosing the starting point as $\boldsymbol{\alpha}^{(0)} := \mathbf{0}$ as in the main algorithm, then it is known that $D(\boldsymbol{\alpha}^*) - D(\boldsymbol{\alpha}^{(0)}) \leq 1$ (see e.g. Lemma 20 in [13]).

5 Related Work

Distributed Primal-Dual Methods. Our approach is most closely related to recent work by [2, 3], which generalizes the distributed optimization method for linear SVMs as in [11] to the primal-dual setting considered here (which was introduced by [13]). The difference between our approach and the ‘practical’ method of [2] is that our internal steps directly correspond to coordinate descent iterations on the global dual objective (2), for coordinates in the current block, while in [3, Equation 8] and [2], the inner iterations apply to a slightly different notion of the sub-dual problem defined on the local data. In terms of convergence results, the analysis of [2] only addresses the mini-batch case without local updates, while the more recent paper [3] shows a convergence rate for a variant of COCOA with inner coordinate steps, but under the unrealistic assumption that the data is orthogonal between the different workers. In this case, the optimization problems become independent, so that an even simpler single-round communication scheme summing the individual resulting models \boldsymbol{w} would give an exact solution. Instead, we show a linear convergence rate for the full problem class of smooth losses, without any assumptions on the data, in the same generality as the non-distributed setting of [13].

While the experimental results in all papers [11, 2, 3] are encouraging for this type of method, they do not yet provide a quantitative comparison of the gains in communication efficiency, or compare to the analogous SGD schemes that use the same distribution and communication patterns, which is the main goal of our experiments in Section 6. For the special case of linear SVMs, the first paper to propose the same algorithmic idea was [11], which used LibLinear in the inner iterations. However, the proposed algorithm [11] processes the blocks sequentially (not in the parallel or distributed setting). Also, it is assumed that the subproblems are solved to near optimality on each block before selecting the next, making the method essentially standard block-coordinate descent. While no convergence rate was given, the empirical results in the journal paper [12] suggest that running LibLinear for just one pass through the local data performs well in practice. Here, we prove this, quantify the communication efficiency, and show that fewer local steps can improve the overall performance. For the LASSO case, [7] has proposed a parallel coordinate descent method converging to the true optimum, which could potentially also be interpreted in our framework here.

Mini-Batches. Another closely related avenue of research includes methods that use *mini-batches* to distribute updates. In these methods, a mini-batch, or sample, of the data examples is selected for processing at each iteration. All updates within the mini-batch are computed based on the same fixed parameter vector w , and then these updates are either added or averaged in a reduce step and communicated back to the worker machines. This concept has been studied for both SGD and SDCA, see e.g. [4, 10] for the SVM case. The so-called naive variant of [2] is essentially identical to mini-batch dual coordinate descent, with a slight difference in defining the sub-problems.

As is shown in [2] and below in Section 6, the performance of these algorithms suffers when processing large batch sizes, as they do not take local updates immediately into account. Furthermore, they are very sensitive to the choice of the parameter β_b , which controls the magnitude of combining all updates between $\beta_b := 1$ for (conservatively) averaging, and $\beta_b := b$ for (aggressively) adding the updates (here we denote b as the size of the selected mini-batch, which can be of size up to n). This instability is illustrated by the fact that even the change of $\beta_b := 2$ instead of $\beta_b := 1$ can lead to divergence of coordinate descent (SDCA) in the simple case of just two coordinates [4]. In practice it can be very difficult to choose the correct data-dependent parameter β_b especially for large mini-batch sizes $b \approx n$, as the parameter range spans many orders of magnitude, and directly controls the step size of the resulting algorithm, and therefore the convergence rate [20, 21]. For sparse data, the work of [20, 21] gives some data dependent choices of β_b which are safe.

Known convergence rates for the mini-batch methods degrade linearly with the growing batch size $b \approx \Theta(n)$. More precisely, the improvement in objective function per example processed degrades with a factor of β_b in [4, 20, 21]. In contrast, our convergence rate as shown in Theorem 2 only degrades with the much smaller number of worker machines K , which in practical applications is often several orders of magnitudes smaller than the mini-batch size b .

Single Round of Communication. One extreme is to consider methods with only a single round of communication (e.g. one map-reduce operation), as in [22, 6, 23]. The output of these methods is the average of K individual models, trained only on the local data on each machine. In [22], the authors give conditions on the data and computing environment under which these one-communication algorithms may be sufficient. In general, however, the true optimum of the original problem (1) is *not* the average of these K models, no matter how accurately the subproblems are solved [24].

Naive Distributed Online Methods, Delayed Gradients, and Multi-Core. On the other extreme, a natural way to distribute updates is to let every machine send updates to the master node (sometimes called the “parameter server”) as soon as they are performed. This is what we call the *naive distributed SGD / CD* in our experiments. The amount of communication for such naive distributed online methods is the same as the number of data examples processed. In contrast to this, the number of communicated vectors in our method is divided by H , that is the number of inner local steps performed per outer iteration, which can be $\Theta(n)$.

The early work of [25] introduced the nice framework of gradient updates where the gradients come with some delays, i.e. are based on outdated iterates, and shows some robust convergence rates. In the machine learning setting, [26] and the later work of [27] have provided additional insights into these types of methods. However, these papers study the case of smooth objective functions of a sum structure, and so do not directly apply to general case we consider here. In the same spirit, [5] implements SGD with communication-intense updates after each example processed, allowing asynchronous updates again with some delay. For coordinate descent, the analogous approach was studied in [28]. Both methods [5, 28] are H times less efficient in terms of communication when compared to COCOA, and are designed for *multi-core* shared memory machines (where communication is as fast as memory access). They require the same amount of communication as *naive distributed SGD / CD*, which we include in our experiments in Section 6, and a slightly larger number of iterations due to the asynchronicity. The $1/t$ convergence rate shown in [5] only holds under strong sparsity assumptions on the data. A more recent paper [29] deepens the understanding of such methods, but still only applies to very sparse data. For general data, [30] theoretically shows that $1/\varepsilon^2$ communications rounds of single vectors are enough to obtain ε -quality for linear classifiers, with the rate growing with K^2 in the number of workers. Our new analysis here makes the dependence on $1/\varepsilon$ logarithmic.

