
Parallel Successive Convex Approximation for Nonsmooth Nonconvex Optimization

Meisam Razaviyayn*
meisamr@stanford.edu

Mingyi Hong†
mingyi@iastate.edu

Zhi-Quan Luo‡
luozq@umn.edu

Jong-Shi Pang§
jongship@usc.edu

Abstract

Consider the problem of minimizing the sum of a smooth (possibly non-convex) and a convex (possibly nonsmooth) function involving a large number of variables. A popular approach to solve this problem is the block coordinate descent (BCD) method whereby at each iteration only one variable block is updated while the remaining variables are held fixed. With the recent advances in the developments of the multi-core parallel processing technology, it is desirable to parallelize the BCD method by allowing multiple blocks to be updated simultaneously at each iteration of the algorithm. In this work, we propose an inexact parallel BCD approach where at each iteration, a subset of the variables is updated in parallel by minimizing convex approximations of the original objective function. We investigate the convergence of this parallel BCD method for both randomized and cyclic variable selection rules. We analyze the asymptotic and non-asymptotic convergence behavior of the algorithm for both convex and non-convex objective functions. The numerical experiments suggest that for a special case of Lasso minimization problem, the cyclic block selection rule can outperform the randomized rule.

1 Introduction

Consider the following optimization problem

$$\min_x h(x) \triangleq f(x_1, \dots, x_n) + \sum_{i=1}^n g_i(x_i) \quad \text{s.t. } x_i \in \mathcal{X}_i, i = 1, 2, \dots, n, \quad (1)$$

where $\mathcal{X}_i \subseteq \mathbb{R}^{m_i}$ is a closed convex set; the function $f : \prod_{i=1}^n \mathcal{X}_i \rightarrow \mathbb{R}$ is a smooth function (possibly non-convex); and $g(x) \triangleq \sum_{i=1}^n g_i(x_i)$ is a separable convex function (possibly nonsmooth). The above optimization problem appears in various fields such as machine learning, signal processing, wireless communication, image processing, social networks, and bioinformatics, to name just a few. These optimization problems are typically of huge size and should be solved expeditiously.

A popular approach for solving the above multi-block optimization problem is the block coordinate descent (BCD) approach, where at each iteration of BCD, only one of the block variables is updated, while the remaining blocks are held fixed. Since only one block is updated at each iteration, the per-iteration storage and computational demand of the algorithm is low, which is desirable in huge-size problems. Furthermore, as observed in [1–3], these methods perform particularly well in practice.

*Electrical Engineering Department, Stanford University

†Industrial and Manufacturing Systems Engineering, Iowa State University

‡Department of Electrical and Computer Engineering, University of Minnesota

§Department of Industrial and Systems Engineering, University of Southern California

The availability of high performance multi-core computing platforms makes it increasingly desirable to develop parallel optimization methods. One category of such parallelizable methods is the (proximal) gradient methods. These methods are parallelizable in nature [4–8]; however, they are equivalent to successive minimization of a quadratic approximation of the objective function which may not be tight; and hence suffer from low convergence speed in some practical applications [9].

To take advantage of the BCD method and parallel multi-core technology, different parallel BCD algorithms have been recently proposed in the literature. In particular, the references [10–12] propose parallel coordinate descent minimization methods for ℓ_1 -regularized convex optimization problems. Using the greedy (Gauss-Southwell) update rule, the recent works [9, 13] propose parallel BCD type methods for general composite optimization problems. In contrast, references [2, 14–20] suggest randomized block selection rule, which is more amenable to big data optimization problems, in order to parallelize the BCD method.

Motivated by [1, 9, 15, 21], we propose a parallel inexact BCD method where at each iteration of the algorithm, a subset of the blocks is updated by minimizing locally tight approximations of the objective function. Asymptotic and non-asymptotic convergence analysis of the algorithm is presented in both convex and non-convex cases for different variable block selection rules. The proposed parallel algorithm is synchronous, which is different than the existing lock-free methods in [22, 23].

The contributions of this work are as follows:

- A parallel block coordinate descent method is proposed for non-convex nonsmooth problems. To the best of our knowledge, reference [9] is the only paper in the literature that focuses on parallelizing BCD for non-convex nonsmooth problems. This reference utilizes greedy block selection rule which requires search among all blocks as well as communication among processing nodes in order to find the best blocks to update. This requirement can be demanding in practical scenarios where the communication among nodes are costly or when the number of blocks is huge. In fact, this high computational cost motivated the authors of [9] to develop further inexact update strategies to efficiently alleviating the high computational cost of the greedy block selection rule.
- The proposed parallel BCD algorithm allows both cyclic and randomized block variable selection rules. The deterministic (*cyclic*) update rule is different than the existing parallel randomized or greedy BCD methods in the literature; see, e.g., [2, 9, 13–20]. Based on our numerical experiments, this update rule is beneficial in solving the Lasso problem.
- The proposed method not only works with the constant step-size selection rule, but also with the diminishing step-sizes which is desirable when the Lipschitz constant of the objective function is not known.
- Unlike many existing algorithms in the literature, e.g. [13–15], our parallel BCD algorithm utilizes the general approximation of the original function which includes the linear/proximal approximation of the objective as a special case. The use of general approximation instead of the linear/proximal approximation offers more flexibility and results in efficient algorithms for particular practical problems; see [21, 24] for specific examples.
- We present an iteration complexity analysis of the algorithm for both convex and non-convex scenarios. Unlike the existing non-convex parallel methods in the literature such as [9] which only guarantee the asymptotic behavior of the algorithm, we provide non-asymptotic guarantees on the convergence of the algorithm as well.

