

Adaptive Distributed Stochastic Gradient Descent for Minimizing Delay in the Presence of Stragglers

Serge Kas Hanna*, Rawad Bitar*, Parimal Parag†, Venkat Dasari ‡, and Salim El Rouayheb*

* Department of Electrical and Computer Engineering, Rutgers University, Piscataway, NJ, USA

† Department of Electrical Communication Engineering, Indian Institute of Science, Bengaluru, KA, India

‡ US Army Research Laboratory, Aberdeen Proving Ground, MD, USA

Emails: {serge.k.hanna, rawad.bitar, salim.elrouayheb}@rutgers.edu, parimal@iisc.ac.in, venkateswara.r.dasari.civ@mail.mil

Abstract

We consider the setting where a master wants to run a distributed stochastic gradient descent (SGD) algorithm on n workers each having a subset of the data. Distributed SGD may suffer from the effect of stragglers, i.e., slow or unresponsive workers who cause delays. One solution studied in the literature is to wait at each iteration for the responses of the fastest $k < n$ workers before updating the model, where k is a fixed parameter. The choice of the value of k presents a trade-off between the runtime (i.e., convergence rate) of SGD and the error of the model. Towards optimizing the error-runtime trade-off, we investigate distributed SGD with adaptive k . We first design an adaptive policy for varying k that optimizes this trade-off based on an upper bound on the error as a function of the wall-clock time which we derive. Then, we propose an algorithm for adaptive distributed SGD that is based on a statistical heuristic. We implement our algorithm and provide numerical simulations which confirm our intuition and theoretical analysis.

I. INTRODUCTION

We consider a distributed computation setting in which a master wants to learn a model on a large amount of data in his possession by dividing the computations on n workers. The data at the master consists of a matrix $X \in \mathbb{R}^{m \times d}$ representing m data vectors $\mathbf{x}_\ell$, $\ell = 1, \dots, m$, and a vector $\mathbf{y} \in \mathbb{R}^m$ representing the labels of the rows of X . Define $A \triangleq [X|\mathbf{y}]$ to be the concatenation of X and $\mathbf{y}$. The master would like to find a model $\mathbf{w}^* \in \mathbb{R}^d$ that minimizes a loss function $F(A, \mathbf{w})$, i.e.,

$$\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathbb{R}^d} F(A, \mathbf{w}). \quad (1)$$

This optimization problem can be solved using Gradient Descent (GD), which is an iterative algorithm that consists of the following update at each iteration j ,

$$\mathbf{w}_{j+1} = \mathbf{w}_j - \eta \nabla F(A, \mathbf{w}_j) \triangleq \mathbf{w}_j - \frac{\eta}{m} \sum_{\ell=1}^m \nabla F(\mathbf{a}_\ell, \mathbf{w}_j), \quad (2)$$

where η is the step size, and $\nabla F(A, \mathbf{w})$ is the gradient of $F(A, \mathbf{w})$. To distribute the computations, the master partitions the data equally to n workers. The data partitioning is horizontal, i.e., each worker receives a set of rows of A with all their corresponding columns. Let S_i be the sub-matrix of A sent to worker i . Each worker computes a partial gradient defined as

$$\nabla F(S_i, \mathbf{w}_j) \triangleq \frac{1}{s} \sum_{\mathbf{a}_\ell \in S_i} \nabla F(\mathbf{a}_\ell, \mathbf{w}_j), \quad (3)$$

where $s = m/n$ is the number of rows in S_i (assuming n divides m). The master computes the average of the received partial gradients to obtain the actual gradient $\nabla F(A, \mathbf{w})$, and then updates $\mathbf{w}_j$.

In this setting, waiting for the partial gradients of all the workers slows down the process as the master has to wait for the stragglers [1], i.e., slow or unresponsive workers, in order to update $\mathbf{w}_j$. Many approaches have been proposed in the literature to alleviate the problem of stragglers. A natural approach is to simply ignore the stragglers and obtain an estimate of the gradient rather than the full gradient, see [2], [3]. This framework emerges from single-node (non-distributed) mini batch stochastic gradient descent (SGD) [4]. Batch SGD is a relaxation of GD in

which $\mathbf{w}_j$ is updated based on a subset (batch) B ($|B| < m$) of data vectors that is chosen uniformly at random from the set of all m data vectors, i.e.,

$$\mathbf{w}_{j+1} = \mathbf{w}_j - \frac{\eta}{|B|} \sum_{\mathbf{a}_\ell \in B} F(\mathbf{a}_\ell, \mathbf{w}_j). \quad (4)$$

It is shown that SGD converges to $\mathbf{w}^*$ under mild assumptions on the loss function $F(A, \mathbf{w})$, but may require a larger number of iterations as compared to GD [4]–[10].

