
Parallel and Distributed Block-Coordinate Frank-Wolfe Algorithms

Yu-Xiang Wang [†]
Veeranjaneyulu Sadhanala [†]
Wei Dai [†]
Willie Neiswanger [†]
Suvrit Sra [‡]
Eric P. Xing [†]

YUXIANGW@CS.CMU.EDU
VSADHANA@CS.CMU.EDU
WDAI@CS.CMU.EDU
WILLIE@CS.CMU.EDU
SUVRIT@MIT.EDU
EPXING@CS.CMU.EDU

[†] Carnegie Mellon University, 5000 Forbes Ave, PA 15213, USA

[‡] Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, MA 02139, USA

Abstract

We study parallel and distributed Frank-Wolfe algorithms; the former on shared memory machines with mini-batching, and the latter in a delayed update framework. In both cases, we perform computations asynchronously whenever possible. We assume block-separable constraints as in Block-Coordinate Frank-Wolfe (BCFW) method (Lacoste-Julien et al., 2013), but our analysis subsumes BCFW and reveals problem-dependent quantities that govern the speedups of our methods over BCFW. A notable feature of our algorithms is that they do not depend on worst-case bounded delays, but only (mildly) on *expected* delays, making them robust to stragglers and faulty worker threads. We present experiments on structural SVM and Group Fused Lasso, and observe significant speedups over competing state-of-the-art (and synchronous) methods.

1. Introduction

The classical Frank-Wolfe (FW) algorithm (Frank & Wolfe, 1956) has witnessed a huge surge of interest recently (Ahipasaoglu et al., 2008; Clarkson, 2010; Jaggi, 2011; 2013). The FW algorithm iteratively minimizes a smooth function f (typically convex) over a compact convex set $\mathcal{M} \subset \mathbb{R}^m$. Unlike methods based on projection, FW uses just a *linear oracle* that solves $\min_{x \in \mathcal{M}} \langle x, g \rangle$, which can be much simpler and faster than projection.

This feature underlies the great popularity of FW, which

has by now witnessed several extensions such as regularized FW (Bredies et al., 2009; Harchaoui et al., 2015; Zhang et al., 2013), linearly convergent special cases (Garber & Hazan, 2013; Lacoste-Julien & Jaggi, 2015), stochastic versions (Hazan & Kale, 2012; Lafond et al., 2015; Ouyang & Gray, 2010), and a randomized block-coordinate FW (Lacoste-Julien et al., 2013).

Despite this progress, parallel and distributed FW variants are barely known. We fill this gap and develop new asynchronous FW algorithms, for the particular setting where the constraint set \mathcal{M} is *block-separable*; thus, we solve

$$\min_x f(x) \text{ s.t. } x = [x_{(1)}, \dots, x_{(n)}] \in \prod_{i=1}^n \mathcal{M}_i, \quad (1)$$

where $\mathcal{M}_i \subset \mathbb{R}^{m_i}$ ($1 \leq i \leq n$) is a compact convex set and $x_{(i)}$ are coordinate partitions of x . This setting for FW was considered in Lacoste-Julien et al. (2013), who introduced the Block-Coordinate Frank-Wolfe (BCFW) method.

Such problems arise in many applications, notably, structural SVMs (Lacoste-Julien et al., 2013), routing (LeBlanc et al., 1975), group fused lasso (Alaíz et al., 2013; Bleakley & Vert, 2011), trace-norm based tensor completion (Liu et al., 2013), reduced rank nonparametric regression (Foygel et al., 2012), and structured submodular minimization (Jegelka et al., 2013), among others.

A standard approach to solve (1) is via block-coordinate (gradient) descent (BCD), which forms a local quadratic model for a block of variables, and then solves a *projection* subproblem (Beck & Tetruashvili, 2013; Nesterov, 2012; Richtárik & Takáč, 2015). However, for many problems, including the ones noted above, projection can be expensive (e.g., projecting onto the trace norm ball, onto base polytopes Fujishige & Isotani, 2011), and in some cases even computationally intractable (Collins et al., 2008).

Frank-Wolfe methods excel in such scenarios as they rely

only on linear oracles that solve $\min_{s \in \mathcal{M}} \langle s, \nabla f(\cdot) \rangle$. For $\mathcal{M} = \prod_i \mathcal{M}_i$, this breaks into the n independent problems

$$\min_{s_{(i)} \in \mathcal{M}_i} \langle s_{(i)}, \nabla_{(i)} f(x) \rangle, \quad 1 \leq i \leq n, \quad (2)$$

where $\nabla_{(i)}$ denotes the gradient w.r.t. coordinates $x_{(i)}$. It is obvious that these n subproblems can be solved in parallel (an idea dating back to at least as early as LeBlanc et al., 1975). However, having to update all the coordinates at each iteration is expensive, hampering the use of FW on big-data problems.

This drawback is partially ameliorated by BCFW (Lacoste-Julien et al., 2013), which randomly selects a block \mathcal{M}_i at each iteration and performs FW updates. But these updates are *strictly* sequential, and do not take advantage of modern multicore architectures or of distributed clusters.

Contributions. Our main contributions are the following:

- Asynchronous Parallel block-coordinate Frank-Wolfe algorithms (AP-BCFW) for both shared-memory and distributed settings. Moreover, AP-BCFW depends only (mildly) on the *expected* delay, therefore is robust to stragglers and faulty worker threads.
- An analysis of the primal and primal-dual convergence of AP-BCFW and its variants for any minibatch size and potentially unbounded maximum delay. When the maximum delay is actually bounded, we show stronger results using results from load-balancing on max-load bounds.
- Insightful deterministic conditions under which minibatching *provably* improves the convergence rate for a class of problems (sometimes by orders of magnitude).
- Experiments that demonstrate on real data how our algorithm solves a structural SVM problem several times faster than the state-of-the-art.

In short, our results contribute towards making FW more attractive for big-data applications. To add perspective, we compare our methods to closely related works below; we refer the reader to Freund & Grigas (2014); Jaggi (2013); Lacoste-Julien et al. (2013); Zhang et al. (2012) for additional notes and references.

BCFW and Structural SVM. Our algorithm AP-BCFW extends and generalizes BCFW to parallel computation. Our analysis follows the structure in (Lacoste-Julien et al., 2013), but uses different stepsizes that must be carefully chosen. Our results contain BCFW as a special case. Lacoste-Julien et al. (2013) primarily focus on more explicit (and stronger) guarantee for BCFW on structural SVM, while we mainly focus on a more general class of problems; the particular subroutine needed by structural SVM requires special treatment though (see Appendix C).

Parallelization of sequential algorithms. The idea of parallelizing sequential optimization algorithms is not new. It dates back to (Tsitsiklis et al., 1986) for stochastic gradient methods; more recently Lee et al. (2014); Liu et al. (2014); Richtárik & Takáč (2015) study parallelization of BCD. The conditions under which these parallel BCD methods succeed, e.g., expected separable overapproximation (ESO), and coordinate Lipschitz conditions, bear a close resemblance to our conditions in Section 3.2, but are not the same due to differences in how solutions are updated and what subproblems arise. In particular, our conditions are *affine invariant*. We provide detailed comparisons to parallel coordinate descent in Appendix D.5.

Asynchronous algorithms. Asynchronous algorithms that allow delayed parameter updates have been proposed earlier for stochastic gradient descent (Niu et al., 2011) and parallel BCD (Liu et al., 2014). We propose the first asynchronous algorithm for Frank-Wolfe. Our asynchronous scheme not only permits delayed minibatch updates, but also allows the updates for coordinate blocks *within each* minibatch to have different delays. Therefore, each update may not be a solution of (2) for any single x . Moreover, we obtain strictly better dependence on the delay parameter than predecessors (e.g., an *exponential improvement* over Liu et al. (2014)) possibly due to a sharper analysis.

Other related work. While preparing our manuscript, we discovered the preprint (Bellet et al., 2014) which also studies distributed Frank-Wolfe. We note that (Bellet et al., 2014) focuses on Lasso type problems and communication costs, and hence, is not directly comparable to our results.

Notation. We briefly summarize our notation now. The vector $x \in \mathbb{R}^m$ denotes the parameter vector, possibly split into n coordinate blocks. For block $i = 1, \dots, n$, $E_i \in \mathbb{R}^{m \times m_i}$ is the projection matrix which projects $x \in \mathbb{R}^m$ down to $x_{(i)} \in \mathbb{R}^{m_i}$; thus $x_{(i)} = E_i x$. The adjoint operator E_i^* maps $\mathbb{R}^{m_i} \rightarrow \mathbb{R}^m$, thus $x_{[i]} = E_i^* x_{(i)}$ is x with zeros in all dimensions except $x_{(i)}$ (note the subscript $x_{[i]}$). We denote the size of a minibatch by τ , and the number of parallel workers (threads) by T . Unless otherwise stated, k denotes the iteration/epoch counter and γ denotes a stepsize. Finally, C_f^τ (and other such constants) denotes some curvature measure associated with function f and minibatch size τ . Such constants are important in our analysis, and will be described in greater detail in the main text.

2. Algorithm

In this section, we present an asynchronous parallel block-coordinate Frank-Wolfe algorithm (AP-BCFW) to solve (1). Our algorithm is designed to run fully asynchronously on either a shared-memory multicore architecture or on a distributed system.

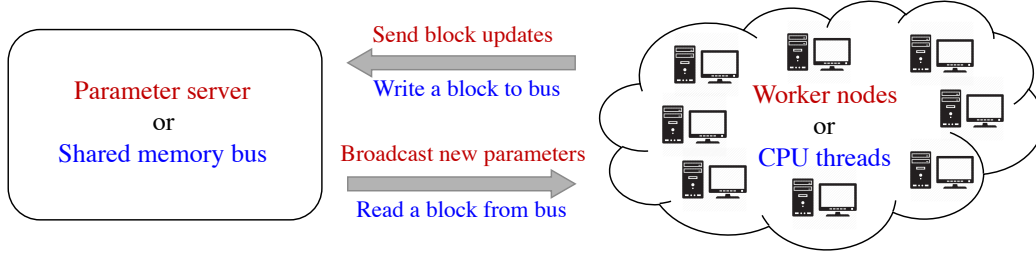


Figure 1. Illustration of the AP-BCFW in the distributed (in red) and share-memory settings (in blue). The “cloud” of all worker nodes (or CPU threads) is abstracted into an oracle that keeps feeding the server (or writing to the memory bus) with updates from solving possibly approximate (and/or delayed) solutions to (2) on iid uniform random blocks.

Algorithm 1 AP-BCFW: Asynchronous Parallel Block-Coordinate Frank-Wolfe (distributed)

Input: An initial feasible $x^{(0)}$, mini-batch size τ , a “Cloud” oracle \mathcal{O} satisfying Assumptions A1, A2.

0. Broadcast $x^{(0)}$ to all workers in \mathcal{O} .

for $k = 1, 2, \dots$ (k is the iteration number) **do**

1. Keep receiving $(i, s_{(i)})$ from \mathcal{O} until we have τ disjoint blocks (overwrite if collision¹). Denote the index set by S .

2. Update $x^{(k)} = x^{(k-1)} + \gamma_k \sum_{i \in S} (s_{[i]} - x_{[i]}^{(k-1)})$ with $\gamma_k = \frac{2n\tau}{\tau^2 k + 2n}$ or via line-search.

3. Broadcast $x^{(k)}$ (or just $x^{(k)} - x^{(k-1)}$) to \mathcal{O} .

4. Break if converged.

end for

Output: $x^{(k)}$.

For the shared-memory model, the computational work is divided amongst worker threads, each of which has access to a pool of coordinates that it may work on, as well as to the shared parameters. This setup matches the system assumptions in (Liu et al., 2014; Niu et al., 2011; Richtárik & Takáč, 2015), and most modern multicore machines permit such an arrangement.

On a distributed system, the parameter server (Dai et al., 2013; Li et al., 2013) broadcasts the most recent parameter vector periodically to each worker and workers keep sending updates to the parameter vector after solving the subroutines corresponding to a randomly chosen parameter. In either setting, we do not wait for slower workers or synchronize the parameters at any point of the algorithm, therefore many updates sent from the workers could be calculated based on a delayed parameter.

For convenience, we treat the pool of all workers as a single “cloud” oracle \mathcal{O} that keeps sending updates of form $\{i, s_{(i)}\}$ to the server, where i selects a block and $s_{(i)}$ is an approximate solution to (2) at the current parameter. Moreover, we assume that

A1. The sequence of i from \mathcal{O} is sampled i.i.d. uniformly from $\{1, 2, \dots, n\}$.

¹We bound the probability of collisions in Appendix D.2.

Assumption A1 is critical as it ensures Step 2 in the algorithm is an unbiased approximation of the batch FW. This assumption allows the workers to be arbitrarily heterogeneous as long as they each sample blocks i.i.d. uniformly and the time for each worker to produce $s_{(i)}$ does not depend on the block index i . Admittedly, this could be troublesome for some applications with heterogeneous blocks, we describe ways to enforce A1 in Appendix D.1.

An advantage of this oracle abstraction is its potential applicability well beyond the per-worker i.i.d. scheme. In practice, each worker might only have access to a small subset of $[n]$ and might be doing sequential sampling with periodic reshuffling. At the aggregate level, however, the oracle assumptions might still be reasonable approximations, especially if the number of workers T is large.

Both distributed and shared-memory settings can be captured under this oracle as is illustrated in Figure 1. Pseudocode of our scheme is given in Algorithm 1.

3. Analysis

The three key questions pertaining to Algorithm 1 are:

- *Does it converge?*
- *How fast? How much faster than BCFW ($\tau = 1$)?*
- *How do delayed updates affect the convergence?*

We answer the first two questions in Sections 3.1 and 3.2. Specifically, we show that AP-BCFW converges at a $O(1/k)$ rate. Our analysis reveals that the speedup of AP-BCFW over BCFW via parallelization is problem dependent. Intuitively, we show that speedups due to mini-batching ($\tau > 1$) depend on the average “coupling” of the objective function f across different coordinate blocks. For example, if f has a block symmetric diagonally dominant Hessian, then AP-BCFW converges $\tau/2$ times faster. We address the third question in Section 3.3, where we establish convergence results that depend only mildly on the “expected” delay κ . The bound is proportional to κ when we

allow the delay to grow unboundedly, and proportional to $\sqrt{\kappa}$ when the delay is bounded by a small κ_{\max} .