2 Parallel Successive Convex Approximation

As stated in the introduction section, a popular approach for solving (1) is the BCD method where at iteration $r + 1$ of the algorithm, the block variable x_i is updated by solving the following subproblem

$$x_i^{r+1} = \arg \min_{x_i \in \mathcal{X}_i} h(x_1^r, \dots, x_{i-1}^r, x_i, x_{i+1}^r, \dots, x_n^r). \quad (2)$$

In many practical problems, the update rule (2) is not in closed form and hence not computationally cheap. One popular approach is to replace the function $h(\cdot)$ with a well-chosen local convex

approximation $\tilde{h}_i(x_i, x^r)$ in (2). That is, at iteration $r + 1$, the block variable x_i is updated by

$$x_i^{r+1} = \arg \min_{x_i \in \mathcal{X}_i} \tilde{h}_i(x_i, x^r), \quad (3)$$

where $\tilde{h}_i(x_i, x^r)$ is a convex (possibly upper-bound) approximation of the function $h(\cdot)$ with respect to the i -th block around the current iteration x^r . This approach, also known as *block successive convex approximation* or *block successive upper-bound minimization* [21], has been widely used in different applications; see [21, 24] for more details and different useful approximation functions. In this work, we assume that the approximation function $\tilde{h}_i(\cdot, \cdot)$ is of the following form:

$$\tilde{h}_i(x_i, y) = \tilde{f}_i(x_i, y) + g_i(x_i). \quad (4)$$

Here $\tilde{f}_i(\cdot, y)$ is an approximation of the function $f(\cdot)$ around the point y with respect to the i -th block. We further assume that $\tilde{f}_i(x_i, y) : \mathcal{X}_i \times \mathcal{X} \rightarrow \mathbb{R}$ satisfies the following assumptions:

- $\tilde{f}_i(\cdot, y)$ is continuously differentiable and uniformly strongly convex with parameter τ , i.e., $\tilde{f}_i(x_i, y) \geq \tilde{f}_i(x'_i, y) + \langle \nabla_{x_i} \tilde{f}_i(x'_i, y), x_i - x'_i \rangle + \frac{\tau}{2} \|x_i - x'_i\|^2$, $\forall x_i, x'_i \in \mathcal{X}_i, \forall y \in \mathcal{X}$
- *Gradient consistency assumption*: $\nabla_{x_i} \tilde{f}_i(x_i, x) = \nabla_{x_i} f(x)$, $\forall x \in \mathcal{X}$
- $\nabla_{x_i} \tilde{f}_i(x_i, \cdot)$ is Lipschitz continuous on \mathcal{X} for all $x_i \in \mathcal{X}_i$ with constant \tilde{L} , i.e., $\|\nabla_{x_i} \tilde{f}_i(x_i, y) - \nabla_{x_i} \tilde{f}_i(x_i, z)\| \leq \tilde{L} \|y - z\|$, $\forall y, z \in \mathcal{X}, \forall x_i \in \mathcal{X}_i, \forall i$.

For instance, the following traditional proximal/quadratic approximations of $f(\cdot)$ satisfy the above assumptions when the feasible set is compact and $f(\cdot)$ is twice continuously differentiable:

- $\tilde{f}(x_i, y) = \langle \nabla_{y_i} f(y), x_i - y_i \rangle + \frac{\alpha}{2} \|x_i - y_i\|^2$.
- $\tilde{f}(x_i, y) = f(x_i, y_{-i}) + \frac{\alpha}{2} \|x_i - y_i\|^2$, for α large enough.

For other practical useful approximations of $f(\cdot)$ and the stochastic/incremental counterparts, see [21, 25, 26].

With the recent advances in the development of parallel processing machines, it is desirable to take the advantage of multi-core machines by updating multiple blocks simultaneously in (3). Unfortunately, naively updating multiple blocks simultaneously using the approach (3) does not result in a convergent algorithm. Hence, we suggest to modify the update rule by using a well-chosen step-size. More precisely, we propose Algorithm 1 for solving the optimization problem (1).

Algorithm 1 Parallel Successive Convex Approximation (PSCA) Algorithm

```

find a feasible point  $x^0 \in \mathcal{X}$  and set  $r = 0$ 
for  $r = 0, 1, 2, \dots$  do
  choose a subset  $S^r \subseteq \{1, \dots, n\}$ 
  calculate  $\hat{x}_i^r = \arg \min_{x_i \in \mathcal{X}_i} \tilde{h}_i(x_i, x^r)$ ,  $\forall i \in S^r$ 
  set  $x_i^{r+1} = x_i^r + \gamma^r (\hat{x}_i^r - x_i^r)$ ,  $\forall i \in S^r$ , and set  $x_i^{r+1} = x_i^r$ ,  $\forall i \notin S^r$ 
end for