6 Experiments

In this section, we compare CoCoA to traditional mini-batch versions of stochastic dual coordinate ascent and stochastic gradient descent, as well as the locally-updating version of stochastic gradient descent. We implement mini-batch SDCA (denoted mini-batch-CD) as described in [4, 2]. The SGD-based methods are mini-batch and locally-updating versions of Pegasos [16], differing only in whether the primal vector is updated locally on each inner iteration or not, and whether the resulting combination/communication of the updates is by an average over the total size KH of the mini-batch (mini-batch-SGD) or just over the number of machines K (local-SGD). For each algorithm, we additionally study the effect of scaling the average by a parameter β_K , as first described in [4], while noting that it is a benefit to avoid having to tune this data-dependent parameter.

We apply these algorithms to standard hinge loss ℓ_2 -regularized support vector machines, using implementations written in Spark on m1.large Amazon EC2 instances [1]. Though this non-smooth case is not yet covered in our theoretical analysis, we still see remarkable empirical performance. Our results indicate that CoCoA is able to converge to .001-accurate solutions nearly $25\times$ as fast compared the other algorithms, when all use $\beta_K = 1$. The datasets used in these analyses are summarized in Table 1, and were distributed among $K = 4, 8,$ and 32 nodes, respectively. We use the same regularization parameters as specified in [16, 17].

Table 1: Datasets for Empirical Study

| Dataset | Training (n) | Features (d) | Sparsity | λ | Workers (K) |
|----------|------------------|------------------|----------|-----------|-----------------|
| cov | 522,911 | 54 | 22.22% | $1e-6$ | 4 |
| rcv1 | 677,399 | 47,236 | 0.16% | $1e-6$ | 8 |
| imagenet | 32,751 | 160,000 | 100% | $1e-5$ | 32 |

In comparing each algorithm and dataset, we analyze progress in primal objective value as a function of both time (Figure 1) and communication (Figure 2). For all competing methods, we present the result for the batch size (H) that yields the best performance in terms of reduction in objective value over time. For the locally-updating methods (CoCoA and local-SGD), these tend to be larger batch sizes corresponding to processing almost all of the local data at each outer step. For the non-locally updating mini-batch methods, (mini-batch SDCA [4] and mini-batch SGD [16]), these typically correspond to smaller values of H , as averaging the solutions to guarantee safe convergence becomes less of an impediment for smaller batch sizes.




Figure 1: Primal Suboptimality vs. Time for Best Mini-Batch Sizes (H): For $\beta_K = 1$, CoCoA converges more quickly than all other algorithms, even when accounting for different batch sizes.




Figure 2: Primal Suboptimality vs. # of Communicated Vectors for Best Mini-Batch Sizes (H): A clear correlation is evident between the number of communicated vectors and wall-time to convergence (Figure 1).

First, we note that there is a clear correlation between the wall-time spent processing each dataset and the number of vectors communicated, indicating that communication has a significant effect on convergence speed. We see clearly that CoCoA is able to converge to a more accurate solution in all datasets much faster than the other methods. On average, CoCoA reaches a .001-accurate solution for these datasets 25x faster than the best competitor. This is a testament to the algorithm’s ability to avoid communication while still making significant global progress by efficiently combining the local updates of each iteration. The improvements are robust for both regimes $n \gg d$ and $n \ll d$.




Figure 3: Effect of H on CoCoA. Figure 4: Best β_K Scaling Values for $H = 1e5$ and $H = 100$.

In Figure 3 we explore the effect of H , the computation-communication trade-off factor, on the convergence of CoCoA for the Cov dataset on a cluster of 4 nodes. As described above, increasing H decreases communication but also affects the convergence properties of the algorithm. In Figure 4, we attempt to scale the averaging step of each algorithm by using various β_K values, for two different batch sizes on the Cov dataset ($H = 1e5$ and $H = 100$). We see that though β_K has a larger impact on the smaller batch size, it is still not enough to improve the mini-batch algorithms beyond what is achieved by CoCoA and local-SGD.

7 Conclusion

We have presented a communication-efficient framework for distributed dual coordinate ascent algorithms that can be used to solve large-scale regularized loss minimization problems. This is crucial in settings where datasets must be distributed across multiple machines, and where communication amongst nodes is costly. We have shown that the proposed algorithm performs competitively on real-world, large-scale distributed datasets, and have presented the first theoretical analysis of this algorithm that achieves competitive convergence rates without making additional assumptions on the data itself.

It remains open to obtain improved convergence rates for more aggressive updates corresponding to $\beta_K > 1$, which might be suitable for using the ‘safe’ updates techniques of [4] and the related expected separable over-approximations of [18, 19], here applied to K instead of n blocks. Furthermore, it remains open to show convergence rates for local SGD in the same communication efficient setting as described here.

Acknowledgments. We thank Shivaram Venkataraman, Ameet Talwalkar, and Peter Richtárik for fruitful discussions. MJ acknowledges support by the Simons Institute for the Theory of Computing.

References

- [1] Matei Zaharia, Mosharaf Chowdhury, Tathagata Das, Ankur Dave, Murphy McCauley, Michael J Franklin, Scott Shenker, and Ion Stoica. Resilient Distributed Datasets: A Fault-Tolerant Abstraction for In-Memory Cluster Computing. *NSDI*, 2012.
- [2] Tianbao Yang. Trading Computation for Communication: Distributed Stochastic Dual Coordinate Ascent. *NIPS*, 2013.
- [3] Tianbao Yang, Shenghuo Zhu, Rong Jin, and Yuanqing Lin. On Theoretical Analysis of Distributed Stochastic Dual Coordinate Ascent. *arXiv:1312.1031*, December 2013.
- [4] Martin Takáč, Avleen Bijral, Peter Richtárik, and Nathan Srebro. Mini-Batch Primal and Dual Methods for SVMs. *ICML*, 2013.