3.1. Main convergence results

We begin by defining a few quantities needed for our analysis. The first key quantity—also key to the analysis of several other FW methods—is the notion of **curvature**. Since AP-BCFW updates a subset of coordinate blocks at a time, we define *set curvature* for an index set $S \subseteq [n]$ as

$$C_f^{(S)} := \sup_{\substack{x \in \mathcal{M}, s_{(S)} \in \mathcal{M}^{(S)}, \\ \gamma \in [0, 1], \\ y = x + \gamma(s_{[S]} - x_{[S]})}} \frac{2}{\gamma^2} (f(y) - f(x) - \langle y_{(S)} - x_{(S)}, \nabla_{(S)} f(x) \rangle). \quad (3)$$

For index sets of size τ , we define the *expected set curvature* over a uniform choice of subsets as

$$C_f^\tau := \mathbb{E}_{S:|S|=\tau} [C_f^{(S)}] = \binom{n}{\tau}^{-1} \sum_{S \subset [n], |S|=\tau} C_f^{(S)}. \quad (4)$$

These curvature definitions are closely related to the global curvature constant C_f of Jaggi (2013) and the coordinate curvature $C_f^{(i)}$ and product curvature C_f^\otimes of Lacoste-Julien et al. (2013). Lemma 1 makes this relation more precise.

Lemma 1 (Curvature relations). *Suppose $S \subseteq [n]$ with cardinality $|S| = \tau$ and $i \in S$. Then,*

- i) $C_f^{(i)} \leq C_f^{(S)} \leq C_f$;
- ii) $\frac{1}{n} C_f^\otimes = C_f^1 \leq C_f^\tau \leq C_f^n = C_f$.

How the expected set curvature C_f^τ scales with τ is critical to bounding the speedup we can expect over BCFW; we provide a detailed analysis of this speedup in Section 3.2.

The next key object is an **approximate linear minimizer**. As in (Jaggi, 2013; Lacoste-Julien et al., 2013), we also allow the core computational subroutine that solves (2) to yield an approximate minimizer $s_{(i)}$. Formally, we assume:

A2. There is a constant $\delta \geq 0$, such that for every $k \geq 1$, the chosen minibatch $S \subset [n]$ of size τ and the corresponding blocks $s_{(S)} := (s_{(i)})_{i \in S}$ from \mathcal{O} obey

$$\mathbb{E} \left[\langle s_{(S)}, \nabla_{(S)} f^{(k)} \rangle - \min_{s' \in \mathcal{M}^{(S)}} \langle s', \nabla_{(S)} f^{(k)} \rangle \right] \leq \frac{\delta \gamma_k C_f^\tau}{2}. \quad (5)$$

where the expectation is taken over the random sequence of minibatch indices and corresponding updates from \mathcal{O} in the entire history up to step k .

Assumption A2 is strictly weaker than what is required in (Jaggi, 2013; Lacoste-Julien et al., 2013), as we only need the approximation to hold *in expectation*. With these definitions in hand, we are ready to state our convergence result.

Theorem 2 (Primal Convergence). *Say we use a “Cloud” oracle \mathcal{O} that generates a sequence of updates satisfying*

A1 and A2. Then, for each $k \geq 0$, the iterations in Algorithm 1 and its line search variant obey

$$\mathbb{E}[f(x^{(k)})] - f(x^*) \leq \frac{2nC}{\tau^2 k + 2n},$$

where $C = nC_f^\tau(1 + \delta) + f(x^{(0)}) - f(x^*)$.

At a first glance, the $n^2 C_f^\tau$ term in the numerator might seem bizarre, but as we will see in the next section, C_f^τ can be as small as $O(\frac{\tau}{n^2})$. This is the scale of the constant one should keep in mind to compare the rate to other methods, e.g., coordinate descent. Also note that so far this convergence result does not explicitly work for delayed updates, which we will analyze in Section 3.3 separately via the approximation parameter δ from (5).

For FW methods, one can also easily obtain a convergence guarantee in an appropriate primal-dual sense. To this end, we introduce our version of the **surrogate duality gap** (Jaggi, 2013); we define this as

$$g(x) := \max_{s \in \mathcal{M}} \langle x - s, \nabla f(x) \rangle \quad (6)$$

$$= \sum_{i=1}^n \max_{s_{(i)} \in \mathcal{M}^{(i)}} \langle x_{(i)} - s_{(i)}, \nabla_{(i)} f(x) \rangle = \sum_{i=1}^n g^{(i)}(x).$$

To see why (6) is actually a duality gap, note that since f is convex, the linearization $f(x) + \langle s - x, \nabla f(x) \rangle$ is always smaller than the function evaluated at any s , so that

$$g(x) \geq \langle x - x^*, \nabla f(x) \rangle \geq f(x) - f(x^*).$$

This duality gap is obtained for “free” in batch FW, but not in BCFW or AP-BCFW. Here, we only have an unbiased estimate $\frac{n}{|S|} \sum_{i \in S} g^{(i)}(x)$. For large τ , this estimate is close to $g(x)$ with high probability (McDiarmid’s Inequality), and can still be useful as a stopping criterion.

Theorem 3 (Primal-Dual Convergence). *Assume \mathcal{O} satisfies A1 and A2. Define the expected surrogate duality gap $g_k := \mathbb{E}g(x^{(k)})$ and weighted average $\bar{g}_k := \frac{2}{K(K+1)} \sum_{k=1}^K k g_k$ for the sequence of parameters $x^{(k)}$ in Algorithm 1. Then for every $K \geq 1$, there exists $k^* \in [K]$ such that*

$$g_{k^*} \leq \bar{g}_K \leq \frac{6nC}{\tau^2(K+1)},$$

with the same C in Theorem 2.

Relation with FW and BCFW: The above convergence guarantees can be thought of as an interpolation between BCFW and batch FW. If we take $\tau = 1$, they give exactly the convergence guarantee for BCFW (Lacoste-Julien et al., 2013, Theorem 2), and if we take $\tau = n$, we can drop $f(x^{(0)}) - f(x^*)$ from C (with a small modification in the analysis) and recover the classic batch guarantee as in Jaggi (2013).

Dependence on initialization: Unlike classic FW, the convergence rate of our method depends on the initialization. When $h_0 := f(x^{(0)}) - f(x^*) \geq nC_f^\tau$ and $\tau^2 < n$, the convergence is slower by a factor of $\frac{n}{\tau^2}$. The same concern was also raised in (Lacoste-Julien et al., 2013) with $\tau = 1$. We can actually remove the $f(x^{(0)}) - f(x^*)$ from C as long as we know that $h_0 \leq nC_f^\tau$. By Lemma 1, the expected set curvature C_f^τ increases with τ , so the fast convergence region becomes larger when we increase τ . In addition, if we pick $\tau^2 > n$, the rate of convergence is not affected by initialization anymore.

Speedup: The reader may have noticed the $n^2C_f^\tau$ term in the numerator. This is undesirable as n can be large (for instance, in structural SVM n is the total number of data points). The saving grace in BCFW is that when $\tau = 1$, C_f^τ is as small as $O(n^{-2})$ (see Lacoste-Julien et al., 2013, Lemmas A1 and A2), and it is easy to check that the dependence on n is the same even for $\tau > 1$. What really matters is how much speedup one can achieve over BCFW, and this relies critically on how C_f^τ depends on τ . Analyzing this dependence is our main focus in the next section.

3.2. Effect of parallelism / mini-batching

To understand when mini-batching is meaningful and to quantify its speedup, below we take a more careful look at the expected set curvature C_f^τ . In particular, we analyze and present a set of insightful conditions that govern its dependence on τ . The key idea is to quantify how strongly different coordinate blocks interact with each other.

To begin, assume that there exists a positive semidefinite matrix H such that for any $x, y \in \mathcal{M}$

$$f(y) \leq f(x) + \langle y - x, \nabla f(x) \rangle + \frac{1}{2}(y - x)^T H (y - x). \quad (7)$$

The matrix H may be viewed as a generalization of the gradient's Lipschitz constant (a scalar) to a matrix. For quadratic functions $f(x) = \frac{1}{2}x^T Qx + c^T x$, we can take $H = Q$. For twice differentiable functions, we can choose $H \in \{K \mid K \succeq \nabla^2 f(x), \forall x \in \mathcal{M}\}$.

Since $x = [x_1, \dots, x_n]$ (we write x_i instead of $x_{(i)}$ for brevity), we separate H into $n \times n$ blocks; so H_{ij} represents the block corresponding to x_i and x_j such that we can take the product $x_i^T H_{ij} x_j$. Now, we define a *bound-ness* parameter B_i for every i , and an *incoherence condition* with parameter μ_{ij} for every block coordinate pair $\mathcal{M}_i, \mathcal{M}_j$ such that

$$B_i = \sup_{x_i \in \mathcal{M}_i} x_i^T H_{ii} x_i, \quad \mu_{ij} = \sup_{x_i \in \mathcal{M}_i, x_j \in \mathcal{M}_j} x_i^T H_{ij} x_j,$$

$$B = \mathbb{E}_{i \sim \text{Unif}(\{[n]\})} B_i, \quad \mu = \mathbb{E}_{(i,j) \sim \text{Unif}(\{(i,j) \in [n]^2, i \neq j\})} \mu_{ij}.$$

Using these quantities, we obtain the following bound on the expected set-curvature.

Theorem 4. $C_f^\tau \leq 4(\tau B + \tau(\tau - 1)\mu)$ for any $\tau \in [n]$.

It is clear that when the incoherence term μ is large, the expected set curvature C_f^τ is proportional to τ^2 , and when μ is close to 0, then C_f^τ is proportional to τ . In other words, when the interaction between coordinates block is small, one gains from parallelizing BCFW. This is analogous to the situation in parallel coordinate descent (Liu et al., 2014; Richtárik & Takáč, 2015) and we will compare the rate of convergence explicitly with them in Appendix D.5.

Corollary 5. Consider a matrix M with B_i on the diagonal and μ_{ij} on the off-diagonal. If M is symmetric diagonally dominant (SDD), i.e., the sum of absolute off-diagonal entries in each row is no greater than the diagonal entry, then C_f^τ is proportional to τ .

The above result depends on the parameters B and μ . In Appendix D.4, we provide two concrete examples (multi-class classification with structural SVM and graph fused lasso) where we can express B and μ as problem-dependent quantities and provide explicit upper bounds of C_f^τ . In both examples, we show that choosing larger τ yields faster convergence (at least up to some point).

3.3. Convergence with delayed updates

Often due to the delays in communication, some updates pushed back by workers are calculated based on delayed parameters that were broadcast earlier. Dropping these updates or enforcing synchronization will create a huge system overhead especially when the size of the minibatch is small. Ideally, we want to just accept the delayed updates as if they were correct, and broadcast new parameters to workers without locking the updates. The question is, does this idea work?

In this section, we model delays from updates to be i.i.d. from an unknown distribution that can depend on k , but not on blocks. Under these assumptions, we show that the effect of delayed updates can be treated as an approximate oracle that satisfies A2 in (5) with some specific constant δ that depends on the expected delay κ and the maximum delay parameter κ_{\max} (when it exists). This allows us to invoke results in Section 3.1 to establish convergence for delayed updates. The results also depend on the following diameter and gradient Lipschitz constant for a norm $\|\cdot\|$

$$D_{\|\cdot\|}^{(S)} = \sup_{x, y \in \mathcal{M}^{(S)}} \|x - y\|,$$

$$L_{\|\cdot\|}^{(S)} = \sup_{\substack{x, y \in \mathcal{M}, y = x + s \\ \|s\| \leq \gamma, \\ s \in \text{span}(\mathcal{M}^{(S)})}} \frac{1}{\gamma^2} (f(y) - f(x) - \langle y - x, \nabla f(x) \rangle),$$

$$D_{\|\cdot\|}^\tau = \max_{S \subset [n] \mid |S|=m} D_{\|\cdot\|}^{(S)}, \text{ and } L_{\|\cdot\|}^\tau = \max_{S \subset [n] \mid |S|=m} L_{\|\cdot\|}^{(S)}.$$

Theorem 6 (Delayed Updates as Approximate Oracle). *For each norm $\|\cdot\|$ of choice, let $D_{\|\cdot\|}^\tau$ and $L_{\|\cdot\|}^\tau$ be defined above. Let the a random variable of delay be \varkappa and let $\kappa := \mathbb{E}\varkappa$ be the expected delay from any worker. Moreover, assume that the algorithm drops any updates with delay greater than $k/2$ at iteration k . Then for the version of the algorithm without line-search, the delayed oracle will produce $s \in \mathcal{M}^{(S)}$ such that (5) holds with*

$$\delta = 4\kappa\tau L_{\|\cdot\|}^1 D_{\|\cdot\|}^1 D_{\|\cdot\|}^\tau / (C_f^\tau). \quad (8)$$

Furthermore, if we assume that there is a κ_{\max} such that $\mathbb{P}(\varkappa \leq \kappa_{\max}) = 1$ for all k , then (5) holds with $\delta =$

$$c_{n,\tau\kappa_{\max}} \frac{4\tau L_{\|\cdot\|}^1 D_{\|\cdot\|}^1 \mathbb{E}D_{\|\cdot\|}^{\varkappa\tau}}{C_f^\tau} \text{ where}$$

$$c_{n,\tau\kappa_{\max}} = \begin{cases} \frac{3 \log n}{\log(n/(\tau\kappa_{\max}))} & \text{if } \kappa_{\max}\tau < n/\log n, \\ O(\log n) & \text{if } \kappa_{\max}\tau = O(n \log n), \\ \frac{(1+o(1))\tau\kappa_{\max}}{n} & \text{if } \kappa_{\max}\tau \gg n \log n. \end{cases} \quad (9)$$

The results above imply that AP-BCFW (without line-search) converges in both primal optimality and in duality gap according to Theorems 2 and 3 with the same $O(1/k)$ rate. Comparing to versions that solve (2) exactly, the delayed version has an additional additive factor in the numerator of form

$$4n\kappa\tau L_{\|\cdot\|}^1 D_{\|\cdot\|}^1 D_{\|\cdot\|}^\tau \quad \text{or} \quad O\left(\tau L_{\|\cdot\|}^1 D_{\|\cdot\|}^1 \mathbb{E}D_{\|\cdot\|}^{\varkappa\tau} \log n\right)$$

with the additional assumption that $\kappa_{\max} = O(n \log n/\tau)$.