```

The procedure of selecting the subset S^r is intentionally left unspecified in Algorithm 1. This selection could be based on different rules. Reference [9] suggests the greedy variable selection rule where at each iteration of the algorithm in [9], the best response of all the variables are calculated and at the end, only the block variables with the largest amount of improvement are updated. A drawback of this approach is the overhead caused by the calculation of all of the best responses at each iteration; this overhead is especially computationally demanding when the size of the problem is huge. In contrast to [9], we suggest the following *randomized* or *cyclic* variable selection rules:

- **Cyclic**: Given the partition $\{\mathcal{T}_0, \dots, \mathcal{T}_{m-1}\}$ of the set $\{1, 2, \dots, n\}$ with $\mathcal{T}_i \cap \mathcal{T}_j = \emptyset$, $\forall i \neq j$ and $\bigcup_{\ell=0}^{m-1} \mathcal{T}_\ell = \{1, 2, \dots, n\}$, we say the choice of the variable selection is *cyclic* if

$$S^{mr+\ell} = \mathcal{T}_\ell, \quad \forall \ell = 0, 1, \dots, m-1 \text{ and } \forall r,$$

- **Randomized:** The variable selection rule is called *randomized* if at each iteration the variables are chosen randomly from the previous iterations so that

$$Pr(j \in S^r \mid x^r, x^{r-1}, \dots, x^0) = p_j^r \geq p_{\min} > 0, \quad \forall j = 1, 2, \dots, n, \forall r.$$

3 Convergence Analysis: Asymptotic Behavior

We first make a standard assumption that $\nabla f(\cdot)$ is Lipschitz continuous with constant $L_{\nabla f}$, i.e.,

$$\|\nabla f(x) - \nabla f(y)\| \leq L_{\nabla f} \|x - y\|,$$

and assume that $-\infty < \inf_{x \in \mathcal{X}} h(x)$. Let us also define \bar{x} to be a *stationary point* of (1) if $\exists d \in \partial g(\bar{x})$ such that $\langle \nabla f(\bar{x}) + d, x - \bar{x} \rangle \geq 0, \forall x \in \mathcal{X}$, i.e., the first order optimality condition is satisfied at the point \bar{x} . The following lemma will help us to study the asymptotic convergence of the PSCA algorithm.

Lemma 1 [9, Lemma 2] *Define the mapping $\hat{x}(\cdot) : \mathcal{X} \mapsto \mathcal{X}$ as $\hat{x}(y) = (\hat{x}_i(y))_{i=1}^n$ with $\hat{x}_i(y) = \arg \min_{x_i \in \mathcal{X}_i} \tilde{h}_i(x_i, y)$. Then the mapping $\hat{x}(\cdot)$ is Lipschitz continuous with constant $\hat{L} = \frac{\sqrt{n}\tilde{L}}{\tau}$, i.e.,*

$$\|\hat{x}(y) - \hat{x}(z)\| \leq \hat{L} \|y - z\|, \quad \forall y, z \in \mathcal{X}.$$

Having derived the above result, we are now ready to state our first result which studies the limiting behavior of the PSCA algorithm. This result is based on the sufficient decrease of the objective function which has been also exploited in [9] for greedy variable selection rule.

Theorem 1 *Assume $\gamma^r \in (0, 1]$, $\sum_{r=1}^{\infty} \gamma^r = +\infty$, and that $\limsup_{r \rightarrow \infty} \gamma^r < \bar{\gamma} \triangleq \min\{\frac{\tau}{L_{\nabla f}}, \frac{\tau}{\tau + \tilde{L}\sqrt{n}}\}$. Suppose either cyclic or randomized block selection rule is employed. For cyclic update rule, assume further that $\{\gamma^r\}_{r=1}^{\infty}$ is a monotonically decreasing sequence. Then every limit point of the iterates is a stationary point of (1) – deterministically for cyclic update rule and almost surely for randomized block selection rule.*

Proof Using the standard sufficient decrease argument (see the supplementary materials), one can show that

$$h(x^{r+1}) \leq h(x^r) + \frac{\gamma^r(-\tau + \gamma^r L_{\nabla f})}{2} \|\hat{x}^r - x^r\|_{S^r}^2. \quad (5)$$

Since $\limsup_{r \rightarrow \infty} \gamma^r < \bar{\gamma}$, for sufficiently large r , there exists $\beta > 0$ such that

$$h(x^{r+1}) \leq h(x^r) - \beta \gamma^r \|\hat{x}^r - x^r\|_{S^r}^2. \quad (6)$$

Taking the conditional expectation from both sides implies

$$\mathbb{E}[h(x^{r+1}) \mid x^r] \leq h(x^r) - \beta \gamma^r \mathbb{E} \left[\sum_{i=1}^n R_i^r \|\hat{x}_i^r - x_i^r\|^2 \mid x^r \right], \quad (7)$$

where R_i^r is a Bernoulli random variable which is one if $i \in S^r$ and it is zero otherwise. Clearly, $\mathbb{E}[R_i^r \mid x^r] = p_i^r$ and therefore,

$$\mathbb{E}[h(x^{r+1}) \mid x^r] \leq h(x^r) - \beta \gamma^r p_{\min} \|\hat{x}^r - x^r\|^2, \quad \forall r. \quad (8)$$