- [5] Feng Niu, Benjamin Recht, Christopher Ré, and Stephen J Wright. Hogwild!: A Lock-Free Approach to Parallelizing Stochastic Gradient Descent. *NIPS*, 2011.
- [6] Martin A Zinkevich, Markus Weimer, Alex J Smola, and Lihong Li. Parallelized Stochastic Gradient Descent. *NIPS 23*, 2010.
- [7] Joseph K Bradley, Aapo Kyrola, Danny Bickson, and Carlos Guestrin. Parallel Coordinate Descent for L1-Regularized Loss Minimization. *ICML*, 2011.
- [8] Jakub Mareček, Peter Richtárik, and Martin Takáč. Distributed Block Coordinate Descent for Minimizing Partially Separable Functions. *arXiv:1408.2467*, June 2014.
- [9] Ion Necoara and Dragos Clipici. Efficient parallel coordinate descent algorithm for convex optimization problems with separable constraints: Application to distributed MPC. *Journal of Process Control*, 23(3):243–253, 2013.
- [10] Martin Takáč, Peter Richtárik, and Nathan Srebro. Primal-Dual Parallel Coordinate Descent for Machine Learning Optimization. *Manuscript*, 2014.
- [11] Hsiang-Fu Yu, Cho-Jui Hsieh, Kai-Wei Chang, and Chih-Jen Lin. Large linear classification when data cannot fit in memory. *the 16th ACM SIGKDD international conference*, page 833, 2010.
- [12] Hsiang-Fu Yu, Cho-Jui Hsieh, Kai-Wei Chang, and Chih-Jen Lin. Large Linear Classification When Data Cannot Fit in Memory. *ACM Transactions on Knowledge Discovery from Data*, 5(4):1–23, 2012.
- [13] Shai Shalev-Shwartz and Tong Zhang. Stochastic Dual Coordinate Ascent Methods for Regularized Loss Minimization. *JMLR*, 14:567–599, 2013.
- [14] Herbert Robbins and Sutton Monro. A Stochastic Approximation Method. *The Annals of Mathematical Statistics*, 22(3):400–407, 1951.
- [15] Léon Bottou. Large-Scale Machine Learning with Stochastic Gradient Descent. *COMPSTAT'2010 - Proceedings of the 19th International Conference on Computational Statistics*, pages 177–187, 2010.
- [16] Shai Shalev-Shwartz, Yoram Singer, Nathan Srebro, and Andrew Cotter. Pegasos: Primal Estimated Sub-Gradient Solver for SVM. *Mathematical Programming*, 127(1):3–30, 2010.
- [17] Cho-Jui Hsieh, Kai-Wei Chang, Chih-Jen Lin, S Sathya Keerthi, and S Sundararajan. A Dual Coordinate Descent Method for Large-scale Linear SVM. *ICML*, 2008.
- [18] Peter Richtárik and Martin Takáč. Iteration complexity of randomized block-coordinate descent methods for minimizing a composite function. *Mathematical Programming*, 144(1-2):1–38, April 2014.
- [19] Peter Richtárik and Martin Takáč. Parallel Coordinate Descent Methods for Big Data Optimization. *arXiv:1212.0873*, 2012.
- [20] Peter Richtárik and Martin Takáč. Distributed Coordinate Descent Method for Learning with Big Data. *arXiv:1310.2059*, 2013.
- [21] Olivier Fercoq, Zheng Qu, Peter Richtárik, and Martin Takáč. Fast Distributed Coordinate Descent for Non-Strongly Convex Losses. *IEEE Workshop on Machine Learning for Signal Processing*, May 2014.
- [22] Yuchen Zhang, John C Duchi, and Martin J Wainwright. Communication-Efficient Algorithms for Statistical Optimization. *JMLR*, 14:3321–3363, November 2013.
- [23] Gideon Mann, Ryan McDonald, Mehryar Mohri, Nathan Silberman, and Daniel D Walker. Efficient Large-Scale Distributed Training of Conditional Maximum Entropy Models. *NIPS*, 1231–1239, 2009.
- [24] Ohad Shamir, Nathan Srebro, and Tong Zhang. Communication-Efficient Distributed Optimization using an Approximate Newton-type Method. *ICML*, 32(1):1000–1008, 2014.
- [25] John N Tsitsiklis, Dimitri P Bertsekas, and Michael Athans. Distributed asynchronous deterministic and stochastic gradient optimization algorithms. *IEEE Trans. on Automatic Control*, 31(9):803–812, 1986.
- [26] Ofer Dekel, Ran Gilad-Bachrach, Ohad Shamir, and Lin Xiao. Optimal Distributed Online Prediction Using Mini-Batches. *JMLR*, 13:165–202, 2012.
- [27] Alekh Agarwal and John C Duchi. Distributed Delayed Stochastic Optimization. *NIPS*, 873–881, 2011.
- [28] Ji Liu, Stephen J Wright, Christopher Ré, Victor Bittorf, and Srikrishna Sridhar. An Asynchronous Parallel Stochastic Coordinate Descent Algorithm. *ICML*, 2014.
- [29] John C Duchi, Michael I Jordan, and H Brendan McMahan. Estimation, Optimization, and Parallelism when Data is Sparse. *NIPS*, 2013.
- [30] Maria-Florina Balcan, Avrim Blum, Shai Fine, and Yishay Mansour. Distributed Learning, Communication Complexity and Privacy. *COLT*, 23:26.1–26.22, 2012.

A Proof of Theorem 2 – Main Convergence Theorem

In the following, for a given vector $\alpha_{[k]} \in \mathbb{R}^{n_k}$, we write $\alpha_{\langle [k] \rangle} \in \mathbb{R}^n$ for its zero-padded version which coincides with $\alpha_{[k]}$ on the k -th coordinate block, and is zero everywhere else.

Theorem' 2. *Assume that Algorithm 1 is run for T outer iterations on K worker machines, with the procedure LOCALDUALMETHOD having local geometric improvement Θ , and let $\beta_K := 1$. Further, assume the loss functions ℓ_i are $(1/\gamma)$ -smooth. Then the following geometric convergence rate holds for the global (dual) objective:*

$$\mathbf{E}[D(\alpha^*) - D(\alpha^{(T)})] \leq \left(1 - (1 - \Theta) \frac{1}{K} \frac{\lambda n \gamma}{\sigma + \lambda n \gamma}\right)^T \left(D(\alpha^*) - D(\alpha^{(0)})\right).$$

Here σ is any real number satisfying

$$\sigma \geq \sigma_{\min} := \max_{\alpha \in \mathbb{R}^n} \lambda^2 n^2 \frac{\sum_{k=1}^K \|A_{[k]} \alpha_{[k]}\|^2 - \|A\alpha\|^2}{\|\alpha\|^2} \geq 0.$$

