

Analysis and Implementation of an Asynchronous Optimization Algorithm for the Parameter Server

Arda Aytekin *Student Member, IEEE*, Hamid Reza Feyzmahdavian, *Student Member, IEEE*,
and Mikael Johansson *Member, IEEE*

Abstract

This paper presents an asynchronous incremental aggregated gradient algorithm and its implementation in a parameter server framework for solving regularized optimization problems. The algorithm can handle both general convex (possibly non-smooth) regularizers and general convex constraints. When the empirical data loss is strongly convex, we establish linear convergence rate, give explicit expressions for step-size choices that guarantee convergence to the optimum, and bound the associated convergence factors. The expressions have an explicit dependence on the degree of asynchrony and recover classical results under synchronous operation. Simulations and implementations on commercial compute clouds validate our findings.

Index Terms

asynchronous, proximal, incremental, aggregated gradient, linear convergence.

I. INTRODUCTION

MACHINE learning and optimization theory have enjoyed a fruitful symbiosis over the last decade. On the one hand, since many machine learning tasks can be posed as optimization problems, advances in large-scale optimization (e.g. [1]) have had an immediate and profound impact on machine learning research. On the other hand, the challenges of dealing with huge data sets, often spread over multiple sites, have inspired the machine learning community to develop novel optimization algorithms [2], improve the theory for asynchronous computations [3], and introduce new programming models for parallel and distributed optimization [4].

In this paper, we consider machine learning in the parameter server framework [4]. This is a master-worker architecture, where a central server maintains the current parameter iterates and queries worker nodes for gradients of the loss evaluated on their data. In this setting, we focus on problems on the form

$$\begin{aligned} & \underset{x}{\text{minimize}} && \sum_{n=1}^N f_n(x) + h(x) \\ & \text{subject to} && x \in \mathbb{R}^d. \end{aligned}$$

Here, the first part of the objective function typically models the empirical data loss and the second term is a regularizer (for example, an ℓ_1 penalty to promote sparsity of the solution). Regularized optimization problems arise in many applications in machine learning, signal processing, and statistical estimation. Examples include Tikhonov and elastic net regularization, Lasso, sparse logistic regression, and support vector machines.

In the parameter server framework, Li *et al.* [4] analyzed a parallel and asynchronous proximal gradient method for non-convex problems and established conditions for convergence to a critical point. Agarwal and Duchi [5], and more recently Feyzmahdavian *et al.* [6], developed parallel mini-batch optimization algorithms based on asynchronous incremental gradient methods. When the loss functions are strongly convex, which is often the case, it has recently been observed that incremental aggregated methods outperform incremental gradient descent and are, in addition, able to converge to the true optimum even with a constant step-size. Gurbuzbalaban *et al.* [7] established linear convergence for an incremental aggregated gradient method suitable for implementation in the parameter server framework. However, the analysis does not allow for any regularization term, nor any additional convex constraints.

This paper presents an asynchronous proximal incremental aggregated gradient algorithm and its implementation in the parameter server framework. Our algorithm can handle both general convex regularizers and convex constraints. We establish linear convergence when the empirical data loss is strongly convex, give explicit expressions for step-size choices that guarantee convergence to the global optimum and bound the associated convergence factors. These expressions have an explicit dependence on the degree of asynchrony and recover classical results under synchronous operation. We believe that this is a practically and theoretically important addition to existing optimization algorithms for the parameter server architecture.

A. Prior work

Incremental gradient methods for smooth optimization problems have a long tradition, most notably in the training of neural networks via back-propagation. In contrast to gradient methods, which compute the full gradient of the loss function before updating the iterate, incremental gradient methods evaluate the gradients of a single, or possibly a few, component functions in each iteration. Incremental gradient methods can be computationally more efficient than traditional gradient methods since each step is cheaper but makes a comparable progress on average. However, for global convergence, the step-size needs to diminish to zero, which can lead to slow convergence [8]. If a constant step-size is used, only convergence to an approximate solution can be guaranteed in general [9].

Recently, Blatt, Hero, and Gauchman [10] proposed a method, the *incremental aggregated gradient* (IAG), that also computes the gradient of a single component function at each iteration. But rather than updating the iterate based on this information, it uses the sum of the most recently evaluated gradients of all component functions. Compared to the basic incremental gradient methods, IAG has the advantage that global convergence can be achieved using a constant step-size when each component function is convex quadratic. Later, Gurbuzbalaban, Ozdaglar, and Parillo [7] proved linear convergence for IAG in a more general setting when component functions are strongly convex. In a more recent work, Vanli, Gurbuzbalaban and Ozdaglar [11] analyzed the global convergence rate of proximal incremental aggregated gradient methods, where they can provide the linear convergence rate only after sufficiently many iterations. Our result differs from theirs in that we provide the linear convergence rate of the algorithm without any constraints on the iteration count and we extend the result to the general distance functions.

There has been some recent work on the stochastic version of the IAG method (called stochastic average gradient, or SAG) where we sample the component function to update instead of using a cyclic order [12]–[14]. Unlike the IAG method where the linear convergence rate depends on the number of passes through the data, the SAG method achieves a linear convergence rate that depends on the number of iterations. Further, when the number of training examples is sufficiently large, the SAG method allows the use of a very large step-size, which leads to improved theoretical and empirical performance.