Thus $\{h(x^r)\}$ is a supermartingale with respect to the natural history; and by the supermartingale convergence theorem [27, Proposition 4.2], $h(x^r)$ converges and we have

$$\sum_{r=1}^{\infty} \gamma^r \|\hat{x}^r - x^r\|^2 < \infty, \quad \text{almost surely.} \quad (9)$$

Let us now restrict our analysis to the set of probability one for which $h(x^r)$ converges and $\sum_{r=1}^{\infty} \gamma^r \|\hat{x}^r - x^r\|^2 < \infty$. Fix a realization in that set. The equation (9) simply implies that, for the fixed realization, $\liminf_{r \rightarrow \infty} \|\hat{x}^r - x^r\| = 0$, since $\sum_r \gamma^r = \infty$. Next we strengthen this result by proving that $\lim_{r \rightarrow \infty} \|\hat{x}^r - x^r\| = 0$. Suppose the contrary that there exists $\delta > 0$ such

that $\Delta^r \triangleq \|\widehat{x}^r - x^r\| \geq 2\delta$ infinitely often. Since $\liminf_{r \rightarrow \infty} \Delta^r = 0$, there exists a subset of indices \mathcal{K} and $\{i_r\}$ such that for any $r \in \mathcal{K}$,

$$\Delta^r < \delta, \quad 2\delta < \Delta^{i_r}, \quad \text{and} \quad \delta \leq \Delta^j \leq 2\delta, \quad \forall j = r+1, \dots, i_r-1. \quad (10)$$

Clearly,

$$\begin{aligned} \delta - \Delta^r &\stackrel{(i)}{\leq} \Delta^{r+1} - \Delta^r = \|\widehat{x}^{r+1} - x^{r+1}\| - \|\widehat{x}^r - x^r\| \stackrel{(ii)}{\leq} \|\widehat{x}^{r+1} - \widehat{x}^r\| + \|x^{r+1} - x^r\| \\ &\stackrel{(iii)}{\leq} (1 + \widehat{L})\|x^{r+1} - x^r\| \stackrel{(iv)}{=} (1 + \widehat{L})\gamma^r \|\widehat{x}^r - x^r\| \leq (1 + \widehat{L})\gamma^r \delta, \end{aligned} \quad (11)$$

where (i) and (ii) are due to (10) and the triangle inequality, respectively. The inequality (iii) is the result of Lemma 1; and (iv) is followed from the algorithm iteration update rule. Since $\limsup_{r \rightarrow \infty} \gamma^r < \frac{1}{1 + \widehat{L}}$, the above inequality implies that there exists an $\alpha > 0$ such that

$$\Delta^r > \alpha, \quad (12)$$

for all r large enough. Furthermore, since the chosen realization satisfies (9), we have that $\lim_{r \rightarrow \infty} \sum_{t=r}^{i_r-1} \gamma^t (\Delta^t)^2 = 0$; which combined with (10) and (12), implies

$$\lim_{r \rightarrow \infty} \sum_{t=r}^{i_r-1} \gamma^t = 0. \quad (13)$$

On the other hand, using the similar reasoning as in above, one can write

$$\begin{aligned} \delta < \Delta^{i_r} - \Delta^r &= \|\widehat{x}^{i_r} - x^{i_r}\| - \|\widehat{x}^r - x^r\| \leq \|\widehat{x}^{i_r} - \widehat{x}^r\| + \|x^{i_r} - x^r\| \\ &\leq (1 + \widehat{L}) \sum_{t=r}^{i_r-1} \gamma^t \|\widehat{x}^t - x^t\| \leq 2\delta(1 + \widehat{L}) \sum_{t=r}^{i_r-1} \gamma^t, \end{aligned}$$

and hence $\liminf_{r \rightarrow \infty} \sum_{t=r}^{i_r-1} \gamma^t > 0$, which contradicts (13). Therefore the contrary assumption does not hold and we must have $\lim_{r \rightarrow \infty} \|\widehat{x}^r - x^r\| = 0$, almost surely. Now consider a limit point \bar{x} with the subsequence $\{x^{r_j}\}_{j=1}^{\infty}$ converging to \bar{x} . Using the definition of \widehat{x}^{r_j} , we have $\lim_{j \rightarrow \infty} \widetilde{h}_i(\widehat{x}_i^{r_j}, x^{r_j}) \leq \widetilde{h}_i(x_i, x^{r_j})$, $\forall x_i \in \mathcal{X}_i$, $\forall i$. Therefore, by letting $j \rightarrow \infty$ and using the fact that $\lim_{r \rightarrow \infty} \|\widehat{x}^r - x^r\| = 0$, almost surely, we obtain $\widetilde{h}_i(\bar{x}_i, \bar{x}) \leq \widetilde{h}_i(x_i, \bar{x})$, $\forall x_i \in \mathcal{X}_i$, $\forall i$, almost surely; which in turn, using the gradient consistency assumption, implies

$$\langle \nabla f(\bar{x}) + d, x - \bar{x} \rangle \geq 0, \quad \forall x \in \mathcal{X}, \quad \text{almost surely,}$$

for some $d \in \partial g(\bar{x})$, which completes the proof for the randomized block selection rule.