Note that (8) depends on the expected delay rather than the maximum delay, and as $k \rightarrow \infty$ we allow the maximum delay to grow unboundedly. This allows the system to automatically deal with heavy-tailed delay distributions and sporadic stragglers. When we do have a small bounded delay, we produce stronger bounds (9) with a multiplier that is either a constant (when $\tau\kappa_{\max} = O(n^{1-\epsilon})$ for any $\epsilon > 0$), proportional to $\log n$ (when $\tau\kappa \leq n$) or proportional to $\frac{\tau\kappa_{\max}}{n}$ (when $\tau\kappa$ is large). The whole expression often has sublinear dependence on the expected delay κ . For instance, we prove in the appendix the following:

Lemma 7. *When $\|\cdot\|$ is Euclidean norm*

$$\mathbb{E}D_{\|\cdot\|}^{\varkappa\tau} \leq D_{\|\cdot\|}^{\lceil \mathbb{E}\varkappa\tau \rceil} \leq \sqrt{\lceil \mathbb{E}\varkappa \rceil} D_{\|\cdot\|}^\tau.$$

The bound is proportional to $\sqrt{\kappa}$ when $\kappa = \Omega(1)$. This is strictly better than Niu et al. (2011) which has quadratic dependence on κ_{\max} and Liu et al. (2014) which has exponential dependence on κ_{\max} . Our mild κ_{\max} dependence for the cases $\tau\kappa_{\max} > n$ suggests that the (9) remains proportional to $\sqrt{\kappa}$ even when we allow the maximum delay parameter to be as large as n/τ or larger without significantly affecting the convergence. Note that this allows some workers to be delayed for several data passes.

Observe that when $\tau = 1$, where the results reduces to a lock-free variant for BCFW, δ becomes proportional to $L_{\|\cdot\|}^1 [D_{\|\cdot\|}^1]^2 / C_f^1$. This is always greater than 1 (see e.g., Jaggi, 2013, Appendix D) but due to the flexibility of choosing the norm, this quantity corresponding to the most favorable norm is typically a small constant. For example, when f is a quadratic function, we show that $C_f^1 = L_{\|\cdot\|}^1 [D_{\|\cdot\|}^1]^2$ (see Appendix D.3). When $\tau > 1$, $\tau L_{\|\cdot\|}^1 D_{\|\cdot\|}^1 D_{\|\cdot\|}^\tau / C_f^\tau$ is often $O(\sqrt{\tau})$ for an appropriately chosen norm. Therefore, (8) and (9) are roughly in the order of $O(\kappa\sqrt{\tau})$ and $O(\sqrt{\kappa\tau})$ respectively².

Lastly, we remark that κ and τ are not independent. When we increase τ , we update the parameters less frequently and κ gets smaller. In a real distributed system, with constant throughput in terms of number of oracle solves per second from all workers. If the average delay is a fixed number in clock time specified by communication time. Then $\tau\kappa$ is roughly a constant regardless how τ is chosen.

4. Experiments


In this section, we experimentally demonstrate performance gains from the three key features of our algorithm: minibatches of data, parallel workers, and asynchrony.

4.1. Minibatches of Data

We conduct simulations to study the effect of mini-batch size τ , where larger τ implies greater degrees of parallelism as each worker can solve one or more subproblems in a mini-batch. In our simulation, we re-use the structural SVM setup from Lacoste-Julien et al. (2013) for a sequence labeling task on a subset of the OCR dataset (Taskar et al., 2004) ($n = 6251, d = 4082$). The dual problem has block-separable probability simplex constraint therefore allowing us to run AP-BCFW, and each subproblem can be solved efficiently using the Viterbi algorithm (more details are included in Appendix C). The speedup on this dataset is shown in Figure 2(a). For this dataset, we use $\lambda = 1$ with weighted averaging and line-search throughout (no delay is allowed). We measure the speedup for a particular $\tau > 1$ in terms of the number of iterations (Algorithm 1) required to converge relative to $\tau = 1$, which corresponds to BCFW. Figure 2(a) shows that AP-BCFW achieves linear speedup for mini-batch size up to $\tau \approx 50$. Further speedup is sensitive to the convergence criteria.

In our simulation for Group Fused Lasso, we generate a piecewise constant dataset of size ($n = 100, d = 10$, in Eq. 2) with Gaussian noise. We use $\lambda = 0.01$ and a primal suboptimality threshold as our convergence criterion. At each iteration, we solve τ subproblems (i.e. the mini-batch size). Figure 2(b) shows the speed-up over $\tau = 1$ (BCFW).

²For details, see our discussion in Appendix D.3



(a) Structural SVM ($n=6251$) (b) Group Fused Lasso ($n=100$)

Figure 2. Performance improvement with τ for (a) Structural SVM on the OCR dataset (Lacoste-Julien et al., 2013; Taskar et al., 2004) and (b) Group Fused Lasso on a synthetic dataset. f^* denotes primal optimum (the “primal” problem is actually referring to the dual problem in both cases). The performance metric here is the number of iterations to achieve ϵ -suboptimality.

Similar to the structural SVM, the speedup is almost perfect for small τ ($\tau \leq 55$) but tapers off for large τ to varying degrees depending on the convergence thresholds.

4.2. Shared Memory Parallel Workers

We implement AP-BCFW for the structural SVM in a multicore shared-memory system using the full OCR dataset ($n = 6877$). All shared-memory experiments were implemented in C++ and conducted on a 16-core machine with Intel(R) Xeon(R) CPU E5-2450 2.10GHz processors and 128G RAM. We first fix the number of workers at $T = 8$ and vary the mini-batch size τ . Figure 3(a) shows the absolute convergence (i.e. the convergence per second). We note that AP-BCFW outperforms single-threaded BCFW under all investigated τ , showing the efficacy of parallelization. Within AP-BCFW, convergence improves with increasing mini-batch sizes up to $\tau = 3T$, but worsens when $\tau = 5T$ as the error from the large mini-batch size dominates additional computation. The optimal τ for a given number of workers (T) depends on both the dataset (how “coupled” are the coordinates) and also system implementations (how costly is the synchronization).

Since speedup for a given T depends on τ , we search for the optimal τ across multiples of T to find the best speedup for each T . Figure 3(b) shows faster convergence of AP-BCFW over BCFW ($T = 1$) when $T > 1$ workers are available. It is important to note that the x-axis is *wall-clock time* rather than the number of epochs.

Figure 3(c) shows the speedup with varying T . AP-BCFW achieves near-linear speed up for smaller T . The speedup curve tapers off for larger T for two reasons: (1) Large T incurs higher system overheads, and thus needs larger τ to utilize CPU efficiently; (2) Larger τ incurs errors as shown in Fig. 2(a). If the subproblems were more time-consuming to solve, the affect of system overhead would be reduced. We simulate harder subproblems by simply

solving them $m \sim \text{Uniform}(5, 15)$ times instead of just once. The speedup is nearly perfect as shown in Figure 3(d). Again, we observe that a more generous convergence threshold produces higher speedup, suggesting that resource scheduling could be useful (e.g., allocate more CPUs initially and fewer as algorithm converges).

We repeated the experiment on a larger synthetic dataset with $n = 103155$, $d = 4082$ created from the above mentioned OCR data as follows: for each of the 6877 words, generate 15 words with noisy images for characters, where the noise is introduced by flipping the bits of the images with probability 0.05 independently. The speedup with parallelization, shown in Figure 4, essentially follows the same pattern as it did in Figures 3c, 3b for original data.

4.3. Performance gain with asynchronous updates

We compare AP-BCFW³ with a synchronous version of the algorithm (SP-BCFW) where the server assigns τ/T subproblems to each worker, then waits for and accumulates the solutions before proceeding to the next iteration. We simulate workers of varying slow-downs in our shared-memory setup by assigning a *return probability* $p_i \in (0, 1]$ to each worker w_i . After solving each subproblem, worker w_i reports the solution to the server with probability p_i . Thus, a worker with $p_i = 0.8$ will drop 20% of the updates on average corresponding to 20% slow-down.

We use $T = 14$ workers for the experiments in this section. We first simulate the scenario with just one straggler with return probability $p \in (0, 1]$ while the other workers run at full speed ($p = 1$). Figure 5(a) shows that the average time per effective datapass (over 20 passes and 5 runs) of AP-BCFW stays almost unchanged with slowdown factor $1/p$ of the straggler, whereas it increases linearly for SP-BCFW. This is because AP-BCFW relies on the average available worker processing power, while SP-BCFW is only as fast as the slowest worker.

Next, we simulate a heterogeneous environment where the workers have varying speeds. While varying a parameter $\theta \in [0, 1]$, we set $p_i = \theta + i/T$ for $i = 1, \dots, T$. Figure 5(b) shows that AP-BCFW slows down only by a factor of 1.4 compared to the no-straggler case. Assuming that the server and worker each take about half the (wall-clock) time on average per epoch, we would expect the run time to increase by 50% if the average worker speed halves, which is the case if $\theta = 0$ (i.e., $\frac{1}{\theta} \rightarrow \infty$). Thus, a factor of 1.4 is reasonable. The performance of SP-BCFW is almost identical to that in the previous experiment as its speed is determined by the slowest worker. Our experiments show that AP-BCFW is robust to stragglers and system heterogeneity.

³The version that has no delayed updates, but allows workers to asynchronously solve subproblems within each mini-batch.

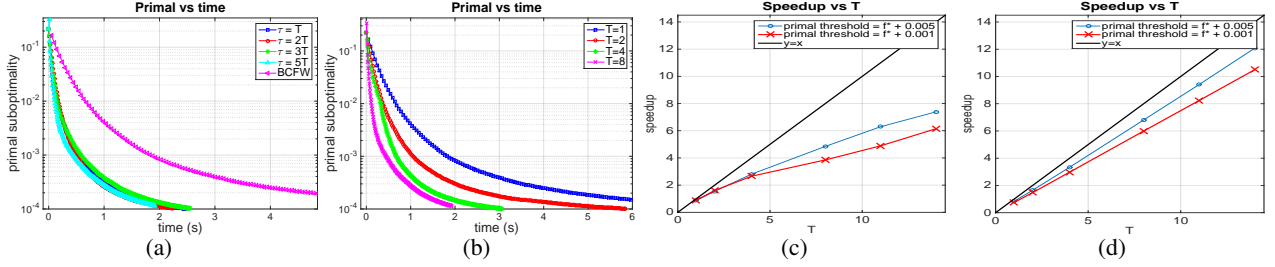


Figure 3. From left: (a) Primal suboptimality vs wall-clock time using 8 workers ($T = 8$) and various mini-batch sizes τ . (b) Primal suboptimality vs wall-clock time for varying T with best τ chosen for each T separately. (c) Speedup via parallelization with the best τ chosen among multiples of T ($T, 2T, \dots$) for each T . (d) The same with longer subproblems.

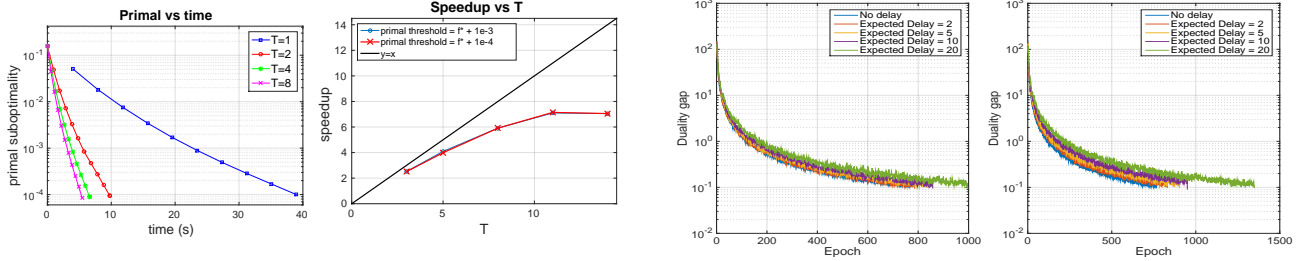


Figure 4. Speedup with parallelization on a synthetic OCR dataset. Left plot shows the decay of primal suboptimality and the right one shows the speedup.




Figure 5. Average time per data pass in asynchronous and synchronous modes for two cases: one worker is slow with return probability p (left); workers have return probabilities (p_i s) uniformly in $[\theta, 1]$ (right). Times normalized separately for AP-BCFW, SP-BCFW w.r.t. to where workers run at full speed.

4.4. Convergence under unbounded heavy-tailed delay

In this section, we illustrate the mild effect of delay on convergence by randomly drawing an independent delay variable for each worker. For simplicity, we use $\tau = 1$ (BCFW) on the group fused lasso problem from Section 4.1. We sample \varkappa using either a Poisson distribution or a heavy-tailed Pareto distribution (round to the nearest integer). The Pareto distribution is chosen with shape parameter $\alpha = 2$ and scale parameter $x_m = \kappa/2$ such that $\mathbb{E}\varkappa = \kappa$ and $\text{Var}\varkappa = \infty$. During the experiment, at iteration k , any updates that were based on a delay greater than $k/2$ are dropped (as our theory stipulates). The results are shown in Figure 6. Observe that for both cases, the impact of delay is rather mild. With expected delays up to 20, the algorithm only takes fewer than twice as many iterations to converge.




Figure 6. Illustrations of the convergence BCFW with delayed updates. On the left, we have the delay sampled from a Poisson distribution. The figure on the right is for delay sampled from a Pareto distribution. We run each problem until the duality gap reaches 0.1.

5. Conclusion

In this paper, we propose an asynchronous parallel generalization of the block-coordinate Frank-Wolfe method (Lacoste-Julien et al., 2013), analyze its convergence and provide intuitive conditions under which it has a provable speed-up over BCFW. We also show that the method is resilient to delayed updates in the distributed setting. The convergence bound depends only linearly on the expected delay and possibly sublinearly if the delay is bounded, yielding an exponential improvement over the dependence on the same parameter in parallel coordinate descent (Liu et al., 2014). The asynchronous updates allow our method to be robust to stragglers and node failure as the speed of AP-BCFW depends on *average* worker speed instead of the slowest. We demonstrate the effectiveness of the algorithm in structural SVM and Group Fused Lasso with both controlled simulation and real-data experiments on a multi-core workstation. For the structural SVM, it leads to a speed-up over the state-of-the-art BCFW by an order of magnitude using 16 parallel processors. As a projection-free FW method, we expect our algorithm to be very competitive in large-scale constrained optimization problems, especially when projections are expensive. Future work includes analysis for the strongly convex setting, the non-convex setting and ultimately releasing a general purpose software package for practitioners to deploy in Big Data applications.