Proof. From the definition of the update performed by Algorithm 1 (for the setting of $\beta_K := 1$), we have $\alpha^{(t+1)} = \alpha^{(t)} + \frac{1}{K} \sum_{k=1}^K \Delta \alpha_{\langle [k] \rangle}$. Let us estimate the change of objective function after one outer iteration. Then using concavity of D we have

$$\begin{aligned} D(\alpha^{(t+1)}) &= D\left(\alpha^{(t)} + \frac{1}{K} \sum_{k=1}^K \Delta \alpha_{\langle [k] \rangle}\right) = D\left(\frac{1}{K} \sum_{k=1}^K (\alpha^{(t)} + \Delta \alpha_{\langle [k] \rangle})\right) \\ &\geq \frac{1}{K} \sum_{k=1}^K D(\alpha^{(t)} + \Delta \alpha_{\langle [k] \rangle}). \end{aligned}$$

Subtracting $D(\alpha^{(t)})$ from both sides and denoting by $\hat{\alpha}_{[k]}^*$ the local maximizer as in (3) we obtain

$$\begin{aligned} D(\alpha^{(t+1)}) - D(\alpha^{(t)}) &\geq \frac{1}{K} \sum_{k=1}^K [D(\alpha^{(t)} + \Delta \alpha_{\langle [k] \rangle}) - D(\alpha^{(t)})] \\ &= \frac{1}{K} \sum_{k=1}^K \left[D(\alpha^{(t)} + \Delta \alpha_{\langle [k] \rangle}) - D((\alpha_{[t]}^{(1)}, \dots, \hat{\alpha}_{[k]}^*, \dots, \alpha_{[t]}^{(K)})) \right] \\ &\quad + D((\alpha_{[t]}^{(1)}, \dots, \hat{\alpha}_{[k]}^*, \dots, \alpha_{[t]}^{(K)})) - D(\alpha^{(t)}) \\ &\stackrel{(3)}{=} \frac{1}{K} \sum_{k=1}^K [\varepsilon_{D,k}(\alpha^{(t)}) - \varepsilon_{D,k}(\alpha^{(t)} + \Delta \alpha_{\langle [k] \rangle})]. \end{aligned}$$

Considering the expectation of this quantity, we are now ready to use Assumption 1 on the *local geometric improvement* of the inner procedure. We have

$$\begin{aligned} \mathbf{E}[D(\alpha^{(t+1)}) - D(\alpha^{(t)}) \mid \alpha^{(t)}] &\geq \frac{1}{K} \sum_{k=1}^K \mathbf{E}[\varepsilon_{D,k}(\alpha^{(t)}) - \varepsilon_{D,k}(\alpha^{(t)} + \Delta \alpha_{\langle [k] \rangle}) \mid \alpha^{(t)}] \\ &\stackrel{(4)}{\geq} \frac{1}{K} (1 - \Theta) \sum_{k=1}^K \varepsilon_{D,k}(\alpha^{(t)}). \end{aligned}$$

It remains to bound $\sum_{k=1}^K \varepsilon_{D,k}(\alpha^{(t)})$.

$$\begin{aligned} \sum_{k=1}^K \varepsilon_{D,k}(\alpha^{(t)}) &\stackrel{(3)}{=} \max_{\hat{\alpha} \in \mathbb{R}^n} \left\{ \sum_{k=1}^K [D((\alpha_{[1]}, \dots, \hat{\alpha}_{[k]}, \dots, \alpha_{[K]})) - D((\alpha_{[1]}, \dots, \alpha_{[k]}, \dots, \alpha_{[K]}))] \right\} \\ &\stackrel{(2)}{=} \max_{\hat{\alpha} \in \mathbb{R}^n} \left\{ \frac{1}{n} \sum_{i=1}^n (-\ell_i^*(-\hat{\alpha}_i) + \ell_i^*(-\alpha_i^{(t)})) \right. \\ &\quad \left. + \frac{\lambda}{2} \sum_{k=1}^K \left[-\|A\alpha^{(t)} + A_{[k]}(\hat{\alpha}_{[k]} - \alpha_{[k]}^{(t)})\|^2 + \|A\alpha^{(t)}\|^2 \right] \right\} \\ &= \max_{\hat{\alpha} \in \mathbb{R}^n} \left\{ D(\hat{\alpha}) - D(\alpha^{(t)}) + \frac{\lambda}{2} \|A\hat{\alpha}\|^2 - \frac{\lambda}{2} \|A\alpha^{(t)}\|^2 \right. \\ &\quad \left. + \frac{\lambda}{2} \sum_{k=1}^K \left[-\|A\alpha^{(t)} + A_{[k]}(\hat{\alpha}_{[k]} - \alpha_{[k]}^{(t)})\|^2 + \|A\alpha^{(t)}\|^2 \right] \right\} \\ &= \max_{\hat{\alpha} \in \mathbb{R}^n} \left\{ D(\hat{\alpha}) - D(\alpha^{(t)}) + \frac{\lambda}{2} \left[\|A\hat{\alpha}\|^2 - \|A\alpha^{(t)}\|^2 \right] \right. \\ &\quad \left. + \frac{\lambda}{2} \left[-2(A\alpha^{(t)})^T A(\hat{\alpha} - \alpha^{(t)}) - \sum_{k=1}^K \|A_{[k]}(\hat{\alpha}_{[k]} - \alpha_{[k]}^{(t)})\|^2 \right] \right\} \end{aligned}$$

$$\begin{aligned}
&= \max_{\hat{\alpha} \in \mathbb{R}^n} \left\{ D(\hat{\alpha}) - D(\alpha^{(t)}) + \frac{\lambda}{2} \left[\|A\hat{\alpha}\|^2 + \|A\alpha^{(t)}\|^2 - 2(A\alpha^{(t)})^T A\hat{\alpha} \right] \right. \\
&\quad \left. + \frac{\lambda}{2} \left[-\sum_{k=1}^K \|A_{[k]}(\hat{\alpha}_{[k]} - \alpha_{[k]}^{(t)})\|^2 \right] \right\} \\
&= \max_{\hat{\alpha} \in \mathbb{R}^n} \left\{ D(\hat{\alpha}) - D(\alpha^{(t)}) + \frac{\lambda}{2} \left[\|A(\alpha^{(t)} - \hat{\alpha})\|^2 \right] \right. \\
&\quad \left. + \frac{\lambda}{2} \left[-\sum_{k=1}^K \|A_{[k]}(\hat{\alpha}_{[k]} - \alpha_{[k]}^{(t)})\|^2 \right] \right\} \\
&\geq \max_{\hat{\alpha} \in \mathbb{R}^n} \left\{ D(\hat{\alpha}) - D(\alpha^{(t)}) - \frac{\sigma}{2\lambda n^2} \|\hat{\alpha} - \alpha^{(t)}\|^2 \right\},
\end{aligned}$$