Now consider the cyclic update rule with a limit point \bar{x} . Due to the sufficient decrease bound (6), we have $\lim_{r \rightarrow \infty} h(x^r) = h(\bar{x})$. Furthermore, by taking the summation over (6), we obtain $\sum_{r=1}^{\infty} \gamma^r \|\widehat{x}^r - x^r\|_{S^r}^2 < \infty$. Consider a fixed block i and define $\{r_k\}_{k=1}^{\infty}$ to be the subsequence of iterations that block i is updated in. Clearly, $\sum_{k=1}^{\infty} \gamma^{r_k} \|\widehat{x}_i^{r_k} - x_i^{r_k}\|^2 < \infty$ and $\sum_{k=1}^{\infty} \gamma^{r_k} = \infty$, since $\{\gamma^r\}$ is monotonically decreasing. Therefore, $\liminf_{k \rightarrow \infty} \|\widehat{x}_i^{r_k} - x_i^{r_k}\| = 0$. Repeating the above argument with some slight modifications, which are omitted due to lack of space, we can show that $\lim_{k \rightarrow \infty} \|\widehat{x}_i^{r_k} - x_i^{r_k}\| = 0$ implying that the limit point \bar{x} is a stationary point of (1). ■

Remark 1 *Theorem 1 covers both diminishing and constant step-size selection rule; or the combination of the two, i.e., decreasing the step-size until it is less than the constant $\bar{\gamma}$. It is also worth noting that the diminishing step-size rule is especially useful when the knowledge of the problem's constants L, \widehat{L} , and τ is not available.*

4 Convergence Analysis: Iteration Complexity

In this section, we present iteration complexity analysis of the algorithm for both convex and non-convex cases.

4.1 Convex Case

When the function $f(\cdot)$ is convex, the overall objective function will become convex; and as a result of Theorem 1, if a limit point exists, it is a global minimizer of (1). In this scenario, it is desirable to derive the iteration complexity bounds of the algorithm. Note that our algorithm employs linear combination of the two consecutive points at each iteration and hence it is different than the existing algorithms in [2, 14–20]. Therefore, not only in the cyclic case, but also in the randomized scenario, the iteration complexity analysis of PSCA is different than the existing results and should be investigated. Let us make the following assumptions for our iteration complexity analysis:

- The step-size is constant with $\gamma^r = \gamma < \frac{\tau}{L_{\nabla f}}, \forall r$.
- The level set $\{x \mid h(x) \leq h(x^0)\}$ is compact and the next two assumptions hold in this set.
- The nonsmooth function $g(\cdot)$ is Lipschitz continuous, i.e., $|g(x) - g(y)| \leq L_g \|x - y\|, \forall x, y \in \mathcal{X}$. This assumption is satisfied in many practical problems such as (group) Lasso.
- The gradient of the approximation function $\tilde{f}_i(\cdot, y)$ is uniformly Lipschitz with constant L_i , i.e., $\|\nabla_{x_i} \tilde{f}_i(x_i, y) - \nabla_{x'_i} \tilde{f}_i(x'_i, y)\| \leq L_i \|x_i - x'_i\|, \forall x_i, x'_i \in \mathcal{X}_i$.

Lemma 2 (Sufficient Descent) *There exists $\hat{\beta}, \tilde{\beta} > 0$, such that for all $r \geq 1$, we have*

- *For randomized rule: $\mathbb{E}[h(x^{r+1}) \mid x^r] \leq h(x^r) - \hat{\beta} \|\hat{x}^r - x^r\|^2$.*
- *For cyclic rule: $h(x^{m(r+1)}) \leq h(x^{mr}) - \tilde{\beta} \|x^{m(r+1)} - x^{mr}\|^2$.*

Proof The above result is an immediate consequence of (6) with $\hat{\beta} \triangleq \beta \gamma p_{\min}$ and $\tilde{\beta} \triangleq \frac{\beta}{\gamma}$. ■

Due to the bounded level set assumption, there must exist constants $\hat{Q}, Q, R > 0$ such that

$$\|\nabla f(x^r)\| \leq Q, \quad \|\nabla_{x_i} \tilde{f}_i(\hat{x}^r, x^r)\| \leq \hat{Q}, \quad \|x^r - x^*\| \leq R, \quad (14)$$

for all x^r . Next we use the constants Q, \hat{Q} and R to bound the cost-to-go in the algorithm.

Lemma 3 (Cost-to-go Estimate) *For all $r \geq 1$, we have*

- *For randomized rule: $(\mathbb{E}[h(x^{r+1}) \mid x^r] - h(x^*))^2 \leq 2((Q + L_g)^2 + nL^2R^2) \|\hat{x}^r - x^r\|^2$*
- *For cyclic rule: $(h(x^{m(r+1)}) - h(x^*))^2 \leq 3n \frac{\theta(1-\gamma)^2}{\gamma^2} \|x^{m(r+1)} - x^{mr}\|^2$*

for any optimal point x^* , where $L \triangleq \max_i \{L_i\}$ and $\theta \triangleq L_g^2 + \hat{Q}^2 + 2nR^2\tilde{L}^2 \frac{\gamma^2}{(1-\gamma)^2} + 2R^2L^2$.

Proof Please see the supplementary materials for the proof.

Lemma 2 and Lemma 3 lead to the iteration complexity bound in the following theorem. The proof steps of this result are similar to the ones in [28] and therefore omitted here for space reasons.