Acknowledgements

We thank the AC and anonymous reviewers for helpful suggestions that led to significant improvement of the paper. YW was supported by NSF Award BCS-0941518 to CMU Statistics and a grant by Singapore NSF under its International Research Centre @ Singapore Funding Initiative administered by the IDM Programme Office. VS was partially supported by NSF grant DMS-1309174. WD was supported by DARPA XDATA FA87501220324. SS acknowledges partial support from NSF-IIS-1409802.

References

- Ahipasaoglu, Damla S, Sun, Peng, and Todd, Michael J. Linear convergence of a modified frank-wolfe algorithm for computing minimum-volume enclosing ellipsoids. *Optimisation Methods and Software*, 23(1):5–19, 2008.
- Alaíz, Carlos M, Barbero, Álvaro, and Dorronsoro, José R. Group fused lasso. In *Artificial Neural Networks and Machine Learning–ICANN 2013*, pp. 66–73. Springer, 2013.
- Beck, Amir and Tetrushvili, Luba. On the convergence of block coordinate descent type methods. *SIAM Journal on Optimization*, 23(4):2037–2060, 2013.
- Bellet, Aurélien, Liang, Yingyu, Garakani, Alireza Bagheri, Balcan, Maria-Florina, and Sha, Fei. Distributed Frank-Wolfe algorithm: A unified framework for communication-efficient sparse learning. *CoRR*, abs/1404.2644, 2014.
- Bleakley, Kevin and Vert, Jean-Philippe. The group fused Lasso for multiple change-point detection. arXiv preprint arXiv:1106.4199, 2011.
- Bredies, Kristian, Lorenz, Dirk A, and Maass, Peter. A generalized conditional gradient method and its connection to an iterative shrinkage method. *Computational Optimization and Applications*, 42(2):173–193, 2009.
- Clarkson, Kenneth L. Coresets, sparse greedy approximation, and the Frank-Wolfe algorithm. *ACM Transactions on Algorithms (TALG)*, 6(4):63, 2010.
- Collins, Michael, Globerson, Amir, Koo, Terry, Carreras, Xavier, and Bartlett, Peter L. Exponentiated gradient algorithms for conditional random fields and max-margin markov networks. *JMLR*, 9:1775–1822, 2008.
- Dai, Wei, Wei, Jinliang, Zheng, Xun, Kim, Jin Kyu, Lee, Seunghak, Yin, Junming, Ho, Qirong, and Xing, Eric P. Petuum: a framework for iterative-convergent distributed ML. *arXiv:1312.7651*, 2013.
- Fercoq, Olivier and Richtárik, Peter. Accelerated, parallel, and proximal coordinate descent. *SIAM Journal on Optimization*, 25(4):1997–2023, 2015.
- Foygel, Rina, Horrell, Michael, Drton, Mathias, and Lafferty, John D. Nonparametric reduced rank regression. In *NIPS’12*, pp. 1628–1636, 2012.
- Frank, Marguerite and Wolfe, Philip. An algorithm for quadratic programming. *Naval Research Logistics Quarterly*, 3(1-2):95–110, 1956.
- Freund, Robert M. and Grigas, Paul. New analysis and results for the frank-wolfe method. *Mathematical Programming*, 155(1):199–230, 2014. ISSN 1436-4646.
- Fujishige, Satoru and Isotani, Shiguo. A submodular function minimization algorithm based on the minimum-norm base. *Pacific Journal of Optimization*, 7(1):3–17, 2011.
- Garber, Dan and Hazan, Elad. A linearly convergent conditional gradient algorithm with applications to online and stochastic optimization. *arXiv:1301.4666*, 2013.
- Harchaoui, Zaid, Juditsky, Anatoli, and Nemirovski, Arkadi. Conditional gradient algorithms for norm-regularized smooth convex optimization. *Mathematical Programming*, 152(1-2):75–112, 2015. ISSN 0025-5610.
- Hazan, Elad and Kale, Satyen. Projection-free online learning. In *ICML’12*, 2012.
- Jaggi, Martin. *Sparse convex optimization methods for machine learning*. PhD thesis, Diss., Eidgenössische Technische Hochschule ETH Zürich, Nr. 20013, 2011, 2011.
- Jaggi, Martin. Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In *ICML’13*, pp. 427–435, 2013.
- Jegelka, Stefanie, Bach, Francis, and Sra, Suvrit. Reflection methods for user-friendly submodular optimization. In *NIPS’13*, pp. 1313–1321, 2013.
- Lacoste-Julien, Simon and Jaggi, Martin. On the global linear convergence of Frank-Wolfe optimization variants. In *NIPS’15*, pp. 496–504, 2015.
- Lacoste-Julien, Simon, Jaggi, Martin, Schmidt, Mark, and Pletscher, Patrick. Block-coordinate Frank-Wolfe optimization for structural svms. In *ICML’13*, pp. 53–61, 2013.
- Lafond, Jean, Wai, Hoi-To, and Moulines, Eric. Convergence analysis of a stochastic projection-free algorithm. *arXiv:1510.01171*, 2015.

- LeBlanc, Larry J, Morlok, Edward K, and Pierskalla, William P. An efficient approach to solving the road network equilibrium traffic assignment problem. *Transportation Research*, 9(5):309–318, 1975.
- Lee, Seunghak, Kim, Jin Kyu, Zheng, Xun, Ho, Qirong, Gibson, Garth A, and Xing, Eric P. On model parallelization and scheduling strategies for distributed machine learning. In *NIPS'14*, pp. 2834–2842, 2014.
- Li, Mu, Zhou, Li, Yang, Zichao, Li, Aaron, Xia, Fei, Andersen, David G, and Smola, Alexander. Parameter server for distributed machine learning. In *NIPS Workshop: Big Learning*, 2013.
- Liu, Ji, Musialski, Przemyslaw, Wonka, Peter, and Ye, Jieping. Tensor completion for estimating missing values in visual data. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 35(1):208–220, 2013.
- Liu, Ji, Wright, Stephen J, Ré, Christopher, and Bittorf, Victor. An asynchronous parallel stochastic coordinate descent algorithm. *JMLR*, 2014.
- Mitzenmacher, Michael. The power of two choices in randomized load balancing. *Parallel and Distributed Systems, IEEE Transactions on*, 12(10):1094–1104, 2001.
- Nesterov, Yu. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization*, 22(2):341–362, 2012.
- Niu, Feng, Recht, Benjamin, Ré, Christopher, and Wright, Stephen J. Hogwild!: A lock-free approach to parallelizing stochastic gradient descent. In *NIPS'11*, pp. 693–701, 2011.
- Ouyang, Hua and Gray, Alexander G. Fast stochastic Frank-Wolfe algorithms for nonlinear SVMs. In *SDM*, 2010.
- Qu, Zheng and Richtárik, Peter. Coordinate descent with arbitrary sampling ii: Expected separable overapproximation. *arXiv preprint arXiv:1412.8063*, 2014.
- Raab, Martin and Steger, Angelika. Balls into bins - a simple and tight analysis. In *Randomization and Approximation Techniques in Computer Science*, pp. 159–170. Springer, 1998.
- Richtárik, Peter and Takáč, Martin. Parallel coordinate descent methods for big data optimization. *Mathematical Programming*, pp. 1–52, 2015.
- Taskar, Ben, Guestrin, Carlos, and Koller, Daphne. Max-margin Markov networks. In *NIPS'04*, pp. 25–32. MIT Press, 2004.
- Tsitsiklis, John N, Bertsekas, Dimitri P, Athans, Michael, et al. Distributed asynchronous deterministic and stochastic gradient optimization algorithms. *IEEE Transactions on Automatic Control*, 31(9), 1986.
- Yu, Chun-Nam John and Joachims, Thorsten. Learning structural svms with latent variables. In *ICML'09*, pp. 1169–1176. ACM, 2009.
- Zhang, Xinhua, Yu, Yaoliang, and Schuurmans, Dale. Accelerated training for matrix-norm regularization: A boosting approach. In *NIPS'12*, pp. 2915–2923, 2012.
- Zhang, Xinhua, Yu, Yao-Liang, and Schuurmans, Dale. Polar operators for structured sparse estimation. In *NIPS'13*, pp. 82–90, 2013.

This is the supplementary document to the paper: “Parallel and Distributed Block-Coordinate Frank-Wolfe Algorithms”.

A. Convergence analysis

We provide a self-contained convergence proof in this section. The skeleton of our convergence proof follow closely from [Lacoste-Julien et al. \(2013\)](#) and [Jaggi \(2013\)](#). There are a few subtle modification and improvements that we need to add due to our weaker definition of approximate oracle call that is nearly correct only in expectation. The delayed convergence is new and interesting for the best of our knowledge, which uses a simple result in “load balancing” ([Mitzenmacher, 2001](#)).

Note that for the cleanness of the presentation, we focus on the primal and primal-dual convergence of the version of the algorithms with pre-defined step sizes and additive approximate subroutine, it is simple to extend the same analysis for line-search variant and multiplicative approximation.

A.1. Primal Convergence

Lemma 8. *Denote the gap between current $f(x^{(k)})$ and the optimal $f(x^*)$ to be $h(x^{(k)})$. The iterative updates in Algorithm 1 (with arbitrary fixed stepsize γ or by the line search) obey*

$$\mathbb{E}h(x^{(k+1)}) \leq (1 - \frac{\gamma\tau}{n})\mathbb{E}h(x^{(k)}) + \frac{\gamma^2(1+\delta)}{2}C_f^\tau.$$

where the expectation is taken over the joint randomness all the way to iteration $k + 1$.

Proof. Let $x := x^{(k)}$ for notational convenience. We prove the result for Algorithm 1 first. Apply the definition of $C_f^{(S)}$ and then apply the definition of the additive approximation in (5), to get

$$\begin{aligned} f(x_{\text{line-search}}^{(k+1)}) &\leq f(x_\gamma^{(k+1)}) = f(x + \gamma \sum_{i \in S} (s_{[i]} - x_{[i]})) \\ &\leq f(x) + \gamma \sum_{i \in S} \langle s_{[i]} - x_{[i]}, \nabla_{[i]} f(x) \rangle + \frac{\gamma^2}{2} C_f^{(S)} \\ &= f(x) + \gamma \langle s_{[S]} - x_{[S]}, \nabla_{[S]} f(x) \rangle + \frac{\gamma^2}{2} C_f^{(S)} \end{aligned}$$

Subtract $f(x^*)$ on both sides we get:

$$h(x^{(k+1)}) \leq h(x^{(k)}) + \gamma \langle s_{[S]} - x_{[S]}^{(k)}, \nabla_{[S]} f(x^{(k)}) \rangle + \frac{\gamma^2}{2} C_f^{(S)}$$

Now take the expectation over the entire history then apply (5) and definition of the surrogate duality gap (6), we obtain

$$\begin{aligned} \mathbb{E}h(x^{(k+1)}) &\leq \mathbb{E}h(x^{(k)}) + \mathbb{E} \left\{ \gamma \langle s_{[S]} - x_{[S]}^{(k)}, \nabla_{[S]} f(x^{(k)}) \rangle \right\} + \mathbb{E} \frac{\gamma^2}{2} C_f^{(S)} \\ &= \mathbb{E}h(x^{(k)}) + \gamma \mathbb{E} \left\{ \langle s_{[S]}, \nabla_{[S]} f(x^{(k)}) \rangle - \min_{s \in \mathcal{M}^{(S)}} \langle s, \nabla_{[S]} f(x^{(k)}) \rangle \right\} \\ &\quad - \gamma \mathbb{E} \left\{ \langle x_{[S]}^{(k)}, \nabla_{[S]} f(x^{(k)}) \rangle - \min_{s \in \mathcal{M}^{(S)}} \langle s, \nabla_{[S]} f(x^{(k)}) \rangle \right\} + \frac{\gamma^2}{2} C_f^\tau \\ &\leq \mathbb{E}h(x^{(k)}) + \frac{\gamma^2 \delta}{2} C_f^\tau - \gamma \mathbb{E}_{x^k} \mathbb{E}_{S|x^k} \sum_{i \in S} g^{(i)}(x^{(k)}) + \frac{\gamma^2}{2} C_f^\tau \\ &= \mathbb{E}h(x^{(k)}) + \frac{\gamma^2 \delta}{2} C_f^\tau - \gamma \mathbb{E}_{x^k} \frac{\tau}{n} g(x^{(k)}) + \frac{\gamma^2}{2} C_f^\tau \tag{10} \\ &\leq (1 - \frac{\gamma\tau}{n})\mathbb{E}h(x^{(k)}) + \frac{\gamma^2(1+\delta)}{2}C_f^\tau. \end{aligned}$$

The last inequality follows from the property of the surrogate duality gap $g(x^{(k)}) \geq h(x^{(k)})$ due to the fact that $g(x) = f(x) - f^*(\cdot)$. This completes the proof of the descent lemma. \square

Now we are ready to state the proof for Theorem 2.

Proof of Theorem 2. We follow the proof in Theorem C.1 in (Lacoste-Julien et al., 2013) to prove the statement for Algorithm 1. The difference is that we use a different and carefully chosen sequence of step size.