II. NOTATION

We let \mathbb{N} and \mathbb{N}_0 denote the set of natural numbers and the set of natural numbers including zero, respectively. The inner product of two vectors $x, y \in \mathbb{R}^d$ is denoted by $\langle x, y \rangle$. We assume that \mathbb{R}^d is endowed with a norm $\|\cdot\|$ and use $\|\cdot\|_*$ to represent the corresponding dual norm, defined by

$$\|y\|_* = \sup_{\|x\| \leq 1} \langle x, y \rangle .$$

III. PROBLEM DEFINITION

We consider optimization problems on the form

$$\begin{aligned} & \underset{x}{\text{minimize}} && \sum_{n=1}^N f_n(x) + h(x) \\ & \text{subject to} && x \in \mathbb{R}^d, \end{aligned} \tag{1}$$

where x is the decision variable, $f_n(x)$ is convex and differentiable for each $n \in \mathcal{N} := \{1, \dots, N\}$ and $h(x)$ is a proper convex function that may be non-smooth and extended real-valued. The role of the regularization term $h(x)$ is to favor solutions with certain preferred structure. For example, $h(x) = \lambda_1 \|x\|_1$ with $\lambda_1 > 0$ is often used to promote sparsity in solutions, and

$$h(x) = I_{\mathcal{X}}(x) := \begin{cases} 0 & \text{if } x \in \mathcal{X} \subseteq \mathbb{R}^d, \\ +\infty & \text{otherwise} \end{cases}$$

is used to force the possible solutions to lie in the closed convex set \mathcal{X} .

In order to solve (1), we are going to use the proximal incremental aggregated gradient method. In this method, at iteration $k \in \mathbb{N}$, the gradients of all component functions $f_n(x)$, possibly evaluated at stale information $x_{k-\tau_k^n}$, are aggregated

$$g_k = \sum_{n=1}^N \nabla f_n(x_{k-\tau_k^n}) .$$

Then, a proximal step is taken based on the current vector x_k , the aggregated gradient g_k , and the non-smooth term $h(x)$,

$$x_{k+1} = \arg \min_x \left\{ \langle g_k, x - x_k \rangle + \frac{1}{2\alpha} \|x - x_k\|^2 + h(x) \right\}. \tag{2}$$

The algorithm has a natural implementation in the parameter server framework. The master node maintains the iterate x and performs the proximal steps. Whenever a worker node reports new gradients, the master updates the iterate and informs the worker about the new iterate. Pseudo code for a basic parameter server implementation is given in Algorithms 1 and 2.

Algorithm 1 Master procedure

```

1: Data:  $g_w$  for each worker  $w \in \mathcal{W} := \{1, 2, \dots, W\}$ 
2: Input:  $\alpha(L, \mu, \bar{\tau})$ ,  $K > 0$  and  $h(x)$ 
3: Output:  $x_K$ 
4: Initialize:  $k = 0$ 

5: while  $k < K$  do
6:   Wait until a set  $\mathcal{R}$  of workers return their gradients
7:   for all  $w \in \mathcal{W}$  do
8:     if  $w \in \mathcal{R}$  then
9:       Update  $g_w \leftarrow \sum_{n \in \mathcal{N}_w} \nabla f_n(x_{k-\tau_k^w})$ 
10:    else
11:      Keep old  $g_w$ 
12:    end if
13:  end for
14:  Aggregate the incremental gradients  $g_k = \sum_{w \in \mathcal{W}} g_w$ 
15:  Solve (2) with  $g_k$ 
16:  for all  $w \in \mathcal{R}$  do
17:    Send  $x_{k+1}$  to worker  $w$ 
18:  end for
19:  Increment  $k$ 
20: end while
21: Signal EXIT
22: Return  $x_K$ 

```

Algorithm 2 Procedure for each worker w

```

1: Data:  $x$ , and loss functions  $\{f_w(x) : w \in \mathcal{N}_w\}$  with  $\bigcup_{w \in \mathcal{W}} \mathcal{N}_w = \mathcal{N}$  and  $\mathcal{N}_{w_1} \cap \mathcal{N}_{w_2} = \emptyset \forall w_1 \neq w_2 \in \mathcal{W}$ 

2: repeat
3:   Receive  $x \leftarrow x_{k+1}$  from master
4:   Calculate incremental gradient (IG)  $\sum_{n \in \mathcal{N}_w} \nabla f_n(x)$ 
5:   Send IG to master with a delay of  $\tau_k^w$ 
6: until EXIT received

```

To establish convergence of the iterates to the global optimum, we impose the following assumptions on Problem (1):

A1) The function $F(x) := \sum_{n=1}^N f_n(x)$ is μ -strongly convex, *i.e.*,

$$F(x) \geq F(y) + \langle \nabla F(y), x - y \rangle + \frac{\mu}{2} \|x - y\|^2, \quad (3)$$

holds for all $x, y \in \mathbb{R}^d$.