by the definition (7) of the complexity parameter σ . Now, we can conclude the bound as follows:

$$\begin{aligned}
\sum_{k=1}^K \varepsilon_{D,k}(\alpha^{(t)}) &\geq \max_{\hat{\alpha} \in \mathbb{R}^n} \left\{ D(\hat{\alpha}) - D(\alpha^{(t)}) - \frac{\sigma}{2\lambda n^2} \|\hat{\alpha} - \alpha^{(t)}\|^2 \right\} \\
&\geq \max_{\eta \in [0,1]} \left\{ D(\eta\alpha^* + (1-\eta)\alpha^{(t)}) - D(\alpha^{(t)}) - \frac{1}{2} \frac{\sigma\eta^2}{\lambda n^2} \|\alpha^* - \alpha^{(t)}\|^2 \right\} \\
&\geq \max_{\eta \in [0,1]} \left\{ \eta D(\alpha^*) + (1-\eta)D(\alpha^{(t)}) - D(\alpha^{(t)}) \right. \\
&\quad \left. + \frac{\gamma\eta(1-\eta)}{2n} \|\alpha^* - \alpha^{(t)}\|^2 - \frac{1}{2} \frac{\sigma\eta^2}{\lambda n^2} \|\alpha^* - \alpha^{(t)}\|^2 \right\} \\
&\geq \max_{\eta \in [0,1]} \left\{ \eta(D(\alpha^*) - D(\alpha^{(t)})) + \frac{\eta}{2n} \left(\gamma(1-\eta) - \frac{\sigma\eta}{\lambda n} \right) \|\alpha^* - \alpha^{(t)}\|^2 \right\}.
\end{aligned}$$

Choosing $\eta^* := \frac{\lambda n \gamma}{\sigma + \lambda n \gamma} \in [0, 1]$ gives

$$\sum_{k=1}^K \varepsilon_{D,k}(\alpha^{(t)}) \geq \frac{\lambda n \gamma}{\sigma + \lambda n \gamma} (D(\alpha^*) - D(\alpha^{(t)})).$$

Therefore, we have

$$\begin{aligned}
\mathbf{E}[D(\alpha^{(t+1)}) - D(\alpha^*) + D(\alpha^*) - D(\alpha^{(t)}) | \alpha^{(t)}] &\geq (1-\Theta) \frac{1}{K} \frac{\lambda n \gamma}{\sigma + \lambda n \gamma} (D(\alpha^*) - D(\alpha^{(t)})), \\
\mathbf{E}[D(\alpha^*) - D(\alpha^{(t+1)}) | \alpha^{(t)}] &\leq \left[1 - (1-\Theta) \frac{1}{K} \frac{\lambda n \gamma}{\sigma + \lambda n \gamma} \right] (D(\alpha^*) - D(\alpha^{(t)})).
\end{aligned}$$

This implies the claim (6) of the theorem. \square

B Proof of Proposition 1 – Local Improvement of the Inner Optimizer

B.1 Decomposition of the Duality Structure over the Blocks of Coordinates

The core concept in our new analysis technique is the following primal-dual structure on each local coordinate block. Using this, we will show below that all steps of LOCALDUALMETHOD can be interpreted as performing coordinate ascent steps on the *global* dual objective function $D(\alpha)$, with the coordinates changes restricted to the k -th block, as follows.

For concise notation, we write $\{\mathcal{I}_k\}_{k=1}^K$ for the partition of data indices $\{1, 2, \dots, n\}$ between the workers, such that $|\mathcal{I}_k| = n_k$. In other words, each set \mathcal{I}_k consists of the indices of the k -th block of coordinates (those with the dual variables $\alpha_{[k]}$ and datapoints $A_{[k]}$, as available on the k -th worker).

Let us define the *local dual problem* as

$$\max_{\alpha_{[k]} \in \mathbb{R}^{n_k}} \left[D_k(\alpha_{[k]}; \bar{\mathbf{w}}) := -\frac{\lambda}{2} \|\bar{\mathbf{w}} + A_{[k]} \alpha_{[k]}\|^2 - \frac{1}{n} \sum_{i \in \mathcal{I}_k} \ell_i^*(-\alpha_i) + \frac{\lambda}{2} \|\bar{\mathbf{w}}\|^2 \right]. \quad (8)$$

The definition is valid for any fixed vector $\bar{\mathbf{w}}$. Here the reader should think of $\bar{\mathbf{w}}$ as representing the status of those dual variables α_i which are *not* part of the active block k . The idea is the following: For the choice of $\bar{\mathbf{w}} := \sum_{k' \neq k} A_{[k']} \alpha_{[k']}^{(0)}$, it turns out that the local dual objective, as a function of the local variables $\alpha_{[k]}$, is identical the global (dual) objective, up to a constant independent of $\alpha_{[k]}$, or formally

$$D_k(\alpha_{[k]}; \bar{\mathbf{w}}) = D(\alpha_{[1]}^{(0)}, \dots, \alpha_{[k]}, \dots, \alpha_{[K]}^{(0)}) + C' \quad \forall \alpha_{[k]} \in \mathbb{R}^{n_k}.$$

The following proposition observes that the defined local problem has a very similar duality structure as the original problem (1) with its dual (2).

Proposition 4. For $\bar{\mathbf{w}} \in \mathbb{R}^d$, let us define the “local” primal problem on the k -th block as

$$\min_{\mathbf{w}_k \in \mathbb{R}^d} \left[P_k(\mathbf{w}_k; \bar{\mathbf{w}}) := \frac{1}{n} \sum_{i \in \mathcal{I}_k} \ell_i((\bar{\mathbf{w}} + \mathbf{w}_k)^T \mathbf{x}_i) + \frac{\lambda}{2} \|\mathbf{w}_k\|^2 \right]. \quad (9)$$

Then the dual of this formulation (with respect to the variable \mathbf{w}_k) is given by the local dual problem (8) for the k -th coordinate block.