Take $C = h_0 + n(1 + \delta)C_f^\tau$, and denote $\mathbb{E}h(x^{(k)})$ as h_k for short hands. The inequality in Lemma 8 simplifies to

$$h_{k+1} \leq \left(1 - \frac{\gamma\tau}{n}\right) h_k + \frac{\gamma^2}{2n} C.$$

Now we will prove $h_k \leq \frac{2nC}{\tau^2 k + 2n}$ for $\gamma_k = \frac{2n\tau}{\tau^2 k + 2n}$ by induction. The base case $k = 0$ is trivially true since $C > h_0$. Assuming that the claim holds for k , we apply the induction hypothesis and the above inequality is reduced to

$$\begin{aligned} h_{k+1} &\leq \left(1 - \frac{\gamma\tau}{n}\right) h_k + \frac{\gamma^2}{2n} C \leq \frac{2nC}{\tau^2 k + 2n} \left[1 - \frac{\gamma\tau}{n} + \frac{\tau^2 k + 2n}{2n} \frac{\gamma^2}{2n}\right] \\ &= \frac{2nC}{\tau^2 k + 2n} \left[\frac{\tau^2 k + 2n}{\tau^2 k + 2n} - \frac{2n\tau}{\tau^2 k + 2n} \cdot \frac{\tau}{n} + \frac{(2n\tau)^2}{4n^2(\tau^2 k + 2n)}\right] \\ &= \frac{2nC}{\tau^2 k + 2n} \cdot \frac{\tau^2 k + 2n - \tau^2}{\tau^2 k + 2n} \leq \frac{2nC}{\tau^2 k + 2n} \cdot \frac{\tau^2 k + 2n - \tau^2 + \tau^2}{\tau^2 k + 2n + \tau^2} \\ &= \frac{2nC}{\tau^2(k+1) + 2n}. \end{aligned}$$

This completes the induction and hence the proof for the primal convergence for Algorithm 1. \square

A.2. Convergence of the surrogate duality gap

Proof of Theorem 3. We mimic the proof in (Lacoste-Julien et al., 2013, Section C.3) for the analogous result closely, and we will use the same notation for h_k and C as in the proof for primal convergence, moreover denote $g_k = \mathbb{E}g(x^{(k)})$. First from (10) in the proof of Lemma 8, we have

$$h_{k+1} \leq h_k - \frac{\gamma\tau}{n} g_k + \frac{\gamma^2}{2n} C.$$

Rearrange the terms, we get

$$g_k \leq \frac{n}{\gamma\tau} (h_k - h_{k+1}) + \frac{\gamma C}{2\tau}. \quad (11)$$

The idea is that if we take an arbitrary convex combination of $\{g_1, \dots, g_K\}$, the result will be within the convex hull, namely between the minimum and the maximum, hence proven the existence claim in the theorem. By choosing weight $\rho_k := k/S_K$ where normalization constant $S_K = \frac{K(K+1)}{2}$ and taking the convex combination of both side of (11), we have

$$\begin{aligned} \mathbb{E}(\min_{k \in [K]} g_k) &\leq \sum_{k=0}^K \rho_k g_k \leq \frac{n}{\tau} \sum_{k=1}^K \rho_k \left(\frac{h_k}{\gamma_k} - \frac{h_{k+1}}{\gamma_k}\right) + \sum_{k=0}^K \rho_k \gamma_k \frac{C}{2\tau} \\ &= \frac{n}{\tau} \left(\frac{h_0 \rho_0}{\gamma_0} - h_{K+1} \frac{\rho_K}{\gamma_K}\right) + \frac{n}{\tau} \sum_{k=0}^{K-1} h_{k+1} \left(\frac{\rho_{k+1}}{\gamma_{k+1}} - \frac{\rho_k}{\gamma_k}\right) + \sum_{k=0}^K \rho_k \gamma_k \frac{C}{2\tau} \\ &\leq \frac{n}{\tau} \sum_{k=0}^{K-1} h_{k+1} \left(\frac{\rho_{k+1}}{\gamma_{k+1}} - \frac{\rho_k}{\gamma_k}\right) + \sum_{k=0}^K \rho_k \gamma_k \frac{C}{2\tau} \end{aligned} \quad (12)$$

Note that $\rho_0 = 0$, so we simply dropped a negative term in last line. Applying the step size $\gamma_k = 2n\tau/(\tau^2 k + 2n)$, we get

$$\begin{aligned} \frac{\rho_{k+1}}{\gamma_{k+1}} - \frac{\rho_k}{\gamma_k} &= \frac{k+1}{S_K} \frac{\tau^2(k+1)2n}{2n\tau} - \frac{k}{S_K} \frac{\tau^2 k + 2n}{2n\tau} \\ &= \frac{1}{2nS_K\tau} [\tau^2(k+1)^2 + 2n(k+1) - \tau^2 k^2 - 2nk] \\ &= \frac{\tau^2(2k+1) + 2n}{2nS_K\tau}. \end{aligned}$$

Plug the above back into (12) and use the bound $h_{k+1} \leq 2nC/(\tau^2(k+1) + 2n)$, we get

$$\begin{aligned} \mathbb{E}(\min_{k \in [K]} g_k) &\leq \sum_{k=0}^K \rho_k g_k \leq \frac{nC}{\tau^2 S_K} \sum_{k=0}^{K-1} \frac{\tau^2(2k+2) + 2n}{2n} \frac{2n}{\tau(2k+1) + 2n} + \sum_{k=0}^K \frac{k}{S_K} \frac{2n\tau}{\tau^2 k + 2n} \frac{C}{2\tau} \\ &= \frac{nC}{\tau^2 S_K} \left[\sum_{k=0}^{K-1} \left(1 + \frac{\tau^2}{\tau^2(k+1) + 2n}\right) + \sum_{k=1}^K \frac{k\tau^2}{(\tau^2 k + 2n)} \right] \\ &\leq \frac{nC}{\tau^2 S_K} [2K + K] = \frac{2nC}{\tau^2(K+1)} \cdot 3. \end{aligned}$$

This completes the proof for $K \geq 1$. \square

Proof of Convergence with Delayed Gradient The idea is that we are going to treat the updates calculated from the delayed gradients as an additive error and then invoke our convergence results that allow the oracle to be approximate. We will first present a lemma that we will use for the proof of Theorem 6.

Lemma 9. *Let $x \in \mathcal{M}$, $\|\cdot\|$ be a norm, $\text{Diam}(\mathcal{M})_{\|\cdot\|} \leq D$, L be the gradient Lipschitz constant of f with respect to the given norm $\|\cdot\|$. Moreover, let x' be at most κ steps away from x and the largest stepsize in the past κ steps, and*

$$\begin{aligned} x^* &:= \underset{s \in \mathcal{M}}{\text{argmin}} \langle s, \nabla f(x) \rangle \\ \tilde{x} &:= \underset{s \in \mathcal{M}}{\text{argmin}} \langle s, \nabla f(x') \rangle \end{aligned}$$

Then, we have

$$\langle \tilde{s} - x, \nabla f(x) \rangle \leq \langle s^* - x, \nabla f(x) \rangle + \gamma\kappa D^2 L$$

Proof. Because \tilde{s} minimizes $\langle s, \nabla f(\tilde{x}) \rangle$ over $s \in \mathcal{M}$ and s^* is feasible, we can write

$$\langle s^* - \tilde{s}, \nabla f(\tilde{x}) \rangle \geq 0.$$

Using this and Hölder's inequality, we can write

$$\langle \tilde{s} - x, \nabla f(x) \rangle - \langle s^* - x, \nabla f(x) \rangle \leq \langle \tilde{s} - s^*, \nabla f(x) - \nabla f(\tilde{x}) \rangle \leq \|\tilde{s} - s^*\| \|\nabla f(\tilde{x}) - \nabla f(x)\|_* \leq DL \|\tilde{x} - x\|.$$

It remains to bound $\|\tilde{x} - x\|$.

$$\|\tilde{x} - x\| = \left\| \tilde{x} - \tilde{x} - \sum_{i=1}^{\kappa} \gamma_{-i} (s_{-i} - x_{-i}) \right\| \leq \gamma\kappa \max_i \|s_{-i} - x_{-i}\| \leq \gamma\kappa D,$$

where we used the fact that x is at most κ steps away from \tilde{x} . Assume γ_{-i} is the stepsize used and $\langle s_{-i}, x_{-i} \rangle$ are the actual updates that had been performed in the nearest i th parameter update before we get to x . \square

The second lemma that we need is the following.

Lemma 10. *Let \mathcal{M} be a convex set. Let $x_0 \in \mathcal{M}$. Let m be any positive integer. For $i = 1, \dots, m$, let $x_i = x_{i-1} + \gamma_i (s_i - x_{i-1})$ for some $0 \leq \gamma_i \leq 1$ and $s_i \in \mathcal{M}$. Then there exists an $s \in \mathcal{M}$ and $\gamma \leq \sum_{i=1}^m \gamma_i$, such that $x_m = \gamma(s - x_0) + x_0$.*

Proof. We prove by induction. When $m = 1$, $s = s_1$ and $\gamma = \gamma_1$. Assume for any $m = k - 1$, that the claim holds assume the condition is true, then by the recursive formula,

$$\begin{aligned} x_k &= x_{k-1} + \gamma_k (s_k - x_{k-1}) \\ &= x_0 + \gamma(s - x_0) + \gamma_k [s_k - x_0 - \gamma(s - x_0)] \\ &= x_0 - (\gamma + \gamma_k - \gamma_k \gamma) x_0 + (\gamma - \gamma_k \gamma) s + \gamma_k s_k \\ &= x_0 + (\gamma + \gamma_k - \gamma_k \gamma) \left[\frac{\gamma - \gamma_k \gamma}{\gamma + \gamma_k - \gamma_k \gamma} s + \frac{\gamma_k}{\gamma + \gamma_k - \gamma_k \gamma} s_k - x_0 \right] \\ &= x_0 + (\gamma + \gamma_k - \gamma_k \gamma) (s' - x_0) \end{aligned}$$

Note that s' is a convex combination of s_k and s therefore by convexity $s' \in \mathcal{M}$. Substitute $\gamma \leq \sum_{i=1}^{k-1} \gamma_i$, we get

$$\gamma + \gamma_k - \gamma_k \gamma \leq \sum_{i=1}^k \gamma_i.$$

This completes the inductive proof for all m . \square

The third Lemma that we will need is the following characterization of the expected ‘‘max load’’ in randomized load balancing.

Lemma 11 ((Mitzenmacher, 2001; Raab & Steger, 1998)). *Suppose m balls are thrown independently and uniformly at random into n bins. Then, the maximum number of balls in a bin Y satisfies*

$$\mathbb{E}Y \leq \begin{cases} \frac{3 \log n}{\log(n/m)} & \text{if } m < n/\log n, \\ c' \log n & \text{if } m < cn \log n, \\ \frac{m}{n} + O\left(\sqrt{\frac{2m}{n} \log n}\right) & \text{if } m \gg n \log n. \end{cases}$$

where c' is a constant that depends only on c .

Proof of Theorem 6. The proof involves a sharpening of the Lemma 9 for the BCFW and minibatch setting, where $x \in \mathcal{M} = \mathcal{M}^{(1)} \times \dots \times \mathcal{M}^{(n)}$ is a product domain. The proof idea is to exploit this property. Let the current update be on coordinate block index subset S . For each $j \in S$, let the corresponding worker be delayed by \varkappa_j steps, and the corresponding parameter vector be \tilde{x} . \varkappa_j is a random variable.

As in the proof of Lemma 9, we can bound the suboptimality of the approximate subroutine for solving problem j :

$$\begin{aligned} \text{Suboptimality}(\tilde{s}_j) &\leq \langle \tilde{s}_j - s_j^*, \nabla_j f(\tilde{x}) - \nabla_j f(x) \rangle \leq \|\tilde{s}_j - s_j^*\| \|\nabla_j f(\tilde{x}) - \nabla_j f(x)\|_* \\ &\leq D_{\|\cdot\|}^{(j)} L_{\|\cdot\|}^{(j)} \|\tilde{x} - x\| = D_{\|\cdot\|}^{(j)} L_{\|\cdot\|}^{(j)} \left\| \sum_{i=1}^{\varkappa_j} \gamma_{-i} (s_{-i} - x_{-i}) \right\| \\ &\leq D_{\|\cdot\|}^1 L_{\|\cdot\|}^1 \sum_{i=1}^{\varkappa_j} \gamma_{-i} \|(s_{-i} - x_{-i})\| \\ &\leq D_{\|\cdot\|}^1 L_{\|\cdot\|}^1 \sum_{i=1}^{\varkappa_j} \gamma_{-i} D_{\|\cdot\|}^\tau \leq \varkappa_j \gamma_{-\varkappa_j} D_{\|\cdot\|}^1 L_{\|\cdot\|}^1 D_{\|\cdot\|}^\tau. \end{aligned} \tag{13}$$

Let $\kappa := \mathbb{E}\varkappa_j$, take expectation on both sides we get

$$\mathbb{E} \text{Suboptimality}(\tilde{s}_j) \leq \mathbb{E}(\varkappa_j \gamma_{-\varkappa_j}) D_{\|\cdot\|}^1 L_{\|\cdot\|}^1 D_{\|\cdot\|}^\tau$$

Repeat the same argument for each $i \in S$, we get

$$\mathbb{E} \text{Suboptimality}(\tilde{s}) \leq \mathbb{E}(\varkappa_j \gamma_{-\varkappa_j}) \tau D_{\|\cdot\|}^1 L_{\|\cdot\|}^1 D_{\|\cdot\|}^\tau.$$

To put it into the desired format in (5), we solve the following inequality for δ

$$\frac{\gamma \delta C_f^\tau}{2} \geq \mathbb{E}(\varkappa_j \gamma_{-\varkappa_j}) \tau D_{\|\cdot\|}^\tau D_{\|\cdot\|}^1 L_{\|\cdot\|}^1$$

we get

$$\delta \geq \frac{2\tau}{C_f^\tau} \mathbb{E} \left(\frac{\varkappa_j \gamma_{-\varkappa_j}}{\gamma} \right) D_{\|\cdot\|}^\tau D_{\|\cdot\|}^1 L_{\|\cdot\|}^1.$$

By the specification of the stepsizes, we can calculate for each k ,

$$\frac{\gamma_{-\varkappa_j}}{\gamma} = \frac{\tau^2 k + 2n}{\tau^2 (\max(k - \varkappa_j, 0)) + 2n}.$$

Note that we always enforce \varkappa_j to be smaller than $\frac{k}{2}$ (otherwise the update is dropped), we can therefore upper bound $\mathbb{E}\left(\frac{\varkappa_j \gamma - \varkappa_j}{\gamma}\right)$ by 2κ . This gives us the first bound (9) on δ in Theorem 6.