A2) Each f_n is convex with L_n -continuous gradient, that is,

$$\|\nabla f_n(x) - \nabla f_n(y)\|_* \leq L_n \|x - y\| \quad \forall x, y \in \mathbb{R}^d.$$

Note that under this assumption, ∇F is also Lipschitz continuous with $L \leq \sum_{n=1}^N L_n$.

A3) $h(x)$ is sub-differentiable everywhere in its effective domain, that is, for all $x, y \in \{z \in \mathbb{R}^d : h(z) < +\infty\}$,

$$h(x) \geq h(y) + \langle s(y), x - y \rangle \quad \forall s(y) \in \partial h(y). \quad (4)$$

A4) The time-varying delays τ_k^n are bounded, *i.e.*, there is a non-negative integer $\bar{\tau}$ such that

$$\tau_k^n \in \{0, 1, \dots, \bar{\tau}\},$$

hold for all $k \in \mathbb{N}_0$ and $n \in \mathcal{N}$.

IV. MAIN RESULT

First, we provide a lemma which is key to proving our main result.

Lemma 1. Assume that the non-negative sequences $\{V_k\}$ and $\{w_k\}$ satisfy the following inequality:

$$V_{k+1} \leq aV_k - bw_k + c \sum_{j=k-k_0}^k w_j, \quad (5)$$

for some real numbers $a \in (0, 1)$ and $b, c \geq 0$, and some integer $k_0 \in \mathbb{N}_0$. Assume also that $w_k = 0$ for $k < 0$, and that the following holds:

$$\frac{c}{1-a} \frac{1-a^{k_0+1}}{a^{k_0}} \leq b.$$

Then, $V_k \leq a^k V_0$ for all $k \geq 0$.

Proof: To prove the linear convergence of the sequence, we divide both sides of (5) by a^{k+1} and take the sum:

$$\begin{aligned} \sum_{k=0}^K \frac{V_{k+1}}{a^{k+1}} &\leq \sum_{k=0}^K \frac{V_k}{a^k} - b \sum_{k=0}^K \frac{w_k}{a^{k+1}} + c \sum_{k=0}^K \frac{1}{a^{k+1}} \sum_{j=k-k_0}^k w_j \\ &= \sum_{k=0}^K \frac{V_k}{a^k} - b \sum_{k=0}^K \frac{w_k}{a^{k+1}} \\ &\quad + \frac{c}{a} (w_{-k_0} + w_{-k_0+1} + \dots + w_0) \\ &\quad + \frac{c}{a^2} (w_{-k_0+1} + w_{-k_0+2} + \dots + w_1) + \dots \\ &\quad + \frac{c}{a^{K+1}} (w_{K-k_0} + w_{K-k_0+1} + \dots + w_K) \\ &\leq \left(c \left(1 + \frac{1}{a} + \dots + \frac{1}{a^{k_0}} \right) - b \right) \sum_{k=0}^K \frac{w_k}{a^{k+1}} \\ &\quad + \sum_{k=0}^K \frac{V_k}{a^k}, \end{aligned} \quad (6)$$

where we have used the non-negativity of w_k to obtain (6).

If the coefficient of the first sum of the right-hand side of (6) is non-positive, i.e., if

$$c + \frac{c}{a} + \dots + \frac{c}{a^{k_0}} = \frac{c}{1-a} \frac{1-a^{k_0+1}}{a^{k_0}} \leq b,$$

then inequality (6) implies that

$$\frac{V_{K+1}}{a^{K+1}} + \frac{V_K}{a^K} + \dots + \frac{V_1}{a^1} \leq \frac{V_K}{a^K} + \frac{V_{K-1}}{a^{K-1}} + \dots + \frac{V_0}{a^0}.$$

Hence, $V_{K+1} \leq a^{K+1} V_0$ for any $K \geq 0$ and the desired result follows. ■

We are now ready to state and prove our main result.

Theorem 1. Assume that Problem (1) satisfies assumptions A1–A4, and that the step-size α satisfies:

$$\alpha \leq \frac{\left(1 + \frac{\mu}{L} \frac{1}{\bar{\tau}+1} \right)^{\frac{1}{\bar{\tau}+1}} - 1}{\mu},$$

where $L = \sum_{n=1}^N L_n$. Then, the iterates generated by Algorithms 1 and 2 satisfy:

$$\|x_k - x^*\|^2 \leq \left(\frac{1}{\mu\alpha + 1} \right)^k \|x_0 - x^*\|^2.$$

for all $k \geq 0$.

Proof: We start with analyzing each component function $f_n(x)$ to find upper bounds on the function values:

$$\begin{aligned} f_n(x_{k+1}) &\leq f_n(x_{k-\tau_k^n}) + \langle \nabla f_n(x_{k-\tau_k^n}), x_{k+1} - x_{k-\tau_k^n} \rangle \\ &\quad + \frac{L_n}{2} \|x_{k+1} - x_{k-\tau_k^n}\|^2 \\ &\leq f_n(x) + \langle \nabla f_n(x_{k-\tau_k^n}), x_{k+1} - x \rangle \\ &\quad + \frac{L_n}{2} \|x_{k+1} - x_{k-\tau_k^n}\|^2 \quad \forall x, \end{aligned} \quad (7)$$

where the first and second inequalities use L_n -continuity and convexity of $f_n(x)$, respectively. Summing (7) over all component functions, we obtain:

$$\begin{aligned} F(x_{k+1}) &\leq F(x) + \langle g_k, x_{k+1} - x \rangle \\ &\quad + \sum_{n=1}^N \frac{L_n}{2} \|x_{k+1} - x_{k-\tau_k^n}\|^2 \quad \forall x. \end{aligned} \quad (8)$$