Theorem 2 *Define $\sigma \triangleq \frac{(\gamma L_{\nabla f} - \tau)\gamma p_{\min}}{4((Q + L_g)^2 + nL^2R^2)}$ and $\tilde{\sigma} \triangleq \frac{(\gamma L_{\nabla f} - \tau)\gamma}{6n\theta(1-\gamma)^2}$. Then*

- *For randomized update rule: $\mathbb{E}[h(x^r)] - h(x^*) \leq \frac{\max\{4\sigma - 2, h(x^0) - h(x^*), 2\}}{\sigma} \frac{1}{r}$.*
- *For cyclic update rule: $h(x^{mr}) - h(x^*) \leq \frac{\max\{4\tilde{\sigma} - 2, h(x^0) - h(x^*), 2\}}{\tilde{\sigma}} \frac{1}{r}$.*

4.2 Non-convex Case

In this subsection we study the iteration complexity of the proposed randomized algorithm for the general nonconvex function $f(\cdot)$ assuming constant step-size selection rule. This analysis is only for the randomized block selection rule. Since in the nonconvex scenario, the iterates may not converge to the global optimum point, the closeness to the optimal solution cannot be considered for the iteration complexity analysis. Instead, inspired by [29] where the size of the gradient of the objective function is used as a measure of optimality, we consider the size of the objective proximal gradient as a measure of optimality. More precisely, we define

$$\tilde{\nabla}h(x) = x - \arg \min_{y \in \mathcal{X}} \left\{ \langle \nabla f(x), y - x \rangle + g(y) + \frac{1}{2} \|y - x\|^2 \right\}.$$

Clearly, $\tilde{\nabla}h(x) = 0$ when x is a stationary point. Moreover, $\tilde{\nabla}h(\cdot)$ coincides with the gradient of the objective if $g \equiv 0$ and $\mathcal{X} = \mathbb{R}^n$. The following theorem, which studies the decrease rate of $\|\tilde{\nabla}h(x)\|$, could be viewed as an iteration complexity analysis of the randomized PSCA.

Theorem 3 Consider randomized block selection rule. Define T_ϵ to be the first time that $\mathbb{E}[\|\tilde{\nabla}h(x^r)\|^2] \leq \epsilon$. Then $T_\epsilon \leq \kappa/\epsilon$ where $\kappa \triangleq \frac{2(L^2+2L+2)(h(x^0)-h^*)}{\hat{\beta}}$ and $h^* = \min_{x \in \mathcal{X}} h(x)$.

Proof To simplify the presentation of the proof, let us define $\tilde{y}_i^r \triangleq \arg \min_{y_i \in \mathcal{X}_i} \langle \nabla_{x_i} f(x^r), y_i - x_i^r \rangle + g_i(y_i) + \frac{1}{2} \|y_i - x_i^r\|^2$. Clearly, $\tilde{\nabla}h(x^r) = (x_i^r - \tilde{y}_i^r)_{i=1}^n$. The first order optimality condition of the above optimization problem implies

$$\langle \nabla_{x_i} f(x^r) + \tilde{y}_i^r - x_i^r, x_i - \tilde{y}_i^r \rangle + g_i(x_i) - g_i(\tilde{y}_i^r) \geq 0, \quad \forall x_i \in \mathcal{X}_i. \quad (15)$$

Furthermore, based on the definition of \hat{x}_i^r , we have

$$\langle \nabla_{x_i} \tilde{f}_i(\hat{x}_i^r, x^r), x_i - \hat{x}_i^r \rangle + g_i(x_i) - g_i(\hat{x}_i^r) \geq 0, \quad \forall x_i \in \mathcal{X}_i. \quad (16)$$

Plugging in the points \hat{x}_i^r and \tilde{y}_i^r in (15) and (16); and summing up the two equations will yield to

$$\langle \nabla_{x_i} \tilde{f}_i(\hat{x}_i^r, x^r) - \nabla_{x_i} f(x^r) + x_i^r - \tilde{y}_i^r, \tilde{y}_i^r - \hat{x}_i^r \rangle \geq 0.$$

Using the gradient consistency assumption, we can write

$$\langle \nabla_{x_i} \tilde{f}_i(\hat{x}_i^r, x^r) - \nabla_{x_i} \tilde{f}_i(x_i^r, x^r) + x_i^r - \hat{x}_i^r + \hat{x}_i^r - \tilde{y}_i^r, \tilde{y}_i^r - \hat{x}_i^r \rangle \geq 0,$$

or equivalently, $\langle \nabla_{x_i} \tilde{f}_i(\hat{x}_i^r, x^r) - \nabla_{x_i} \tilde{f}_i(x_i^r, x^r) + x_i^r - \hat{x}_i^r, \tilde{y}_i^r - \hat{x}_i^r \rangle \geq \|\hat{x}_i^r - \tilde{y}_i^r\|^2$. Applying Cauchy-Schwarz and the triangle inequality will yield to

$$\left(\|\nabla_{x_i} \tilde{f}_i(\hat{x}_i^r, x^r) - \nabla_{x_i} \tilde{f}_i(x_i^r, x^r)\| + \|x_i^r - \hat{x}_i^r\| \right) \|\tilde{y}_i^r - \hat{x}_i^r\| \geq \|\hat{x}_i^r - \tilde{y}_i^r\|^2.$$