To get the second bound on δ , we start from (13) and bound $\|\tilde{x} - x\|$ differently. Let S be the set of $\tau\varkappa_j$ coordinate blocks that were updated in the past \varkappa_j iterations. In the cases where fewer than $\tau\varkappa_j$ blocks were updated, just arbitrarily pick among the coordinate blocks that were updated 0 times so that $|S| = \tau\varkappa_j$. $\tilde{x} - x$ is supported only on S . Suppose coordinate block $i \in S$ is updated by m times, as below

$$\tilde{x}_{(i)} = \sum_{j=1}^m \gamma_j (s_j - [x_j]_{(i)})$$

for some sequence of $0 \leq \gamma_1, \dots, \gamma_m \leq 1$ and $s_1, \dots, s_m \in \mathcal{M}_i$ and recursively $[x_j]_{(i)} = [x_{j-1}]_{(i)} + \gamma_j (s_j - [x_{j-1}]_{(i)})$ ($x_0 = x$). Apply Lemma 10 for each coordinate block, we know that there exist $s_{(i)} \in \mathcal{M}_i$ in each block $i \in S$ such that

$$\tilde{x}_{(i)} = x_{(i)} + \gamma_{(i)} (s_{(i)} - x_{(i)})$$

with

$$\gamma_{(i)} \leq \sum_{j \in \text{iterations where } i \text{ is updated}} \gamma_j \leq m\gamma_{\max}. \quad (14)$$

Note that $s_{(i)} \in \mathcal{M}_i$ for each $i \in S$ implies that their concatenation $s_{(S)} \in \mathcal{M}_S$. Also $\gamma_{\max} \leq \gamma_{-\varkappa_j}$. Therefore

$$\|\tilde{x} - x\| = \left\| \sum_{i \in S} \gamma_{(i)} (s_{(i)} - x_{(i)}) \right\| \leq m\gamma_{\max} \|s_{(S)} - x_{(S)}\| \leq Y\gamma_{-\varkappa_j} D_{\|\cdot\|}^{\tau\varkappa_j}$$

where Y is a random variable that denotes the number of updates received by the most updated coordinate block (the maximum load). Apply a previously used argument to get $\gamma_{-\varkappa_j} < 2\gamma$, take expectation on both sides, to get the following by the law of total expectations and (14)

$$\begin{aligned} \mathbb{E}\|\tilde{x} - x\| &\leq \mathbb{E}\left(Y\gamma_{-\varkappa_j} D_{\|\cdot\|}^{\tau\varkappa_j}\right) = \mathbb{E}\left[\mathbb{E}\left(Y\gamma_{-\varkappa_j} D_{\|\cdot\|}^{\tau\varkappa_j} \mid \varkappa_j\right)\right] = \mathbb{E}\left[\gamma_{-\varkappa_j} D_{\|\cdot\|}^{\tau\varkappa_j} \mathbb{E}(Y \mid \varkappa_j)\right] \\ &\leq 2\gamma \mathbb{E}\left[D_{\|\cdot\|}^{\tau\varkappa_j} \mathbb{E}(Y \mid \varkappa_j)\right] \end{aligned} \quad (15)$$

This expectation is taken over the entire history of minibatch choice and delay associated with each update. When we condition on \varkappa_j , the conditional expectation of Y becomes the load-balancing problem.

By Lemma 11 when $\kappa_{\max}\tau \leq \frac{n}{\log n}$, it follows from (15) that

$$\mathbb{E}\|\tilde{x} - x\| \leq 2\gamma \mathbb{E} D_{\|\cdot\|}^{\varkappa_j \tau} \frac{3 \log n}{\log(n/\varkappa_j \tau)} \leq \frac{3 \log n}{\log[n/(\tau\kappa_{\max})]} 2\gamma \mathbb{E} D_{\|\cdot\|}^{\varkappa_j \tau}.$$

When $\kappa_{\max}\tau < cn \log n$,

$$\mathbb{E}\|\tilde{x} - x\| \leq 2\gamma \mathbb{E} D_{\|\cdot\|}^{\varkappa_j \tau} O(\log n) \leq O(\log n) 2\gamma \mathbb{E} D_{\|\cdot\|}^{\varkappa_j \tau}.$$

When $\kappa_{\max}\tau \gg n \log n$, then

$$\mathbb{E}\|\tilde{x} - x\| \leq (1 + o(1)) \frac{\tau\kappa_{\max}}{n} 2\gamma \mathbb{E} D_{\|\cdot\|}^{\varkappa_j \tau}.$$

Repeating the above results for each block $j \in S$, and summing them up leads to an upper bound for $\frac{\gamma\delta C_f^\tau}{2}$ and the proof of (9) is complete by solving for δ . \square

B. Proofs of other technical results

Relationship of the curvatures.

Proof of Lemma 1. $C_f^{(S)} \leq C_f$ follows from the fact that

$$\langle y_{[S]} - x_{[S]}, \nabla_{(S)} f(x) \rangle = \langle y_{[S]} - x_{[S]}, \nabla f(x) \rangle,$$

and $s_{[S]} \in \mathcal{M}$. In other words, the arg sup of (3) is a feasible solution in the sup to compute the global C_f . Similar argument holds for the proof $C_f^{(i)} \leq C_f^{(S)}$ as $i \in S$.

In the second part,

$$C_f^\tau = \frac{1}{\binom{n}{\tau}} \sum_{T \subseteq [n], |T|=\tau} C_f^{(T)}.$$

We can evenly partition sets T in the summation into n parts P_j for $j \in [n]$, such that sets in P_j have the element j . Clearly each P_j has a size of $\binom{n}{\tau}/n$. We can use $C_f^{(S)} \geq C_f(j)$ from the first inequality of the lemma, to get the inequality below.

$$C_f^\tau = \frac{1}{\binom{n}{\tau}} \sum_{j \in [n]} \sum_{T \in P_j} C_f^{(T)} \geq \frac{1}{\binom{n}{\tau}} \sum_{j \in [n]} \sum_{T \in P_j} C_f^{(j)} = \frac{1}{\binom{n}{\tau}} \sum_{j \in [n]} \binom{n}{\tau} \frac{1}{n} C_f^{(j)} = \frac{1}{n} C_f^\otimes$$

The relaxation of C_f^τ to C_f is trivial since $C_f^{(T)} \leq C_f$ holds for any $T \subseteq [n]$ from the first part of the lemma. \square

Bounding C_f^τ using expected boundedness and expected incoherence

Proof of Theorem 4. By Definition of H , for any $x, z \in \mathcal{M}$, $\gamma \in [0, 1]$

$$f(x + \gamma(z - x)) \leq f(x) + \gamma(z - x)^T \nabla f(x) + \frac{\gamma^2}{2} (z - x)^T H(z - x).$$

Rearranging the terms we get

$$\frac{2}{\gamma^2} [f(x + \gamma(z - x)) - f(x) - \gamma(z - x)^T \nabla f(x)] \leq (z - x)^T H(z - x)$$

The definition of set curvature (3) is written in an equivalent notation with $z = x_{[S^c]} + s_{[S]}$ and $y = x + \gamma(z - x) = x + \gamma(s_{[S]} - x_{[S]})$. So we know the support of $z - x$ is constrained to be within the coordinate blocks S .

Plugging this into the definition of (3) we get an analog of Equation (2.12) in (Jaggi, 2011) for $C_f^{(S)}$.

$$\begin{aligned} C_f^{(S)} &= \sup_{\substack{x, z \in \mathcal{M}, \gamma \in [0, 1] \\ z_{(S^c)} - x_{(S^c)} = 0}} \frac{2}{\gamma^2} [f(x + \gamma(z - x)) - f(x) - \gamma(z - x)^T \nabla f(x)] \\ &\leq \sup_{\substack{x, z \in \mathcal{M}, \\ z_{(S^c)} - x_{(S^c)} = 0}} (z - x)^T H(z - x) = \sup_{\substack{x, z \in \mathcal{M}, \\ z_{(S^c)} - x_{(S^c)} = 0}} s_{(S)}^T H s_{(S)} \\ &\leq \sup_{w \in \mathcal{M}^{(S)}} (2w^T) H_S (2w) = 4 \left\{ \sup_{w_i \in \mathcal{M}^{(i)} \forall i \in S} \sum_{i \in S} w_i^T H_{ii} w_i + \sum_{i, j \in S, i \neq j} w_i^T H_{ii} w_j \right\} \\ &\leq 4 \left\{ \sum_{i \in S} \sup_z \sup_{w_i} w_i^T H_{ii}(z) w_i + \sum_{i, j \in S, i \neq j} \sup_z \sup_{w_i, w_j} w_i^T H_{ii}(z) w_j \right\} \\ &\leq 4 \left(\sum_{i \in S} B_i + \sum_{i, j \in S, i \neq j} \mu_{ij} \right). \end{aligned}$$

Take expectation for all possible S of size τ and we obtain the lemma statement. \square

Proof of the example with sublinear dependence of κ

Proof of Lemma 7. We first show that a continuous extension of $D_{\|\cdot\|}^\theta$ is concave in θ

$$\begin{aligned} D_{\|\cdot\|}^\theta &= \max_{S \subset [n] | |S|=\theta} \sup_{x, y \in \mathcal{M}^{(S)}} \|x - y\| \\ &= \max_{S \subset [n] | |S|=\theta} \sup_{x, y \in \mathcal{M}^{(S)}} \sqrt{\sum_{i \in S} \|x_{(i)} - y_{(i)}\|^2} \\ &= \sqrt{\max_{S \subset [n] | |S|=\theta} \sum_{i \in S} \sup_{x_{(i)}, y_{(i)} \in \mathcal{M}^{(i)}} \|x_{(i)} - y_{(i)}\|^2} \end{aligned}$$

The supremum is obtained by sorting and the function in the square root is concave function of θ , when we extend the support of this function to \mathbb{R}_+ through linear interpolation. By the composition theorem, the square root of that is also a concave function in τ . We call this function $\tilde{D}_{\|\cdot\|}^\theta$. Note that $\tilde{D}_{\|\cdot\|}^\theta = D_{\|\cdot\|}^\theta$ when $\theta \in [n]$ such that if we take expectation over the any discrete distribution over θ , their expectations are the same. It follows from Jensen's inequality that

$$\begin{aligned} \mathbb{E} D_{\|\cdot\|}^{\mathcal{X}\tau} &= \mathbb{E} \tilde{D}_{\|\cdot\|}^{\mathcal{X}\tau} \leq \tilde{D}_{\|\cdot\|}^{\mathbb{E} \mathcal{X}\tau} \\ &\leq D_{\|\cdot\|}^{\lceil \mathbb{E} \mathcal{X}\tau \rceil} \\ &= \sqrt{\max_{S \subset [n] | |S|=\lceil \mathbb{E} \mathcal{X}\tau \rceil} \sum_{i \in S} \sup_{x_{(i)}, y_{(i)} \in \mathcal{M}^{(i)}} \|x_{(i)} - y_{(i)}\|^2} \\ &\leq \sqrt{\mathbb{E} \mathcal{X}} D_{\|\cdot\|}^\tau. \end{aligned}$$

□

Proof of specific examples

Proof of Example 1. First of all, $H = \lambda A^T A$. Since all columns of A have the same magnitude $\sqrt{2}/n$. By the Holder's inequality and the 1-norm constraint in every block, we know $B_i = \frac{2}{n^2 \lambda}$ for any i therefore $B = \frac{2}{n^2 \lambda}$. Secondly, by well-known upper bound for the area of the spherical cap, which says for any fixed vector z and random vector a on a unit sphere in \mathbb{R}^d ,

$$\mathbb{P}(|\langle z, a \rangle| > \epsilon \|z\|) \leq 2e^{-\frac{d\epsilon^2}{2}},$$

we get

$$\mathbb{P}(\mu_{ij} > 2\sqrt{\frac{20 \log d}{d}}) \leq \frac{2}{d^{10}}.$$

Take union bound over all pairs of labels we get the probability as claimed. □

Proof of Example 2. The matrix $D^T D$ is tridiagonal with 2 on the diagonal and -1 on the off-diagonal. If we vectorize U by concatenating $u = [u_1; \dots; u_{n-1}]$, the Hessian matrix for u will be $H = \Pi I_d \otimes (D^T D) \Pi^T$ where Π is some permutation matrix. Without calculating it explicitly, we can express

$$\begin{aligned} u_S^T H_S u_S &= u_S^T (D^T \otimes 1_d) (D^T \otimes 1_d)^T u_S \\ &= \sum_{i \in S} u_i^T \begin{bmatrix} D_{:,i}^T \\ D_{:,i}^T \\ \vdots \\ D_{:,i}^T \end{bmatrix} [D_{:,i} \ D_{:,i} \ \dots \ D_{:,i}] u_i + \sum_{i, j \in S, i \neq j} u_i^T \begin{bmatrix} D_{:,i}^T \\ D_{:,i}^T \\ \vdots \\ D_{:,i}^T \end{bmatrix} [D_{:,j} \ D_{:,j} \ \dots \ D_{:,j}] u_j. \end{aligned}$$

We note that for any $|i - j| \geq 2$, the second term is 0. Apply the constraint that $\|u_i\|_2 \leq \lambda$ and the fact that the ℓ_2 operator norm of $[D_{:,j} \ D_{:,j} \ \dots \ D_{:,j}]$ is $\sqrt{2d}$, we get $B_i = 2\lambda^2 d$. Similarly, $2(n-2)$ nonzero obeys $\mu_{ij} = \lambda^2 d$. This allows us to obtain an upper bound

$$C_f^\tau \leq 4 \left[2\tau \lambda^2 d + \frac{2(n-2)\tau(\tau-1)}{(n-2)(n-1)} \lambda^2 d \right] \leq 16\tau \lambda^2 d.$$

which scales with τ . □

B.1. Pseudocode for the Multicore Shared Memory Architecture

We present pseudocode for the multicore shared memory setting here. It is the same except that each worker becomes a thread, the network buffer of servers become the a data structure, the workers' network buffer becomes the shared parameter vector and the workers can write to the data structure or the shared parameter vector directly.

Algorithm 2 AP-BCFW: Asynchronous Parallel Block-Coordinate Frank-Wolfe (Shared memory)

-----SERVER THREAD-----

Input: An initial feasible $x^{(0)}$, mini-batch size τ , number of workers T .

0. Write $x^{(0)}$ to shared memory. Declare a container (a queue or a stack).

for $k = 1, 2, \dots$ (k is the iteration number.) **do**

1. Keep popping the container until we have τ updates on τ disjoint blocks. Denote the index set by S .
2. Set step size $\gamma = \frac{2n\tau}{\tau^2 k + 2n}$.
3. Write sparse updates $x^{(k)} = x^{(k-1)} + \gamma \sum_{i \in S} (s_{[i]} - x_{[i]}^{(k-1)})$ into the shared memory.

if converged **then**
Broadcast STOP signal to all threads and break.

end if

end for

Output: $x^{(k)}$.