Observe that optimality condition of (2) implies:

$$\begin{aligned} \langle g_k, x_{k+1} - x \rangle &\leq \frac{1}{\alpha} \langle x_{k+1} - x_k, x - x_{k+1} \rangle \\ &\quad + \langle s(x_{k+1}), x - x_{k+1} \rangle \quad \forall x \in \mathcal{X}. \end{aligned} \quad (9)$$

To find an upper bound on the second term of the right-hand side of (8), we use the three-point equality on (9) to obtain:

$$\begin{aligned} \langle g_k, x_{k+1} - x \rangle &\leq \frac{1}{2\alpha} \|x_k - x\|^2 - \frac{1}{2\alpha} \|x_{k+1} - x_k\|^2 \\ &\quad - \frac{1}{2\alpha} \|x_{k+1} - x\|^2 \\ &\quad + \langle s(x_{k+1}), x - x_{k+1} \rangle \quad \forall x \in \mathcal{X}. \end{aligned} \quad (10)$$

Plugging $y = x_{k+1}$ in (4), and using (4) together with (10) in (8), we obtain the following relation:

$$\begin{aligned} F(x_{k+1}) + h(x_{k+1}) + \frac{1}{2\alpha} \|x_{k+1} - x\|^2 &\leq F(x) + h(x) \\ &\quad + \frac{1}{2\alpha} \|x_k - x\|^2 - \frac{1}{2\alpha} \|x_{k+1} - x_k\|^2 \\ &\quad + \sum_{n=1}^N \frac{L_n}{2} \|x_{k+1} - x_{k-\tau_k^n}\|^2 \quad \forall x \in \mathcal{X}. \end{aligned}$$

Using the strong convexity property (3) on $F(x_{k+1}) + h(x_{k+1})$ above and choosing $x = x^*$ gives:

$$\begin{aligned} &\langle \nabla F(x^*) + s(x^*), x_{k+1} - x^* \rangle + \frac{\mu}{2} \|x_{k+1} - x^*\|^2 \\ &\quad + \frac{1}{2\alpha} \|x_{k+1} - x^*\|^2 \\ &\leq \frac{1}{2\alpha} \|x_k - x^*\|^2 - \frac{1}{2\alpha} \|x_{k+1} - x_k\|^2 \\ &\quad + \sum_{n=1}^N \frac{L_n}{2} \|x_{k+1} - x_{k-\tau_k^n}\|^2. \end{aligned} \quad (11)$$

Due to the optimality condition of (1), there exists a subgradient $s(x^*)$ such that the first term on the left-hand side is non-negative. Using this particular subgradient, we drop the first term. The last term on the right-hand side of the inequality can be further upper-bounded using Jensen's inequality as follows:

$$\begin{aligned} \sum_{n=1}^N \frac{L_n}{2} \|x_{k+1} - x_{k-\tau_k^n}\|^2 &= \sum_{n=1}^N \frac{L_n}{2} \left\| \sum_{j=k-\tau_k^n}^k x_{j+1} - x_j \right\|^2 \\ &\leq \frac{L(\bar{\tau}+1)}{2} \sum_{j=k-\bar{\tau}}^k \|x_{j+1} - x_j\|^2, \end{aligned}$$

where $L = \sum_{n=1}^N L_n$. As a result, rearranging the terms in (11), we obtain:

$$\begin{aligned} \|x_{k+1} - x^*\|^2 &\leq \frac{1}{\mu\alpha + 1} \|x_k - x^*\|^2 \\ &\quad - \frac{1}{\mu\alpha + 1} \|x_{k+1} - x_k\|^2 \\ &\quad + \frac{\alpha(\bar{\tau}+1)L}{\mu\alpha + 1} \sum_{j=k-\bar{\tau}}^k \|x_{j+1} - x_j\|^2. \end{aligned}$$

We note that $\|x_{j+1} - x_j\|^2 = 0$ for all $j < 0$. Using Lemma 1 with $V_k = \|x_{k+1} - x^*\|^2$, $w_k = \|x_{k+1} - x_k\|^2$, $a = b = \frac{1}{\mu\alpha+1}$, $c = \frac{\alpha(\bar{\tau}+1)L}{\mu\alpha+1}$ and $k_0 = \bar{\tau}$ completes the proof. \blacksquare

Remark 1. For the special case of Algorithms 1 and 2 where $\tau_k^n = 0$ for all k, n , Xiao and Zhang [15] have shown that the convergence rate of serial proximal gradient method with a constant step-size $\alpha = \frac{1}{L}$ is

$$\mathcal{O}\left(\left(\frac{L - \mu_F}{L + \mu_h}\right)^k\right)$$

where μ_F and μ_h are strong convexity parameters of $F(x)$ and $h(x)$, respectively. It is clear that in the case that $\bar{\tau} = 0$, the guaranteed bound in Theorem 1 reduces to the one obtained in [15].