Consider the approach where the master updates the model based on the responses of the fastest $k < n$ workers and ignores the remaining stragglers. Henceforth, we call this approach *fastest- k SGD*. The update rule for fastest- k SGD is given by

$$\mathbf{w}_{j+1} = \mathbf{w}_j - \frac{\eta}{k} \sum_{i \in R_j} \nabla F(S_i, \mathbf{w}_j) \triangleq \mathbf{w}_j - \eta \hat{\mathbf{g}}(\mathbf{w}_j), \quad (5)$$

where R_j is the set of the fastest k workers at iteration j ; and $\hat{\mathbf{g}}(\mathbf{w}_j)$ is the average of the partial gradients received by the master at iteration j which is an estimate of the full gradient $\nabla F(A, \mathbf{w}_j)$. Note that if we assume that the response times of the workers are random *iid*, then one can easily show that fastest- k SGD is essentially equivalent to single-node batch SGD since the master updates the model at each iteration based on a uniformly random batch of data vectors belonging to the set of the fastest k workers. Therefore, fastest- k SGD converges to $\mathbf{w}^*$ under the random *iid* assumption on the response times of the workers and the standard assumptions on the loss function $F(A, \mathbf{w})$.

The convergence rate of distributed SGD depends on two factors simultaneously: (i) the error in the model versus the number of iterations; (ii) the time spent per iteration. Therefore, in this work we focus on studying the convergence rate with respect to the wall-clock time rather than the number of iterations. In fastest- k SGD with fixed step size, the choice of the value of k presents a trade-off between the convergence rate and the error floor. Namely, choosing a small value of k will lead to fast convergence since the time spent per iteration would be short, however, this will also result in a low accuracy in the final model, i.e., higher error floor. Towards optimizing this trade-off, we study adaptive policies for fastest- k SGD where the master starts with waiting for a small number of workers k and then gradually increases k to minimize the error as a function of time. Such an optimal adaptive policy would guarantee that the error is minimized at any instant of the wall-clock time. This would be particularly useful in applications where SGD is run with a deadline, since the learning algorithm would achieve the best accuracy within any time restriction.

A. Related work

1) *Distributed SGD*: The works that are closely related to our work are that of [2], [3]. In [3] the authors study fastest- k SGD for a predetermined k . In [2], the authors consider the same setting as [3] and analyze the convergence rate of fastest- k SGD with respect to the number of iterations. In addition to the convergence analysis with respect to the number of iterations, the authors in [2] separately analyze the time spent per iteration as a function of k .

Several works proposed using redundancy when distributing the data to the workers. The master then uses coding theoretical tools to recover the gradient in the presence of a fixed number of stragglers, for example [11]–[24]. In [25]–[27] the authors propose a mixed strategy in which the master distributes the data redundantly and uses coding techniques to obtain the whole gradient for a given number of stragglers. In addition, if more workers than accounted for are stragglers, the master can use the same coding techniques to compute an estimate of the gradient.

Note that the setting of the previously mentioned works, and the setting of interest for our work, focuses on the so-called synchronous SGD in which the workers are all synchronized at each iteration (i.e., have the same model). The literature also studies the asynchronous setting. In asynchronous distributed SGD, whenever a worker finishes its assigned computation, it sends the result to the master who directly updates $\mathbf{w}$ and sends an updated $\mathbf{w}$ to that worker who starts a new computation of the partial gradient while the other workers continue their previous computation, see for example [2], [28]–[32].

2) *Single-node SGD*: Murata [33] showed that irrespective of its convergence speed, the single-node SGD algorithm with fixed step size goes through a transient phase and a stationary phase. In the transient phase, $\mathbf{w}_j$ approaches $\mathbf{w}^*$ exponentially fast in the number of iterations. Whereas, in the stationary phase $\mathbf{w}_j$ oscillates around $\mathbf{w}^*$. Note that if a decreasing step size over the iterations is used, then $\mathbf{w}_j$ converges to $\mathbf{w}^*$ rather than oscillating

around it, however this leads to a long transient phase and hence a lower convergence rate. To detect the phase transition, [34] uses a statistical test based on Pflug's method [35] for stochastic approximation. Detecting the phase transition serves many purposes, such as indicating when to stop the SGD algorithm or when to start implementing further tricks to reduce the distance between $\mathbf{w}_j$ and $\mathbf{w}^*$. In this paper, we build on this line of work to derive the times at which the master should start waiting for more workers in fastest- k SGD.

In another line of work on single-node SGD, the authors in [36] suggested increasing the batch size with the number of iterations as an alternative to decreasing the step size. The results in [36] show that increasing the batch size while keeping a constant step size, leads to near-identical model accuracy as decreasing the step size, but with fewer parameter updates, i.e., shorter training time.