Since the function $\tilde{f}_i(\cdot, x)$ is Lipschitz, we must have

$$\|\hat{x}_i^r - \tilde{y}_i^r\| \leq (1 + L_i) \|x_i^r - \hat{x}_i^r\| \quad (17)$$

Using the inequality (17), the norm of the proximal gradient of the objective can be bounded by

$$\begin{aligned} \|\tilde{\nabla}h(x^r)\|^2 &= \sum_{i=1}^n \|x_i^r - \tilde{y}_i^r\|^2 \leq 2 \sum_{i=1}^n (\|x_i^r - \hat{x}_i^r\|^2 + \|\hat{x}_i^r - \tilde{y}_i^r\|^2) \\ &\leq 2 \sum_{i=1}^n (\|x_i^r - \hat{x}_i^r\|^2 + (1 + L_i)^2 \|x_i^r - \hat{x}_i^r\|^2) \leq 2(2 + 2L + L^2) \|\hat{x}^r - x^r\|^2. \end{aligned}$$

Combining the above inequality with the sufficient decrease bound in (7), one can write

$$\begin{aligned} \sum_{r=0}^T \mathbb{E} \left[\|\tilde{\nabla}h(x^r)\|^2 \right] &\leq \sum_{r=1}^T 2(2 + 2L + L^2) \mathbb{E} \left[\|\hat{x}^r - x^r\|^2 \right] \\ &\leq \sum_{r=0}^T \frac{2(2 + 2L + L^2)}{\hat{\beta}} \mathbb{E} [h(x^r) - h(x^{r+1})] \leq \frac{2(2 + 2L + L^2)}{\hat{\beta}} \mathbb{E} [h(x^0) - h(x^{T+1})] \\ &\leq \frac{2(2 + 2L + L^2)}{\hat{\beta}} [h(x^0) - h^*] = \kappa, \end{aligned}$$

which implies that $T_\epsilon \leq \frac{\kappa}{\epsilon}$. ■

5 Numerical Experiments:

In this short section, we compare the numerical performance of the proposed algorithm with the classical serial BCD methods. The algorithms are evaluated over the following Lasso problem:

$$\min_x \frac{1}{2} \|Ax - b\|_2^2 + \lambda \|x\|_1,$$

where the matrix A is generated according to the Nesterov’s approach [5]. Two problem instances are considered: $A \in \mathbb{R}^{2000 \times 10,000}$ with 1% sparsity level in x^* and $A \in \mathbb{R}^{1000 \times 100,000}$ with 0.1% sparsity level in x^* . The approximation functions are chosen similar to the numerical experiments in [9], i.e., block size is set to one ($m_i = 1, \forall i$) and the approximation function

$$\tilde{f}(x_i, y) = f(x_i, y_{-i}) + \frac{\alpha}{2} \|x_i - y_i\|^2$$

is considered, where $f(x) = \frac{1}{2} \|Ax - b\|_2^2$ is the smooth part of the objective function. We choose constant step-size γ and proximal coefficient α . In general, careful selection of the algorithm parameters results in better numerical convergence rate. The smaller values of step-size γ will result in less zigzag behavior for the convergence path of the algorithm; however, too small step sizes will clearly slow down the convergence speed. Furthermore, in order to make the approximation function sufficiently strongly convex, we need to choose α large enough. However, choosing too large α values enforces the next iterates to stay close to the current iterate and results in slower convergence speed; see the supplementary materials for related examples.

Figure 1 and Figure 2 illustrate the behavior of cyclic and randomized parallel BCD method as compared with their serial counterparts. The serial methods “Cyclic BCD” and “Randomized BCD” are based on the update rule in (2) with the cyclic and randomized block selection rules, respectively. The variable q shows the number of processors and on each processor we update 40 scalar variables in parallel. As can be seen in Figure 1 and Figure 2, parallelization of the BCD algorithm results in more efficient algorithm. However, the computational gain does not grow linearly with the number of processors. In fact, we can see that after some point, the increase in the number of processors lead to slower convergence. This fact is due to the communication overhead among the processing nodes which dominates the computation time; see the supplementary materials for more numerical experiments on this issue.

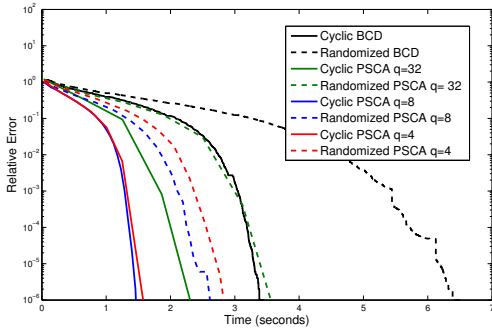


Figure 1: Lasso Problem: $A \in \mathbb{R}^{2,000 \times 10,000}$

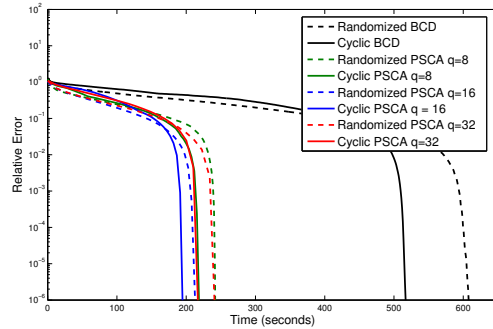


Figure 2: Lasso Problem: $A \in \mathbb{R}^{1,000 \times 100,000}$

Acknowledgments: The authors are grateful to the University of Minnesota Graduate School Doctoral Dissertation Fellowship and AFOSR, grant number FA9550-12-1-0340 for the support during this research.