Proof. The dual of this problem is derived by plugging in the definition of the conjugate function $\ell_i((\bar{\mathbf{w}} + \mathbf{w}_k)^T \mathbf{x}_i) = \max_{\alpha_i} -\alpha_i(\bar{\mathbf{w}} + \mathbf{w}_k)^T \mathbf{x}_i - \ell_i^*(-\alpha_i)$, which gives

$$\begin{aligned} & \min_{\mathbf{w}_k \in \mathbb{R}^d} \frac{1}{n} \sum_{i \in \mathcal{I}_k} \max_{\alpha_i} \left(-\alpha_i(\bar{\mathbf{w}} + \mathbf{w}_k)^T \mathbf{x}_i - \ell_i^*(-\alpha_i) \right) + \frac{\lambda}{2} \|\mathbf{w}_k\|^2 \\ &= \frac{1}{n} \sum_{i \in \mathcal{I}_k} \max_{\alpha_i} -\ell_i^*(-\alpha_i) + \min_{\mathbf{w}_k \in \mathbb{R}^d} \left[-\alpha_i(\bar{\mathbf{w}} + \mathbf{w}_k)^T \mathbf{x}_i + \frac{\lambda}{2} \|\mathbf{w}_k\|^2 \right] \\ &= \max_{\alpha \in \mathbb{R}^{n_k}} \frac{1}{n} \sum_{i \in \mathcal{I}_k} -\ell_i^*(-\alpha_i) + \min_{\mathbf{w}_k \in \mathbb{R}^d} \left[-\frac{1}{n} \sum_{i \in \mathcal{I}_k} \alpha_i(\bar{\mathbf{w}} + \mathbf{w}_k)^T \mathbf{x}_i + \frac{\lambda}{2} \|\mathbf{w}_k\|^2 \right]. \end{aligned}$$

The first-order optimality condition for \mathbf{w}_k , by setting its derivative to zero in the inner minimization, can be written as

$$\mathbf{w}_k^* = \frac{1}{\lambda n} \sum_{i \in \mathcal{I}_k} \alpha_i \mathbf{x}_i = A_{[k]} \boldsymbol{\alpha}_{[k]}$$

plugging this back in, we have that the inner minimization becomes

$$\begin{aligned} & -\lambda \frac{1}{\lambda n} \sum_{i \in \mathcal{I}_k} \alpha_i (\bar{\mathbf{w}} + \mathbf{w}_k^*)^T \mathbf{x}_i + \frac{\lambda}{2} \|\mathbf{w}_k^*\|^2 \\ &= -\lambda (\bar{\mathbf{w}} + \mathbf{w}_k^*)^T \mathbf{w}_k^* + \frac{\lambda}{2} \|\mathbf{w}_k^*\|^2 \end{aligned}$$

Writing the resulting full problem, we obtain precisely the local dual problem (8) for the k -th coordinate block. \square

Using these local subproblems, we have the following nice structure of local and global duality gaps:

$$g_k(\boldsymbol{\alpha}) := P_k(\mathbf{w}_k; \bar{\mathbf{w}}) - D_k(\boldsymbol{\alpha}_{[k]}; \bar{\mathbf{w}}) \quad \text{and} \quad g(\boldsymbol{\alpha}) := P(\mathbf{w}(\boldsymbol{\alpha})) - D(\boldsymbol{\alpha})$$

Here the contributions of all blocks except the active one are collected in $\bar{\mathbf{w}} := \mathbf{w} - \mathbf{w}_k$. Recall that we defined $\mathbf{w}_k := A_{[k]} \boldsymbol{\alpha}_{[k]}$ and $\mathbf{w} := A\boldsymbol{\alpha}$, so that $\mathbf{w} = \sum_{k=1}^K \mathbf{w}_k$.

B.2 Local Convergence of LOCALSDCA

This section is just a small modification of results obtained in [13].

Observe that the coordinate step performed by an iteration of LOCALSDCA can equivalently be written as

$$\Delta \alpha_i^* = \Delta \alpha_i^*(\boldsymbol{\alpha}_{[k]}^{(h-1)}, \bar{\mathbf{w}}) := \operatorname{argmax}_{\Delta \alpha_i} \left(D_k(\boldsymbol{\alpha}_{[k]}^{(h-1)} + e_i \Delta \alpha_i; \bar{\mathbf{w}}) \right). \quad (10)$$

For a vector $\bar{\mathbf{w}} \in \mathbb{R}^d$. In other words, the step will optimize one of the local coordinates with respect to the local dual objective, which is identical to the global dual objective when the coordinates of the other blocks are kept fixed, as we have seen in the previous subsection.

Lemma 5. Assume that ℓ_i^* is γ -strongly convex (where $\gamma \geq 0$). Then for all iterations h of LOCALSDCA and any $s \in [0, 1]$ we have

$$\mathbf{E}[D_k(\boldsymbol{\alpha}_{[k]}^{(h)}; \bar{\mathbf{w}}) - D_k(\boldsymbol{\alpha}_{[k]}^{(h-1)}; \bar{\mathbf{w}})] \geq \frac{s}{n_k} (P_k(\mathbf{w}_k^{(h-1)}; \bar{\mathbf{w}}) - D_k(\boldsymbol{\alpha}_{[k]}^{(h-1)}; \bar{\mathbf{w}})) - \frac{s^2}{2\lambda n^2} G^{(h)}, \quad (11)$$

where $G^{(h)} := \frac{1}{n_k} \sum_{i \in \mathcal{I}_k} \left(\|\mathbf{x}_i\|^2 - \frac{\lambda n \gamma (1-s)}{s} \right) (u_i^{(h)} - \alpha_i^{(h-1)})^2$, $-u_i^{(h-1)} \in \partial \ell_i(\mathbf{x}_i^T \mathbf{w}^{(h-1)})$ and $\mathbf{w}^{(h-1)} = \bar{\mathbf{w}} + A_{[k]} \boldsymbol{\alpha}_{[k]}^{(h-1)}$.