B. Our contributions

We focus on straggler mitigation in synchronous distributed SGD with fixed step size. We consider a setting where the master distributes the data without redundancy. Under standard assumptions on the loss function, and assuming independent and identical distributed random response times for the workers, we give a theoretical bound on the error fastest- k SGD as a function of time rather than the number of iterations. We derive an adaptive policy which shows that this bound on the error can be optimized as a function of time by increasing the value of k at specific times which we explicitly determine in terms of the system parameters. Furthermore, we develop an algorithm for adaptive fastest- k SGD that is based on a statistical heuristic which works while being oblivious to the system parameters. We implement this algorithm and provide numerical simulations which show that the adaptive fastest- k SGD can outperform both non-adaptive fastest- k SGD and asynchronous SGD.

II. PRELIMINARIES

In this paper we focus on fastest- k SGD with fixed step size. In fastest- k SGD, as soon as the master updates $\mathbf{w}$ based on (5), it sends it to all n workers, and all workers start computing their partial gradients based on the new model $\mathbf{w}$, i.e., the ongoing computations of the slowest $n - k$ workers from the previous iteration are automatically cancelled. Note that it is possible to obtain the partial computations of the slow $n - k$ workers to improve the estimate of the full gradient. However, this improvement would come at the expense of higher communication cost between the master and the workers. In our work, we are considering distributed systems with scarce communication resources and abundant computation resource. In such a setting, cancellation of partial computations is justified by the reduced communication cost.

We consider a random straggling model where the time spent by worker i to finish the computation of its partial gradient (i.e., response time) is a random variable X_i , for $i = 1, \dots, n$. We assume that $X_i, i = 1, \dots, n$, are *iid* and independent across iterations. Therefore, the time per iteration for fastest- k SGD is given by the k^{th} order statistic of the random variables $X_1, \dots, X_n$, denoted by $X_{(k)}$. In the previously described setting, the following bound on the error of fastest- k SGD as a function of the number of iterations was shown in [2], [5].

Proposition 1 (Error vs. iterations of fastest- k SGD [2], [5]). *Under certain conditions (stated in the Appendix), the error of fastest- k SGD after j iterations with fixed step size satisfies*

$$\mathbb{E}[F(\mathbf{w}_j) - F^*] \leq \frac{\eta L \sigma^2}{2cks} + (1 - \eta c)^j \left(F(\mathbf{w}_0) - F^* - \frac{\eta L \sigma^2}{2cks} \right),$$

where L and c are the Lipschitz and the strong convexity parameters of the loss function respectively, F^* is the optimal value of the loss function, and σ^2 is the variance bound on the gradient estimate.

The notations used throughout the paper are summarized in Table I.

Variable	Description	Variable	Description
F	loss function	σ^2	variance bound on the gradient estimate
m	total number of data vectors	k	number of workers to wait for in k -sync SGD
d	data dimension	X_i	random response time of worker i
n	number of workers	$X_{(k)}$	k^{th} order statistic of $X_1, \dots, X_n$
s	number of data vectors per worker	μ_k	$\mathbb{E}[X_{(k)}]$: average of $X_{(k)}$
η	step size	σ_k^2	$Var[X_{(k)}]$: variance of $X_{(k)}$
L	Lipschitz constant	t	wall-clock time
c	strong convexity parameter	$J(t)$	number of iterations completed in time t

TABLE I: Summary of the notations used in the paper.

III. THEORETICAL ANALYSIS

In this section, we present our theoretical results. The proofs of these results are given in Section VI. In Lemma 1, by applying techniques from renewal theory, we give a bound on the error of fastest- k SGD as a function of the wall-clock time t rather than the number of iterations. The bound holds with high probability for large t and is based on Proposition 1.

Lemma 1 (Error vs. wall-clock time of fastest- k SGD). *Under the same assumptions as Proposition 1, the error of fastest- k SGD after wall-clock time t with fixed step size satisfies*

$$\mathbb{E}[F(\mathbf{w}_t) - F^* | J(t)] \leq \frac{\eta L \sigma^2}{2cks} + (1 - \eta c)^{\frac{t}{\mu_k}(1-\epsilon)} \left(F(\mathbf{w}_0) - F^* - \frac{\eta L \sigma^2}{2cks} \right), \quad (6)$$

with high probability ($Pr \geq 1 - \frac{\sigma_k^2}{\epsilon^2}(\frac{2}{t\mu_k} + \frac{1}{t^2})$) for large t , where $0 < \epsilon \ll 1$ is a constant error term, $J(t)$ is the number of iterations completed in time t , and μ_k is the average of the k^{th} order statistic $X_{(k)}$.