References

- [1] Y. Nesterov. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization*, 22(2):341–362, 2012.
- [2] P. Richtárik and M. Takáč. Efficient serial and parallel coordinate descent methods for huge-scale truss topology design. In *Operations Research Proceedings*, pages 27–32. Springer, 2012.

- [3] Y. T. Lee and A. Sidford. Efficient accelerated coordinate descent methods and faster algorithms for solving linear systems. In *54th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 147–156. IEEE, 2013.
- [4] I. Necoara and D. Clipici. Efficient parallel coordinate descent algorithm for convex optimization problems with separable constraints: application to distributed MPC. *Journal of Process Control*, 23(3):243–253, 2013.
- [5] Y. Nesterov. Gradient methods for minimizing composite functions. *Mathematical Programming*, 140(1):125–161, 2013.
- [6] P. Tseng and S. Yun. A coordinate gradient descent method for nonsmooth separable minimization. *Mathematical Programming*, 117(1-2):387–423, 2009.
- [7] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.
- [8] S. J. Wright, R. D. Nowak, and M. Figueiredo. Sparse reconstruction by separable approximation. *IEEE Transactions on Signal Processing*, 57(7):2479–2493, 2009.
- [9] F. Facchinei, S. Sagratella, and G. Scutari. Flexible parallel algorithms for big data optimization. *arXiv preprint arXiv:1311.2444*, 2013.
- [10] J. K. Bradley, A. Kyröla, D. Bickson, and C. Guestrin. Parallel coordinate descent for ℓ_1 -regularized loss minimization. *arXiv preprint arXiv:1105.5379*, 2011.
- [11] C. Scherrer, A. Tewari, M. Halappanavar, and D. Haglin. Feature clustering for accelerating parallel coordinate descent. In *NIPS*, pages 28–36, 2012.
- [12] C. Scherrer, M. Halappanavar, A. Tewari, and D. Haglin. Scaling up coordinate descent algorithms for large ℓ_1 regularization problems. *arXiv preprint arXiv:1206.6409*, 2012.
- [13] Z. Peng, M. Yan, and W. Yin. Parallel and distributed sparse optimization. *preprint*, 2013.
- [14] I. Necoara and D. Clipici. Distributed coordinate descent methods for composite minimization. *arXiv preprint arXiv:1312.5302*, 2013.
- [15] P. Richtárik and M. Takáč. Parallel coordinate descent methods for big data optimization. *arXiv preprint arXiv:1212.0873*, 2012.
- [16] P. Richtárik and M. Takáč. On optimal probabilities in stochastic coordinate descent methods. *arXiv preprint arXiv:1310.3438*, 2013.
- [17] O. Fercoq and P. Richtárik. Accelerated, parallel and proximal coordinate descent. *arXiv preprint arXiv:1312.5799*, 2013.
- [18] O. Fercoq, Z. Qu, P. Richtárik, and M. Takáč. Fast distributed coordinate descent for non-strongly convex losses. *arXiv preprint arXiv:1405.5300*, 2014.
- [19] O. Fercoq and P. Richtárik. Smooth minimization of nonsmooth functions with parallel coordinate descent methods. *arXiv preprint arXiv:1309.5885*, 2013.
- [20] A. Patrascu and I. Necoara. A random coordinate descent algorithm for large-scale sparse nonconvex optimization. In *European Control Conference (ECC)*, pages 2789–2794. IEEE, 2013.
- [21] M. Razaviyayn, M. Hong, and Z.-Q. Luo. A unified convergence analysis of block successive minimization methods for nonsmooth optimization. *SIAM Journal on Optimization*, 23(2):1126–1153, 2013.
- [22] F. Niu, B. Recht, C. Ré, and S. J. Wright. Hogwild!: A lock-free approach to parallelizing stochastic gradient descent. *Advances in Neural Information Processing Systems*, 24:693–701, 2011.
- [23] J. Liu, S. J. Wright, C. Ré, and V. Bittorf. An asynchronous parallel stochastic coordinate descent algorithm. *arXiv preprint arXiv:1311.1873*, 2013.
- [24] J. Mairal. Optimization with first-order surrogate functions. *arXiv preprint arXiv:1305.3120*, 2013.
- [25] J. Mairal. Incremental majorization-minimization optimization with application to large-scale machine learning. *arXiv preprint arXiv:1402.4419*, 2014.
- [26] M. Razaviyayn, M. Sanjabi, and Z.-Q. Luo. A stochastic successive minimization method for nonsmooth nonconvex optimization with applications to transceiver design in wireless communication networks. *arXiv preprint arXiv:1307.4457*, 2013.
- [27] D. P. Bertsekas and J. N. Tsitsiklis. *Neuro-dynamic programming*. 1996.
- [28] M. Hong, X. Wang, M. Razaviyayn, and Z.-Q. Luo. Iteration complexity analysis of block coordinate descent methods. *arXiv preprint arXiv:1310.6957*, 2013.
- [29] Y. Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer, 2004.