Proof.

$$\begin{aligned} & n \left[D_k(\boldsymbol{\alpha}_{[k]}^{(h)}; \bar{\mathbf{w}}) - D_k(\boldsymbol{\alpha}_{[k]}^{(h-1)}; \bar{\mathbf{w}}) \right] \\ &= \underbrace{-\ell_i^*(-\alpha_i^{(h)}) - \frac{\lambda n}{2} \|\mathbf{w}^{(h)}\|^2}_A - \underbrace{\left(-\ell_i^*(-\alpha_i^{(h-1)}) - \frac{\lambda n}{2} \|\mathbf{w}^{(h-1)}\|^2 \right)}_B. \end{aligned}$$

By the definition of the update (10) we have for all $s \in [0, 1]$ that

$$\begin{aligned} A &= \max_{\Delta \alpha_i} -\ell_i^*(-\alpha_i^{(h-1)} - \Delta \alpha_i) - \frac{\lambda n}{2} \|\mathbf{w}^{(h-1)} + \frac{1}{\lambda n} \Delta \alpha_i \mathbf{x}_i\|^2 \\ &\geq -\ell_i^*(-\alpha_i^{(h-1)} - s(u_i^{(h)} - \alpha_i^{(h-1)})) - \frac{\lambda n}{2} \|\mathbf{w}^{(h-1)} + \frac{1}{\lambda n} s(u_i^{(h)} - \alpha_i^{(h-1)}) \mathbf{x}_i\|^2. \end{aligned}$$

From strong convexity we have

$$\begin{aligned} & \ell_i^*(-\alpha_i^{(h-1)} - s(u_i^{(h)} - \alpha_i^{(h-1)})) \\ &\leq s\ell_i^*(-u_i^{(h)}) + (1-s)\ell_i^*(-\alpha_i^{(h-1)}) - \frac{\gamma}{2} s(1-s)(u_i^{(h)} - \alpha_i^{(h-1)})^2. \end{aligned} \quad (12)$$

Hence

$$\begin{aligned} A &\stackrel{(12)}{\geq} -s\ell_i^*(-u_i^{(h)}) - (1-s)\ell_i^*(-\alpha_i^{(h-1)}) + \frac{\gamma}{2} s(1-s)(u_i^{(h)} - \alpha_i^{(h-1)})^2 \\ &\quad - \frac{\lambda n}{2} \|\mathbf{w}^{(h-1)} + \frac{1}{\lambda n} s(u_i^{(h)} - \alpha_i^{(h-1)}) \mathbf{x}_i\|^2 \\ &= \underbrace{-s(\ell_i^*(-u_i^{(h)}) + u_i^{(h)} \mathbf{x}_i^T \mathbf{w}^{(h-1)})}_{s\ell(\mathbf{x}_i^T \mathbf{w}^{(h-1)})} + \underbrace{(-\ell_i^*(-\alpha_i^{(h-1)}) - \frac{\lambda n}{2} \|\mathbf{w}^{(h-1)}\|^2)}_B \\ &\quad + \frac{s}{2} \left(\gamma(1-s) - \frac{1}{\lambda n} s \|\mathbf{x}_i\|^2 \right) (u_i^{(h)} - \alpha_i^{(h-1)})^2 + s(\ell_i^*(-\alpha_i^{(h-1)}) + \alpha_i^{(h-1)} \mathbf{x}_i^T \mathbf{w}^{(h-1)}). \end{aligned}$$

Therefore

$$\begin{aligned} A - B &\geq s \left[\ell(\mathbf{x}_i^T \mathbf{w}^{(h-1)}) + \ell_i^*(-\alpha_i^{(h-1)}) + \alpha_i^{(h-1)} \mathbf{x}_i^T \mathbf{w}^{(h-1)} \right. \\ &\quad \left. + \frac{1}{2} \left(\gamma(1-s) - \frac{1}{\lambda n} s \|\mathbf{x}_i\|^2 \right) (u_i^{(h)} - \alpha_i^{(h-1)})^2 \right]. \end{aligned} \quad (13)$$

Recall that our above definition of the local pair of primal and dual problems gives

$$P_k(\mathbf{w}_k; \bar{\mathbf{w}}) - D_k(\boldsymbol{\alpha}_{[k]}; \bar{\mathbf{w}}) = \frac{1}{n} \sum_{i \in \mathcal{I}_k} (\ell_i((\bar{\mathbf{w}} + \mathbf{w}_k)^T \mathbf{x}_i) + \ell_i^*(-\alpha_i) + \alpha_i(\bar{\mathbf{w}} + \mathbf{w}_k)^T \mathbf{x}_i).$$

where $\mathbf{w}_k := A_{[k]} \boldsymbol{\alpha}_{[k]}$.

If we take the expectation of (13) we obtain

$$\begin{aligned} \frac{1}{s} \mathbf{E}[A - B] &\geq \frac{n}{n_k} \frac{1}{n} \sum_{i \in \mathcal{I}_k} \underbrace{\left[\ell(\mathbf{x}_i^T \mathbf{w}^{(h-1)}) + \ell_i^*(-\alpha_i^{(h-1)}) + \alpha_i^{(h-1)} \mathbf{x}_i^T \mathbf{w}^{(h-1)} \right]}_{P_k(\mathbf{w}_k^{(h-1)}; \bar{\mathbf{w}}) - D_k(\boldsymbol{\alpha}_{[k]}^{(h-1)}; \bar{\mathbf{w}})} \\ &\quad - \underbrace{\frac{s}{2\lambda n} \frac{1}{n_k} \sum_{i \in \mathcal{I}_k} \left(\|\mathbf{x}_i\|^2 - \frac{\lambda n \gamma (1-s)}{s} \right)}_{G^{(h)}} (u_i^{(h)} - \alpha_i^{(h-1)})^2. \end{aligned}$$

Therefore, we have obtained the claimed improvement bound

$$\frac{n}{s} \mathbf{E}[D_k(\boldsymbol{\alpha}_{[k]}^{(h)}; \bar{\mathbf{w}}) - D_k(\boldsymbol{\alpha}_{[k]}^{(h-1)}; \bar{\mathbf{w}})] \geq \frac{n}{n_k} (P_k(\mathbf{w}_k^{(h-1)}; \bar{\mathbf{w}}) - D_k(\boldsymbol{\alpha}_{[k]}^{(h-1)}; \bar{\mathbf{w}})) - \frac{s}{2\lambda n} G^{(h)}. \quad \square$$

Proposition' 1. Assume the loss functions ℓ_i are $(1/\gamma)$ -smooth. Then for using LOCALSDCA, Assumption 1 holds with

$$\Theta = \left(1 - \frac{\lambda n \gamma}{1 + \lambda n \gamma \tilde{n}} \right)^H. \quad (14)$$

where $\tilde{n} := \max_k n_k$ is the size of the largest block of coordinates.