-----WORKER THREADS-----

while no STOP signal received **do**

- a. Randomly choose $i \in [n]$.
- b. Calculate partial gradient $\nabla_{(i)} f(x)$ using x in the shared memory and solve (2).
- c. Push $\{i, s_{(i)}\}$ to the container.

end while

The above pseudo code can be further simplified when $\tau = 1$. In particular, we do not need a server any more. Each worker can simply write to the shared memory bus. The probability of two workers writing to the same block is small as we analyzed in Section D.2. The algorithm essentially lock-free as in (Niu et al., 2011) modulo requiring the updates of each coordinate block to be atomic. Niu et al. (2011) is stronger in that it allows each scalar addition to be atomic.

There is an additional restriction due to the fixed predefined sequence of step sizes, which in fact requires a centralized shared counter that is atomic, so that no two threads have simultaneously the same k . In practice, we can simply choose a fixed sequence of stepsize for each worker separately.

Algorithm 3 AP-BCFW: Asynchronous Parallel Block-Coordinate Frank-Wolfe (Lock-Free Shared-Memory)

Input: An initial feasible $x^{(0)}$, number of workers T , a centralized counter.

0. Write $x^{(0)}$ to shared memory.

-----INDEPENDENTLY ON EACH THREAD-----

while not converged **do**

- a. Randomly choose $i \in [n]$.
- b. Calculate partial gradient $\nabla_{(i)} f(x)$ using x in the shared memory and solve (2).
- c. Read centralized counter for k . Set step size $\gamma = \frac{2n}{k+2n}$.
- d. Add $\gamma(s_{(i)} - x_{(i)})$ to block i of the shared memory.
- e. Increment the counter $k = k + 1$.

end while

if converged **then**
Output: $x^{(k)}$. and break.

end if

C. Application to Structural SVM

We briefly review structural SVMs and show how to solve the associated convex optimization problem using our AP-BCFW method.

In structured prediction setting, the task is to predict a structured output $\mathbf{y} \in \mathcal{Y}$, given $\mathbf{x} \in \mathcal{X}$. For example, \mathbf{x} could be the pixels in the picture of a word, \mathbf{y} could be the sequence of characters in the word. A feature map $\phi : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^d$ encodes compatibility between inputs and outputs. A linear classifier parameter \mathbf{w} is learned from data so that $\operatorname{argmax}_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{w}, \phi(\mathbf{x}, \mathbf{y}) \rangle$ gives the output for an input \mathbf{x} . Suppose we have the training data $\{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^n$ to learn \mathbf{w} . Define $\psi_i(\mathbf{y}) := \phi(\mathbf{x}_i, \mathbf{y}_i) - \phi(\mathbf{x}_i, \mathbf{y})$ and let $L_i(\mathbf{y}) := L(\mathbf{y}_i, \mathbf{y})$ denote the loss incurred by predicting \mathbf{y} instead of the correct output \mathbf{y}_i . The classifier parameter \mathbf{w} is learned by solving the optimization problem

$$\begin{aligned} \min_{\mathbf{w}, \xi} \quad & \frac{\lambda}{2} \|\mathbf{w}\|^2 + \frac{1}{n} \sum_{i=1}^n \xi_i \\ \text{s.t.} \quad & \langle \mathbf{w}, \psi_i(\mathbf{y}) \rangle \geq L(\mathbf{y}_i, \mathbf{y}) - \xi_i \quad \forall i, \mathbf{y} \in \mathcal{Y}(\mathbf{x}_i). \end{aligned} \quad (16)$$

We solve the dual of this problem using our method. We introduce some more notation to formulate the dual. Denote $\mathcal{Y}_i := \mathcal{Y}(\mathbf{x}_i)$, the set of possible labels for \mathbf{x}_i . Note that $|\mathcal{Y}_i|$ is exponential in the length of label \mathbf{y}_i . Let $m = \sum_{i=1}^n |\mathcal{Y}_i|$. Let $A \in \mathbb{R}^{d \times m}$ denote a matrix whose m columns are given by $\{\frac{1}{\lambda n} \psi_i(\mathbf{y}) \mid i \in [n], \mathbf{y} \in \mathcal{Y}_i\}$. Let $b \in \mathbb{R}^m$ be a vector given by the entries $\{\frac{1}{n} L_i(\mathbf{y}) \mid i \in [n], \mathbf{y} \in \mathcal{Y}_i\}$. The dual of (16) is given by

$$\begin{aligned} \min_{\alpha \in \mathbb{R}^m} \quad & f(\alpha) := \frac{\lambda}{2} \|A\alpha\|^2 - b^T \alpha \\ \text{s.t.} \quad & \sum_{\mathbf{y} \in \mathcal{Y}_i} \alpha_i(\mathbf{y}) = 1 \quad \forall i \in [n], \alpha \geq 0 \end{aligned} \quad (17)$$

The primal solution \mathbf{w} can be retrieved from the dual solution α from the relation $\mathbf{w} = A\alpha$ obtained from KKT conditions. Also note that the domain \mathcal{M} of (17) is exactly the product of simplices $\mathcal{M} = \Delta_{|\mathcal{Y}_1|} \times \cdots \times \Delta_{|\mathcal{Y}_n|}$.

The subproblem in equation (2) takes a well-known form in the Frank-Wolfe setup for solving (17). The gradient is given by

$$\nabla f(\alpha) = \lambda A^T A \alpha - b = \lambda A^T \mathbf{w} - b$$

whose (i, \mathbf{y}) -th component is given by $\frac{1}{n} (\langle \mathbf{w}, \psi_i(\mathbf{y}) \rangle - L_i(\mathbf{y}))$. Define $H_i(\mathbf{y}; \mathbf{w}) := L_i(\mathbf{y}) - \langle \mathbf{w}, \psi_i(\mathbf{y}) \rangle$ so that the (i, \mathbf{y}) -th component of the gradient is $-\frac{1}{n} H_i(\mathbf{y}; \mathbf{w})$. In the subproblem (2), the domain $\mathcal{M}^{(i)}$ is the simplex $\Delta_{\mathcal{Y}_i}$ and the block gradient $\nabla_{(i)} f(\alpha)$ is linear. So, the objective is minimized at a corner of the simplex $\mathcal{M}^{(i)}$ and the optimum value is simply given by $\min_{\mathbf{y}} \nabla_{(i)} f(\alpha)$ which can be rewritten as $\max_{\mathbf{y}} H_i(\mathbf{y}; \mathbf{w})$. Further, the corner can be explicitly written as the indicator vector $e^{\mathbf{y}_i^*} \in \mathcal{M}^{(i)}$ where $\mathbf{y}_i^* = \operatorname{argmax}_{\mathbf{y}} H_i(\mathbf{y}; \mathbf{w})$. It turns out that this maximization problem can be solved efficiently for several problems. For example, when the output is a sequence of labels, a dynamic programming algorithm like Viterbi can be used.

As mentioned before, m is too large to update the dual variable α directly. So, we make an update to the primal variable $\mathbf{w} = A\alpha$ instead. The Block-Coordinate Frank-Wolfe update for the i -th block maybe written as $\alpha_{(i)}^{k+1} = \alpha_{(i)}^k + \gamma(s_i - \alpha_{(i)}^k)$ where γ is the step-size. Recalling that the optimal s_i is $e^{\mathbf{y}_i^*}$, by multiplying the previous equation by A_i , we arrive at $\mathbf{w}_i^{(k+1)} = \mathbf{w}_i^k + \gamma(A_{i, \mathbf{y}_i^*} - \mathbf{w}_i^{(k)})$ where $\mathbf{w}_i^{(k)} := A_i \alpha_{(i)}^k$. From this definition of $\mathbf{w}_i^{(k)}$, the primal update is obtained by noting that $\mathbf{w}^{(k)} = \sum_i \mathbf{w}_i^{(k)}$. Explicitly, the primal update is given by $\mathbf{w}^{(k+1)} = \mathbf{w}^k + \gamma(A_{i, \mathbf{y}_i^*} - \mathbf{w}_i^{(k)})$. Note that $A_{i, \mathbf{y}_i^*} = \frac{1}{\lambda n} \psi(\mathcal{Y}_i^*)$. This Block-Coordinate version can be easily extended to AP-BCFW. In our shared memory implementation, for OCR dataset, we do the line search computation and $\mathbf{w}_i^{(k)}$ update step on the workers instead of the server because these computations turn out to be expensive enough to make the server the bottleneck even for modest number of workers.

D. Other technical results and discussions

D.1. Oracle assumption and heterogeneous blocks

Recall that our results rely on the oracle assumption that \mathcal{O} provides updates that are iid uniform over $[n]$ (Assumption A1). We discuss the implications and limitations of this assumption and then propose possible solutions.

Consider the setting where \mathcal{O} consists of T possibly heterogeneous workers and each worker samples iid from $[n]$. As we discussed before, A1 holds under the additional condition that the time needed to complete one subroutine solve for Block i by Worker j does not depend on i .

Consider the simple example due to an anonymous reviewer: Let $\tau = 1$, $T = 2$ and there are a total of two blocks. Block 1 takes only a millisecond and Block 2 takes a year to solve for both workers. In this case, the first update received by the server is with probability $3/4$ for Block 1 and only $1/4$ for Block 2.

This could potentially limit the use of our parallel algorithm for applications such as structured predictions where sentences having different lengths, or cases where there are different sparsity level over data points/constraints depending on how we formulate the problem.

This is in fact not a problem unique to us, Assumption A1 is implicitly required in most existing analysis for asynchronous stochastic algorithms (e.g., Liu et al., 2014; Niu et al., 2011). As a result, they all share the same woe. One could argue that parallelization is the wrong problem to address when block subroutines significantly differ with each other. Efforts should be spent on perhaps solving the expensive subproblem in parallel. But still, even mild heterogeneity over blocks invalidates our convergence result.

Henceforth, we propose two simple ways to address this issue and discuss their pros and cons.

Padding: A naive solution is to per-calculate the time-complexity with respect to each block and inject artificial time padding on each user such that all blocks have the same time complexity.

Pre-select S : An alternative is to let the server randomly choose a coordinate subset S of size τ , and the workers can only work on S , either by independently sample from S or work on whichever that is not available.

Neither of the two solutions is completely satisfactory. The padding approach ensures all results in the paper to hold including those for the delayed oracles, but inevitably, the time to complete each block now depends on the most expensive block. The second approach has milder dependence on the worst block, in fact it depends only on the time for the fastest worker to solve the slowest problem in each chosen S . However, it requires sending an updated parameter to all workers in every iteration. It could still be robust to heterogeneous workers when τ is several times larger than T , and when workers work asynchronously within the mini-batch, we prove in Proposition 12 that the number of collisions is small.

Fully asynchronous parallelism over blocks with heterogeneous blocks without dependence on the slowest block remains an important open problem.

D.2. Controlling collisions in distributed setting

In the distributed setting, different workers might end up working on the same slot.

In Algorithm 1, different workers may end up working on the same coordinate block and the server will drop a number of updates in case of collision. The following proposition shows that for this potential redundancy is not excessive is small and for a large range of τ , we also show additional strong concentration to its mean.

Proposition 12. *In the distributed asynchronous update scheme above:*

- i) *The expected number of subroutine calls from all workers to complete each iteration is $\tau + \sum_{i=1}^{\tau-1} \frac{i}{n-i}$.*
- ii) *If $0.02n < \tau < 0.6n$, with probability at least $1 - \exp(-n/60)$, no more than 2τ random draws (2τ subroutine calls in total from all workers) suffice to complete each iteration.*

Proof. The first claim is the well-known coupon collector problem.

The second claim requires an upper bound of the expectation. In expectation, we need $\frac{n}{n-k}$ balls to increase the unique

count from k to $k + 1$. So in expectation we need

$$\begin{aligned} 1 + \frac{n}{n-1} + \frac{n}{n-2} + \dots + \frac{n}{n-\tau+1} &= \tau + \sum_{i=1}^{\tau-1} \frac{i}{n-i} \\ &\leq \tau + \frac{1+2+\dots+(\tau-1)}{n-\tau+1} = \tau + \frac{\tau(\tau-1)}{2(n-\tau+1)} < \tau \left[1 + \frac{1}{2(n/\tau-1)} \right]. \end{aligned}$$

To see the second claim, first defined f_t to be the number of non-empty bins after t random ball throws, which can be consider as a function of the t iid ball throws X_1, X_2, \dots, X_t . It is clear that if we change only one of the X_i , f_t can be changed by at most 1. Also, note that the probability that any one bin being filled is $1 - (1 - \frac{1}{n})^t$, so $\mathbb{E}f_t = n \left[1 - (1 - \frac{1}{n})^t \right]$.

By the McDiarmid's inequality, $\mathbb{P}[f_t < \mathbb{E}f_t - \epsilon] \leq \exp \left[-\frac{2\epsilon^2}{t} \right]$. Take $t = 2\tau$, and $\epsilon = \mathbb{E}f_{2\tau} - \tau$, then

$$\begin{aligned} \mathbb{P}[f_{2\tau} < \tau] &\leq \exp \left[-\frac{1}{\tau} \left(n \left[1 - \left(1 - \frac{1}{n} \right)^{2\tau} \right] - \tau \right)^2 \right] \leq \exp \left[-\frac{1}{\tau} \left(n \left[1 - e^{-\frac{2\tau}{n}} \right] - \tau \right)^2 \right] \\ &= \exp \left[-n \cdot \frac{n}{\tau} \left(1 - e^{-\frac{2\tau}{n}} - \frac{\tau}{n} \right)^2 \right] \leq \exp[-Cn], \end{aligned}$$

where C is some constant which is the smaller of the two evaluations of the function $\frac{n}{\tau} \left(1 - e^{-\frac{2\tau}{n}} - \frac{\tau}{n} \right)^2$ at $\tau = 0.02n$ and $\tau = 0.6n$ (where the function is concave between the two). As a matter of fact, C can be taken as $\frac{1}{60}$.