V. PROXIMAL INCREMENTAL AGGREGATE DESCENT WITH GENERAL DISTANCE FUNCTIONS

The update rule of our algorithm can be easily extended to a non-Euclidean setting, by replacing the Euclidean squared distance in (2) with a general Bregman distance function. We first define a Bregman distance function, also referred to as a prox-function.

Definition 1. A function $\omega: \mathbb{R}^d \rightarrow \mathbb{R}$ is called a distance generating function with modulus $\mu_\omega > 0$ if ω is continuously differentiable and μ_ω -strongly convex with respect to $\|\cdot\|$. Every distance generating function introduces a corresponding Bregman distance function given by

$$D_\omega(x, x') := \omega(x') - \omega(x) - \langle \nabla \omega(x), x' - x \rangle .$$

For example, if we choose $\omega(x) = \frac{1}{2} \|x\|_2^2$, which is 1-strongly convex with respect to the l_2 -norm, that would result in $D_\omega(x, x') = \frac{1}{2} \|x' - x\|_2^2$. Another common example of distance generating functions is the entropy function

$$\omega(x) = \sum_{i=1}^d x_i \log(x_i) ,$$

which is 1-strongly convex with respect to the l_1 -norm over the standard simplex

$$\Delta := \left\{ x \in \mathbb{R}^d : \sum_{i=1}^d x_i = 1, x \geq 0 \right\} ,$$

and its associated Bregman distance function is

$$D_\omega(x, x') = \sum_{i=1}^d x'_i \log\left(\frac{x'_i}{x_i}\right) .$$

The main motivation to use a generalized distance generating function rather than the usual Euclidean distance function is to design an optimization algorithm that can take advantage of the geometry of the feasible set.

The associated convergence result now reads as follows.

Corollary 1. Consider using the following proximal gradient method to solve (1):

$$\begin{aligned} x_{k+1} &= \arg \min_{x \in \mathcal{X}} \left\{ \langle g_k, x - x_k \rangle + \frac{1}{\alpha} D_\omega(x, x_k) + h(x) \right\}, \\ g_k &= \sum_{n=1}^N \nabla f_n(x_{k-\tau_k^n}) . \end{aligned} \quad (12)$$

Assume that $D_\omega(\cdot, \cdot)$ satisfies:

$$\frac{\mu_\omega}{2} \|x - y\|^2 \leq D_\omega(x, y) \leq \frac{L_\omega}{2} \|x - y\|^2 . \quad (13)$$

Assume also that the problem satisfies assumptions A1–A4, and that the step-size α satisfies:

$$\alpha \leq \frac{L_\omega \left(1 + \frac{\mu}{L} \frac{1}{\bar{\tau}+1} \frac{\mu_\omega}{L_\omega}\right)^{\frac{1}{\bar{\tau}+1}} - 1}{\mu} ,$$

where $L = \sum_{n=1}^N L_n$. Then, the iterates generated by the method satisfy:

$$D_\omega(x^*, x_k) \leq \left(\frac{L_\omega}{\mu\alpha + L_\omega}\right)^k D_\omega(x^*, x_0) .$$

Proof: The analysis is similar to that of Theorem 1. This time, the optimality condition of (12) implies:

$$\begin{aligned} \langle g_k, x_{k+1} - x \rangle &\leq \frac{1}{\alpha} \langle \nabla\omega(x_{k+1}) - \nabla\omega(x_k), x - x_{k+1} \rangle \\ &\quad + \langle s(x_{k+1}), x - x_{k+1} \rangle \quad \forall x \in \mathcal{X}. \end{aligned} \quad (14)$$

Using the following four-point equality

$$\begin{aligned} D_\omega(a, d) - D_\omega(c, d) - D_\omega(a, b) + D_\omega(c, b) &= \\ \langle \nabla\omega(b) - \nabla\omega(d), a - c \rangle, \end{aligned}$$

in (14) with $a = x$, $b = c = x_{k+1}$ and $d = x_k$, and following the steps of the proof of Theorem 1, we obtain:

$$\begin{aligned} \frac{\mu}{2} \|x_{k+1} - x^*\|^2 + \frac{1}{\alpha} D_\omega(x^*, x_{k+1}) &\leq \\ \frac{1}{\alpha} D_\omega(x^*, x_k) - \frac{1}{\alpha} D_\omega(x_{k+1}, x_k) &+ \\ + \frac{L(\bar{\tau} + 1)}{2} \sum_{j=k-\bar{\tau}}^k \|x_{j+1} - x_j\|^2. \end{aligned}$$

This time, using the upper and lower bounds of (13) on the left and right hand-side of the above inequality, respectively, and rearranging the terms, we arrive at:

$$\begin{aligned} D_\omega(x^*, x_{k+1}) &\leq \frac{L_\omega}{\mu\alpha + L_\omega} D_\omega(x^*, x_k) \\ &\quad - \frac{L_\omega}{\mu\alpha + L_\omega} D_\omega(x_{k+1}, x_k) \\ &\quad + \frac{\alpha L(\bar{\tau} + 1)}{\mu\alpha + L_\omega} \frac{L_\omega}{\mu_\omega} \sum_{j=k-\bar{\tau}}^k D_\omega(x_{k+1}, x_j) \end{aligned}$$