Lemma 6 (Local Convergence on the Subproblem). For any $\boldsymbol{\alpha}_{[k]}^{(0)} \in \mathbb{R}^{n_k}$ and $\bar{\mathbf{w}} \in \mathbb{R}^d$ let us define

$$\boldsymbol{\alpha}_{[k]}^{(*)} := \operatorname{argmax}_{\boldsymbol{\alpha}_{[k]} \in \mathbb{R}^{n_k}} D_k(\boldsymbol{\alpha}_{[k]}; \bar{\mathbf{w}}). \quad (15)$$

If LOCALSDCA is used for H iterations on block k , then

$$\mathbf{E} [D_k(\boldsymbol{\alpha}_{[k]}^{(*)}; \bar{\mathbf{w}}) - D_k(\boldsymbol{\alpha}_{[k]}^{(H)}; \bar{\mathbf{w}})] \leq \left(1 - \frac{s}{n_k} \right)^H \left(D_k(\boldsymbol{\alpha}_{[k]}^{(*)}; \bar{\mathbf{w}}) - D_k(\boldsymbol{\alpha}_{[k]}^{(0)}; \bar{\mathbf{w}}) \right). \quad (16)$$

Proof. We will use Lemma 5 with $s := \frac{\lambda n \gamma}{1 + \lambda n \gamma} \in [0, 1]$. Because $\|\mathbf{x}_i\| \leq 1$, we have $G^{(h)} \leq 0$. Therefore

$$\mathbf{E} [D_k(\boldsymbol{\alpha}_{[k]}^{(h)}; \bar{\mathbf{w}}) - D_k(\boldsymbol{\alpha}_{[k]}^{(h-1)}; \bar{\mathbf{w}})] \geq \frac{s}{n_k} \left(P_k(\mathbf{w}_k^{(h-1)}; \bar{\mathbf{w}}) - D_k(\boldsymbol{\alpha}_{[k]}^{(h-1)}; \bar{\mathbf{w}}) \right). \quad (17)$$

Following the proof of Theorem 5 in [13] we obtain the claimed bound. \square

Now, to prove Proposition 1 it is enough to observe that for fixed k and any $\boldsymbol{\alpha} = (\boldsymbol{\alpha}_{[1]}, \dots, \boldsymbol{\alpha}_{[K]}) \in \mathbb{R}^n$, assuming we define $\bar{\mathbf{w}} = \sum_{k' \neq k} A_{[k']} \boldsymbol{\alpha}_{[k']}$, it holds that

$$\begin{aligned} \varepsilon_{D,k}(\boldsymbol{\alpha}) &\stackrel{(3)}{=} \max_{\hat{\boldsymbol{\alpha}}_{[k]} \in \mathbb{R}^{n_k}} D((\boldsymbol{\alpha}_{[1]}, \dots, \hat{\boldsymbol{\alpha}}_{[k]}, \dots, \boldsymbol{\alpha}_{[K]})) - D((\boldsymbol{\alpha}_{[1]}, \dots, \boldsymbol{\alpha}_{[k]}, \dots, \boldsymbol{\alpha}_{[K]})) \\ &= D((\boldsymbol{\alpha}_{[1]}, \dots, \boldsymbol{\alpha}_{[k]}^{(*)}, \dots, \boldsymbol{\alpha}_{[K]})) - D((\boldsymbol{\alpha}_{[1]}, \dots, \boldsymbol{\alpha}_{[k]}, \dots, \boldsymbol{\alpha}_{[K]})) \\ &= D_k(\boldsymbol{\alpha}_{[k]}^{(*)}; \bar{\mathbf{w}}) - D_k(\boldsymbol{\alpha}_{[k]}; \bar{\mathbf{w}}), \end{aligned}$$

where $\boldsymbol{\alpha}_{[k]}^{(*)}$ is defined by (15).

C Proof of Lemma 3 – The Problem Complexity Parameter σ_{\min}

Lemma' 3. If $K = 1$ then $\sigma_{\min} = 0$. For any $K \geq 1$, when assuming $\|\mathbf{x}_i\| \leq 1 \forall i$, we have

$$0 \leq \sigma_{\min} \leq \tilde{n}.$$

Moreover, if datapoints between different workers are orthogonal, i.e. $(A^T A)_{i,j} = 0 \forall i, j$ such that i and j do not belong to the same part, then $\sigma_{\min} = 0$.

Proof. If $K = 1$ then $\alpha_{[K]} \equiv \alpha$ and hence $\sigma_{\min} = 0$. For a non-trivial case, when $K > 1$ we have

$$\begin{aligned}
\sigma_{\min} &\stackrel{(7)}{=} \max_{\alpha \in \mathbb{R}^n} \lambda^2 n^2 \frac{\sum_{k=1}^K \|A_{[k]} \alpha_{[k]}\|^2 - \|A \alpha\|^2}{\|\alpha\|^2} \\
&= \max_{\alpha \in \mathbb{R}^n} \frac{\sum_{k=1}^K \|X_{[k]} \alpha_{[k]}\|^2 - \|X \alpha\|^2}{\|\alpha\|^2} \\
&\leq \max_{\alpha \in \mathbb{R}^n} \frac{\sum_{k=1}^K \|X_{[k]} \alpha_{[k]}\|^2}{\sum_{k=1}^K \|\alpha_{[k]}\|^2} \leq \max_{\alpha \in \mathbb{R}^n} \frac{\sum_{k=1}^K \tilde{n} \|\alpha_{[k]}\|^2}{\sum_{k=1}^K \|\alpha_{[k]}\|^2} \leq \tilde{n},
\end{aligned}$$

where we have used the definition of the rescaled data matrix $A = \frac{1}{\lambda^n} X$. The last inequality follows from the fact that since the datapoints have bounded norm $\|\mathbf{x}_i\| \leq 1$, so that for all parts $k \in \{1, \dots, K\}$ we have

$$\max_{\alpha_{[k]} \in \mathbb{R}^{n_k}} \frac{\|X_{[k]} \alpha_{[k]}\|^2}{\|\alpha_{[k]}\|^2} = \|X_{[k]}\|_{\text{op}}^2 \leq \|X_{[k]}\|_{\text{frob}}^2 \leq n_k \leq \tilde{n}, \text{ where } \tilde{n} := \max_k n_k.$$

The case when $A^T A$ is block-diagonal (with respect to the partition) is trivial. \square