Let g_τ be the number of balls that one throws that fills τ bins, the result is proven by noting that

$$\mathbb{P}(g_\tau \leq 2\tau) = \mathbb{P}(f_{2\tau} \geq \tau) \geq 1 - \exp[-Cn].$$

□

D.3. Curvature and Lipschitz Constant

In this section, we illustrate the relationship between the coordinate curvature constant, coordinate gradient Lipschitz conditions, and work out the typical size of the constants in Theorem 6. For the sake of discussion, we will focus on the quadratic function $f(x) = \frac{x^T A x}{2} + b^T x$. We start by showing that for quadratic function. The constant that one can get via choosing a specific norm can actually match the curvature constant. To be completely explicit, we define gradient Lipschitz constant $L_{\|\cdot\|}$ with respect to a norm $\|\cdot\|$, this requires that for any x, y ,

$$\|\nabla f(y) - \nabla f(x)\|_* \leq L\|x - y\|.$$

where $\|\cdot\|_*$ is the dual norm.

Proposition 13. *For quadratic functions with Hessian $A \succeq 0$, there exists a norm $\|\cdot\|$ such that the curvature constant $C_f = [D_{\|\cdot\|}]^2 L_{\|\cdot\|}$.*

Proof. We will show that this norm is simply the A -norm, $\|\cdot\|_A = \sqrt{(\cdot)^T A (\cdot)}$. The upper bound $C_f \leq [D_{\|\cdot\|_A}]^2 L_{\|\cdot\|_A}$ is a direct application of the result in Jaggi (2013, Appendix D). To show a lower bound it suffices to construct $s, x \in \mathcal{M}, \gamma \in [0, 1]$ and $y = \gamma s + (1 - \gamma)x$ such that

$$\frac{2}{\gamma^2} (f(y) - f(x) - \langle y - x, \nabla f(x) \rangle) = [D_{\|\cdot\|_A}]^2 L_{\|\cdot\|_A}.$$

For quadratic functions,

$$\frac{2}{\gamma^2} [f(y) - f(x) - \langle y - x, \nabla f(x) \rangle] = \frac{1}{2} (y - x)^T A (y - x) = \frac{1}{\gamma^2} \|y - x\|_A^2$$

Take $\gamma = 1$ and y, x on the boundary of \mathcal{M} such that $\|y - x\|_A = D_{\|\cdot\|_A}$, as a result, we get $C_f \geq D_{\|\cdot\|_A}^2$. It remains to show that the gradient Lipschitz constant with respect to \mathcal{A} -norm is 1, which directly follows from the Taylor expansion. \square

Similar arguments work for $C_f^{(i)}$ and $C_f^{(S)}$ under the same norm. Clearly, this means that the corresponding restriction of the subset domain has $A_{i,i}$ -norm or $A_{(S)}$ -norm.

We now consider the approximation constants due to the delays in Theorem 6, and work out more explicit bounds for quadratic functions and carefully chosen norm. Recall that the simple bound (8) has constant δ in the order of

$$\frac{\kappa\tau L_{\|\cdot\|}^1 D_{\|\cdot\|}^1 D_{\|\cdot\|}^\tau}{C_f^\tau}.$$

Suppose we use the A -norm, then $L_{\|\cdot\|}^1 = L_{\|\cdot\|}^\tau = 1$, and $C_f^\tau = [D_{\|\cdot\|}^\tau]^2$, the bound can be reduced to

$$\delta = O\left(\frac{\tau D_{\|\cdot\|}^1}{D_{\|\cdot\|}^\tau}\right) = O(\kappa\sqrt{\tau}).$$

where the last step requires \mathcal{M}_i to be all equivalent and A to be block-diagonal with identical $A_{(i)}$.

Similarly the strong bound (9) has constant δ in the order of

$$\delta = \tilde{O}\left(\frac{\tau L_{\|\cdot\|}^1 D_{\|\cdot\|}^1 D_{\|\cdot\|}^{\kappa\tau}}{C_f^\tau}\right) = \tilde{O}\left(\frac{\tau D_{\|\cdot\|}^1 D_{\|\cdot\|}^{\kappa\tau}}{[D_{\|\cdot\|}^\tau]^2}\right) = \tilde{O}(\sqrt{\kappa\tau})$$

Again, the last step requires a strong assumption that \mathcal{M}_i to be all equivalent and A to be block-diagonal with identical diagonal blocks. While these calculations only apply to specific case of a quadratic function with a lot of symmetry, we conjecture that in general the flexibility of choosing the norm will allow the ratio of these boundedness constants and C_f^τ to be a well-controlled constant and the typical dependence on the system parameter τ and κ should stay within the same ballpark.

D.4. Examples and illustrations

In this section, we now derive specific instances of the Theorem 4 for the structural SVM and Group Fused Lasso. For the structural SVM, a simple generalization of [Lacoste-Julien et al. \(2013, Lemmas A.1, A.2\)](#) shows that in the worst case, using $\tau > 1$ offers no gain at all. Fortunately, if we consider more specific problems, using larger τ does yield faster convergence. We provide two such examples below.

Example 1 (Structural SVM for multi-label classification (with random data)). We describe the application to structural SVMs in detail in Section C (please see this section for details on notation). Here, we describe the convergence rate for this application. According to [\(Yu & Joachims, 2009\)](#), the compatibility function $\phi(x, y)$ for multiclass classification will be $[0, \dots, 0, x^T, 0, \dots, 0]^T / \lambda n$ where the only nonzero block that we fill with the feature vector is the (y) th block. So $\psi_i(x_i, j) = \phi(x_i, y_i) - \phi(x_i, j)$ looks like $[0, \dots, 0, x_i^T, 0, \dots, 0, -x_i^T, 0, \dots, 0]^T / \lambda n$. This already ensures that $B = \frac{2}{n^2\lambda}$ provided x_i lie on a unit sphere. Suppose we have K classes and each class has a unique feature vector drawn randomly from a unit sphere in \mathbb{R}^d ; furthermore, for simplicity assume we always draw $\tau < K$ data points with τ distinct labels⁴ $\mu \leq \sqrt{\frac{c \log d}{d} \frac{2}{n^2\lambda}}$, for some constant c . In addition, if $d \geq \tau^2 \sqrt{c \log d}$, then with high probability

$$C_f^\tau \leq \frac{8\tau + 8\tau^2 \sqrt{\frac{c \log d}{d}}}{n^2\lambda} = O\left(\frac{c\tau}{n^2\lambda}\right),$$

which yields a convergence rate $O\left(\frac{R^2}{\lambda\tau k}\right)$, where $R :=$

$\max_{i \in [n], y \in \mathcal{Y}_i} \|\psi_i(y)\|_2$ using notation from Lemmas A.1 and A.2 of [Lacoste-Julien et al. \(2013\)](#).

This analysis suggests that a good rule-of-thumb is that we should choose τ to be at most the number of categories for the classification. If each class is a mixture of random draws from the unit sphere, then we can choose τ to be the underlying number of mixture components.

⁴This is an oversimplification but it offers a rough rule-of-thumb. In practice, C_f^τ should be in the same ballpark as our estimate here.

Example 2 (Group Fused Lasso). The Group Fused Lasso aims to solve (typically for $q = 2$)

$$\min_X \frac{1}{2} \|X - Y\|_F^2 + \lambda \|XD\|_{1,q}, \quad q > 1, \quad (18)$$

where $X, Y \in \mathbb{R}^{d \times n}$, and column y_t of Y is an observed noisy d -dimensional feature vector at time $1 \leq t \leq n$. The matrix $D \in \mathbb{R}^{n \times (n-1)}$ is the differencing matrix that takes the difference of feature vectors at adjacent time points (columns). The formulation aims to filter the trend that has some piecewise constant structures. The dual to (18) is

$$\begin{aligned} \max_U & -\frac{1}{2} \|UD^T\|_F^2 + \text{tr} UD^T Y^T \\ \text{s.t.} & \|U_{:,t}\|_p \leq \lambda, \quad \forall t = 1, \dots, n-1, \end{aligned}$$

where p is conjugate to q , i.e., $1/p + 1/q = 1$. This block-constrained problem fits our structure (1). For this problem, we find that $B \leq 2\lambda^2 d$ and $\mu \leq 2\lambda^2 d/(n-1)$, which yields

$$C_f^\tau \leq 16\tau\lambda^2 d.$$

Consequently, the rate of convergence becomes $O(\frac{n^2\lambda^2 d}{\tau k})$. In this case, batch FW will have a better rate of convergence than BCFW⁵.

Example 3 (Structural SVM worst-case bound). For structural SVM with arbitrary data (including even pathological/trivial data), using notation from Lemmas A.1 and A.2 of [Lacoste-Julien et al. \(2013\)](#), define $R := \max_{i \in [n], y \in \mathcal{Y}_i} \|\psi_i(y)\|_2$. Then we can provide an upper bound

$$B, \mu \leq \frac{R^2}{\lambda n^2} \implies C_f^\tau \leq \frac{4\tau^2 R^2}{\lambda n^2}. \quad (19)$$

In this case, for any $\tau = 1, \dots, n$, the rate of convergence will be the same $O(\frac{R^2}{\lambda k})$.

An illustration for the group fused lasso Figure 7 shows a typically application for group fused lasso (filtering piecewise constant multivariate signals whose change points are grouped together).

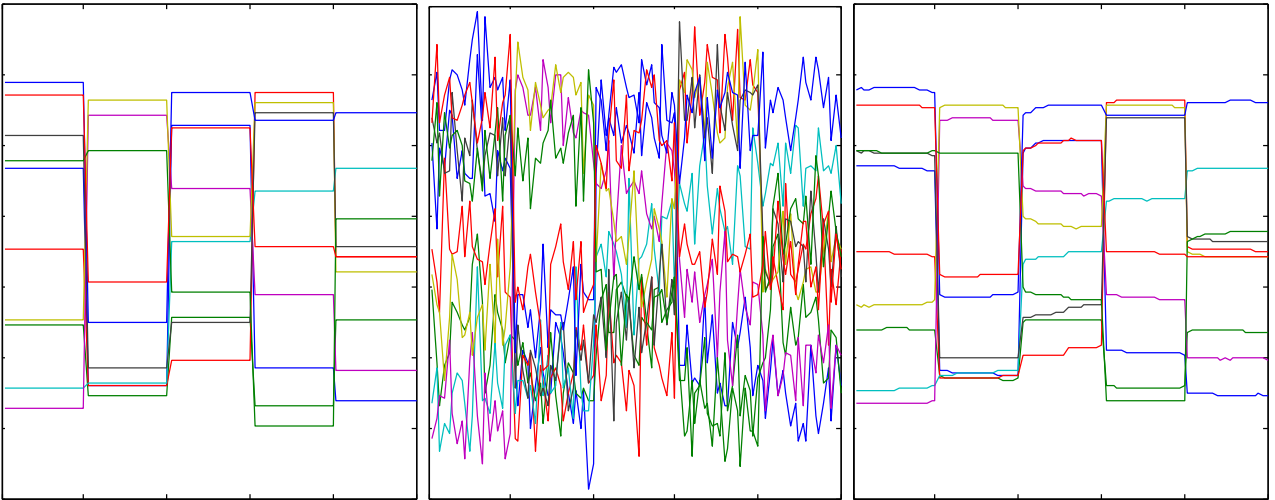


Figure 7. Illustration of the signal data used in the Fused Lasso experiments. We show the original signal (left), the noisy signal given to the algorithm (middle), and the signal recovered after performing the fused lasso optimization (right).

⁵Observe that C_f^τ does not have an n^2 term in the denominator to cancel out the numerator. This is because the objective function is not appropriately scaled with n like it does in the structural SVM formulation.

D.5. Comparison to parallel block coordinate descent

With some understanding on C_f^τ , we can now explicitly compare the rate of convergence in Theorem 2 with parallel BCD (Liu et al., 2014; Richtárik & Takáč, 2015) under the assumption of $\mu = O(B/\tau)$ — a fair and equally favorable case to all of these methods. We acknowledge that more general treatments of ESO property in more recent extensions of Richtárik & Takáč (2015) in a similar flavor as our (7) (see e.g., Qu & Richtárik 2014) but similar results are not available for the asynchronous version. To facilitate comparison, we will convert the constants in all three methods to block coordinate gradient Lipschitz constant L_i , which obeys

$$f(x + s_{[i]}) \leq f(x) + \langle s_{[i]}, \nabla f(x) \rangle + L_i \|s_{[i]}\|^2, \quad (20)$$

for any $x \in \mathcal{M}$, $s_{(i)} \in \mathcal{M}_i$. Observe that $B_i \leq 4L_i \text{diam}(\mathcal{M}_i)^2 = L_i \max_{x_i^*, x_i \in \mathcal{M}_i} \|x_i - x_i^*\|^2$, so

$$B \leq \frac{1}{n} \sum_i L_i \max_{x_i, x_i^*} \|x_i - x_i^*\| \quad (21)$$

$$\leq \frac{1}{n} \sum_i L_i \max_x \|x - x^*\|^2 = \mathbb{E}_i(L_i) R^2 \quad (22)$$

where $R := \max_x \|x - x^*\|$. The rate of convergence for the three methods (with τ oracle calls considered as one iteration) are given below.

Method	Rate
AP-BCFW (Ours)	$O_p\left(\frac{n\mathbb{E}_i(L_i)R^2}{\tau k}\right)$
P-BCD ⁶	$O_p\left(\frac{n\mathbb{E}_i(L_i)R^2}{\tau k}\right)$
AP-BCD ⁷	$O_p\left(\frac{n \max_i L_i R^2}{\tau k}\right)$

The comparison illustrates that these methods have the same $O(1/k)$ rate and almost the same dependence on n and τ despite the fact that we use a much simpler linear oracle. Nothing comes for free though: Nesterov acceleration does not apply for Frank-Wolfe based methods in general, while a careful implementation of parallel coordinate descents can achieve $O(1/k^2)$ rate without any full-vector interpolation in every iteration (Fercq & Richtárik, 2015). Also, Frank-Wolfe methods usually need additional restrictive conditions or algorithmic steps to get linear convergence for strongly convex problems.

These facts somewhat limits the applicability of our method to cases when projection can be computed as efficiently as (2). However, as is surveyed in (Jaggi, 2013), there are many interesting cases when (2) is much cheaper than projections, e.g., projection onto a nuclear norm ball takes $O(n^3)$ while (2) takes only $O(n^2)$.

Lastly, we note that in the fully asynchronous setting, we obtained an exponential improvement on the dependence of delay comparing to that in (Liu et al., 2014). It is unclear whether this is a unique property of the block-coordinate Frank-Wolfe algorithm or similar results can be obtained for projection based block-coordinate descent.

⁶In (Richtárik & Takáč, 2015, Theorem 19)

⁷In (Liu et al., 2014, Theorem 3)