Applying Lemma 1 with $V_k = D_\omega(x^*, x_{k+1})$, $w_k = D_\omega(x_{k+1}, x_k)$, $a = b = \frac{L_\omega}{\mu\alpha + L_\omega}$, $c = \frac{\alpha L(\bar{\tau} + 1)}{\mu\alpha + L_\omega} \frac{L_\omega}{\mu_\omega}$ and $k_0 = \bar{\tau}$ completes the proof. \blacksquare

VI. NUMERICAL EXAMPLES

In this section, we present numerical examples which verify our theoretical bound in different settings. First, we simulate the implementation of Algorithms 1 and 2 on a parameter server architecture to solve a small, toy problem. Then, we implement the framework on Amazon EC2 and solve a binary classification problem on three different real-world datasets.

A. Toy problem

To verify our theoretical bounds provided in Theorem 1 and Corollary 1, we consider solving (1) with

$$\begin{aligned} f_n(x) &= \begin{cases} (x_n - c)^2 + \frac{1}{2}(x_{n+1} + c)^2, & n = 1, \\ \frac{1}{2}(x_{n-1} + c)^2 + \frac{1}{2}(x_n - c)^2, & n = N, \\ \frac{1}{2}(x_{n-1} + c)^2 + \frac{1}{2}(x_n - c)^2 + \frac{1}{2}(x_{n+1} + c)^2, & \end{cases} \\ h(x) &= \lambda_1 \|x\|_1 + I_{\mathcal{X}}(x), \\ \mathcal{X} &= \{x \geq 0\}, \end{aligned}$$

for some $c \geq 0$. We use $D_\omega(x, x_k) = \frac{1}{2} \|x - x_k\|_p^2$ in the proximal step (12) and consider both $p = 1.5$ and $p = 2$.

It can be verified that $\nabla F(x)$ is $(N + 1)$ -continuous and $F(x)$ is 2-strongly convex, both with respect to $\|\cdot\|_2$, and that the optimizer for the problem is $x^* = \frac{\max(0, c - \lambda_1)}{3} e_1$, where e_n denotes the n^{th} basis vector. Moreover, it can be shown that if $p \in (1, 2]$, then $\mu_\omega = 1$ and $L_\omega = N^{2/p-1}$ satisfy (13) with respect to $\|\cdot\|_2$.

We select the problem parameters $N = 100$, $c = 3$ and $\lambda_1 = 1$. We simulate solving the problem with $W = 4$ workers, where at each iteration k , a worker w is selected uniformly at random to return their gradient, information evaluated on stale information $x_{k-\tau_k^w}$, to the master. Here, at time k , τ_k^w is simply the number of iterations since the last time worker w was selected. Each worker holds $N/W = 25$ component functions, and we tune step-size based on the assumption that $\bar{\tau} = W$. Figure 1 shows the results of a representative simulation. As can be observed, the iterates converge to the optimizer and the theoretical bound derived is valid.

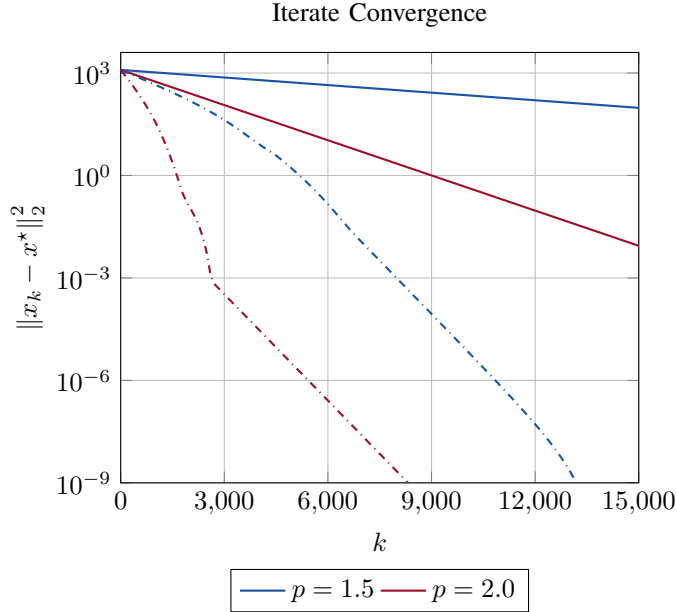


Fig. 1. Convergence of the iterates in toy problem. Solid lines represent our theoretical upper bound, whereas dash-dotted lines represent simulation results.

B. Binary classification on actual datasets

Next, we consider solving a regularized, sparse binary classification problem on three different datasets: `rcv1` (sparse) [16], `url` (sparse) [17] and `epsilon` (dense) [18]. To this end, we implement the parameter server framework in the Julia language, and instantiate it with Problem (1):

$$f_n(x) = \frac{1}{N} \left(\log(1 + \exp(-b_n \langle a_n, x \rangle)) + \frac{1}{2} \lambda_2 \|x\|_2^2 \right),$$

$$h(x) = \lambda_1 \|x\|_1,$$

Here, $a_n \in \mathbb{R}^d$ is the feature vector for sample n , and $b_n \in \{-1, 1\}$ is the corresponding binary label. We pick $\lambda_1 = 10^{-5}$ and $\lambda_2 = 10^{-4}$ for `rcv1` and `epsilon` datasets, and $\lambda_1 = 10^{-3}$ and $\lambda_2 = 10^{-4}$ for `url`. `rcv1` is already normalized to have unit norm in its samples; hence, we normalize `url` and `epsilon` datasets to have comparable problem instances.

`rcv1` is a text categorization test collection from Reuters, having $N = 804414$ documents and $d = 47236$ features (density: 0.16%) for each document. We choose to classify *sports*, *disaster* and *government* related articles from the corpus. `url` is a collection of data for identification of malicious URLs. It has $N = 2396130$ URL samples, each having $d = 64$ real valued features out of a total of 3231961 attributes (density: 18.08%). Finally, `epsilon` is a synthetic, dense dataset, having $N = 500000$ samples and $d = 2000$ features.

It can be verified that $\nabla F(x)$ is $(1/4 \|A\|_2^2 + \lambda_2)$ -Lipschitz continuous with $\|A\|_2^2 = 1$ in all the examples, and $F(x)$ is λ_2 -strongly convex with respect to $\|\cdot\|_2$.

We create three `c4.2xlarge` compute nodes in Amazon’s Elastic Compute Cloud. The compute nodes are physically located in Ireland (`eu`), North Virginia (`us`) and Tokyo (`ap`), respectively. Then, we assign one CPU from each node as workers, resulting in a total of 3 workers, and we pick the master node at KTH in Sweden. We run a small number of iterations of the algorithms to obtain an *a priori* delay distribution of the workers in this setting, and we observe that $\bar{\tau} = 6$.

In Figures 2 and 3, we present the convergence results of our experiments and delay distributions of the workers, respectively. As in the previous example, the iterates converge to the optimizer and the theoretical bound derived is valid. Another observation worth noting is that the denser the datasets become, the smaller the gap between the actual iterates and the theoretical upper bound gets.

VII. DISCUSSIONS AND CONCLUSION

In this paper, we have studied the use of parameter server framework on solving regularized machine learning problems. One class of methods applicable for this framework is the proximal incremental aggregated gradient method. We have shown that when the objective function is strongly convex, the iterates generated by the method converges linearly to the global optimum. We have also given constant step-size rule when the degree of asynchrony in the architecture is known. Moreover, we have validated our theoretical bound by simulating the parameter server architecture on two different problems.

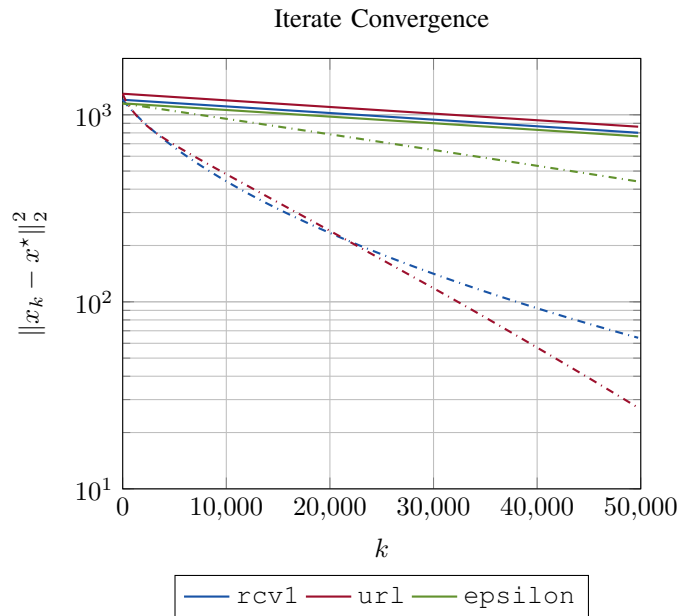


Fig. 2. Convergence of the iterates in Amazon EC2 experiments. Solid lines represent our theoretical upper bound, whereas dash-dotted lines represent experiment results.

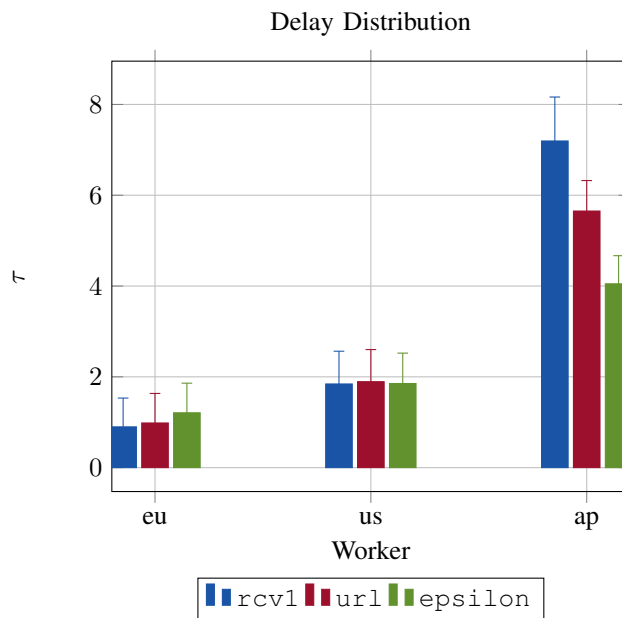


Fig. 3. Worker delays in Amazon EC2 experiments. Bars represent the mean delays, whereas vertical stacked lines represent the standard deviation. For each worker, from left to right, we present the delays obtained in `rcv1`, `url` and `epsilon` experiments, respectively.

REFERENCES

- [1] Y. Nesterov, “Efficiency of coordinate descent methods on huge-scale optimization problems,” *SIAM Journal on Optimization*, vol. 22, no. 2, pp. 341–362, 2012.
- [2] O. Dekel, R. Gilad-Bachrach, O. Shamir, and L. Xiao, “Optimal distributed online prediction using mini-batches,” *Journal of Machine Learning Research*, vol. 13, pp. 165–202, 2012.
- [3] B. Recht, C. Re, S. Wright, and F. Niu, “Hogwild: A lock-free approach to parallelizing stochastic gradient descent,” in *Advances in Neural Information Processing Systems 24*, 2011, pp. 693–701.
- [4] M. Li, L. Zhou, Z. Yang, A. Li, F. Xia, D. G. Andersen, and A. Smola, “Parameter server for distributed machine learning,” in *Big Learning NIPS Workshop*, vol. 1, 2013.
- [5] A. Agarwal and J. C. Duchi, “Distributed delayed stochastic optimization,” in *Advances in Neural Information Processing Systems*, 2011, pp. 873–881.
- [6] H. R. Feyzmahdavian, A. Aytakin, and M. Johansson, “An asynchronous mini-batch algorithm for regularized stochastic optimization,” *IEEE Transactions on Automatic Control*, vol. PP, no. 99, pp. 1–1, 2016.
- [7] M. Gurbuzbalaban, A. Ozdaglar, and P. Parrilo, “On the convergence rate of incremental aggregated gradient algorithms,” jun 2015, arXiv: 1506.02081v1 [math.OC].

- [8] D. P. Bertsekas, “Incremental gradient, subgradient, and proximal methods for convex optimization: A survey,” *Optimization for Machine Learning*, vol. 2010, pp. 1–38, 2011.
- [9] M. V. Solodov, “Incremental gradient algorithms with stepsizes bounded away from zero,” *Computational Optimization and Applications*, vol. 11, no. 1, pp. 23–35, 1998.
- [10] D. Blatt, A. O. Hero, and H. Gauchman, “A convergent incremental gradient method with a constant step size,” *SIAM J. Optim.*, vol. 18, no. 1, pp. 29–51, jan 2007. [Online]. Available: <http://dx.doi.org/10.1137/040615961>
- [11] N. D. Vanli, M. Gurbuzbalaban, and A. Ozdaglar, “Global convergence rate of proximal incremental aggregated gradient methods,” aug 2016, arXiv: 1608.01713v1 [math.OC].
- [12] M. Schmidt, N. Le Roux, and F. Bach, “Minimizing Finite Sums with the Stochastic Average Gradient,” may 2016, arXiv: 1309.2388v2 [math.OC]. [Online]. Available: <https://hal.inria.fr/hal-00860051>
- [13] A. Defazio, F. Bach, and S. Lacoste-Julien, “Saga: A fast incremental gradient method with support for non-strongly convex composite objectives,” in *Advances in Neural Information Processing Systems 27*. Curran Associates, Inc., 2014, pp. 1646–1654. [Online]. Available: <http://papers.nips.cc/paper/5258-saga-a-fast-incremental-gradient-method-with-support-for-non-strongly-convex-composite-objectives.pdf>
- [14] J. Mairal, “Incremental majorization-minimization optimization with application to large-scale machine learning,” *SIAM Journal on Optimization*, vol. 25, no. 2, pp. 829–855, 2015.
- [15] L. Xiao and T. Zhang, “A proximal stochastic gradient method with progressive variance reduction,” *SIAM Journal on Optimization*, vol. 24, no. 4, pp. 2057–2075, 2014. [Online]. Available: <http://dx.doi.org/10.1137/140961791>
- [16] D. D. Lewis, Y. Yang, T. G. Rose, and F. Li, “Rcv1: A new benchmark collection for text categorization research,” *J. Mach. Learn. Res.*, vol. 5, pp. 361–397, Dec. 2004. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1005332.1005345>
- [17] J. Ma, L. K. Saul, S. Savage, and G. M. Voelker, “Identifying suspicious urls: An application of large-scale online learning,” in *Proceedings of the 26th Annual International Conference on Machine Learning (ICML 2009)*, ser. ICML ’09, ACM. Montreal, Quebec: ACM, June 2009, pp. 681–688. [Online]. Available: <http://doi.acm.org/10.1145/1553374.1553462>
- [18] “Pascal large scale learning challenge.” Accessed: 2016-10-07. [Online]. Available: <http://largescale.ml.tu-berlin.de/>