Notice that the first term in (6) is constant (independent of t), whereas the second term decreases exponentially in t ($\eta c < 1$ from [5]). In fact, it is well known that the SGD with constant step size goes first through a transient phase where the error decreases exponentially fast, and then enters a stationary phase where the error oscillates around a constant term [33]. From (6), it is easy to see that the rate of the exponential decrease in the transient phase is governed by the value of $1/\mu_k$. Since μ_k is an increasing function of k , the exponential decrease is fastest for $k = 1$ and slowest for $k = n$. Whereas the stationary phase error which is upper bounded by $\eta L \sigma^2 / 2cks$, is highest for $k = 1$ and lowest for $k = n$. This creates a trade-off between the rate of decrease of the error in the transient phase, and the error floor achieved in the stationary phase. Ultimately, we would like to first have a fast decrease through the transient phase, and then have a low error in the stationary phase. To this end, we look for an adaptive policy for varying k that starts with $k = 1$ and then switches to higher values of k at specific times in order to optimize error-runtime trade-off. Such an adaptive policy guarantees that the error is minimized at every instant of the wall-clock time t .

Since the bound in (6) holds with high probability for large t , we explicitly derive the switching times for adaptive fastest- k SGD that optimize this bound. Note that for the sake of simplicity, we drop the constant error term ϵ in our analysis.

Theorem 1 (Bound-optimal policy for adaptive fastest- k SGD). *The bound-optimal times $t_k, k = 1, \dots, n-1$, at which the master should switch from waiting for the fastest k workers to waiting for the fastest $k+1$ workers are given by*

$$t_k = t_{k-1} + \frac{\mu_k}{-\ln(1 - \eta c)} \times \left[\ln(\mu_{k+1} - \mu_k) - \ln(\eta L \sigma^2 \mu_k) + \ln(2ck(k+1)s(F(\mathbf{w}_{t_{k-1}}) - F^*) - \eta L(k+1)\sigma^2) \right],$$

where $t_0 = 0$.

Example 1 (Theoretical analysis on fastest- k SGD with iid exponential response times). *Suppose $X_i \sim \exp(\mu)$, $i = 1, \dots, n$, is the random time taken by worker i to compute its partial gradient, where X_i 's are iid. The average time spent by the master per iteration in fastest- k SGD is the average of the k^{th} order statistic of the exponential random variables $X_1, \dots, X_n$, given by*

$$\mu_k = \frac{1}{\mu} \sum_{i=n-k+1}^n \frac{1}{i}. \quad (7)$$

Let $n = 5, \mu = 5, \eta = 0.001, \sigma^2 = 10, F(\mathbf{w}_0) - F^* = 100, L = 2, c = 1, s = 10$. We evaluate the bound in Lemma 1 for multiple fixed values of k (non-adaptive) and compare it to adaptive fastest- k SGD if we apply the switching times in Theorem 1. Let $k_{fixed}^*(t)$ and $k_{adaptive}^*(t)$ be the bound-optimal choices of the number of workers to wait for as a function of the wall-clock time for non-adaptive and adaptive fastest- k SGD, respectively. The results are shown in Fig. 1 and 2.




Fig. 1: The upper bound on the error given by (6) as a function of time, evaluated for $k = 1, 2, 3, 4, 5$.




Fig. 2: The bound-optimal choices of the number of workers to wait for $k_{fixed}^*(t)$ and $k_{adaptive}^*(t)$ for non-adaptive and adaptive fastest- k SGD, respectively.

The plots in Fig. 1 and Fig. 2a show that for the time interval $[0, 1098)$, the bound-optimal choice of fixed (non-adaptive) k that minimizes the error is $k = 1$. Namely, if one wants to run non-adaptive fastest- k SGD for less than 1098 time units, then the best strategy is to wait for the fastest single worker before updating the model. It is easy to see from Fig. 1 that any other choice of fixed $k > 1$, leads to a higher error in the time interval $[0, 1098)$. For a runtime greater than 6623 time units, the bound-optimal choice for non-adaptive fastest- k SGD is to wait for the responses of all workers. Whereas for adaptive fastest- k SGD, notice from Fig. 1 and Fig. 2b, that for the time interval $[0, 589)$, the adaptive policy assigns $k = 1$ since it gives the fastest error decrease in the beginning. Then, as the error approaches the stationary phase, the policy increases k to $k = 2$. This allows the error to decrease below the error floor for $k = 1$. The procedure continues until k attains its maximum value of $k = n = 5$. In conclusion, the results demonstrate the adaptive version enables achieving lower error values in less time.

This analysis shows the potential of adaptive strategies in optimizing the error-runtime trade-off. It also suggests that the value of k should be gradually increased throughout the runtime of fastest- k SGD in order to optimize this trade-off. Although Theorem 1 provides useful insights about how to adapt k over time, applying the exact

switching times derived in Theorem 1 may not be effective in practice due to the following reasons: (i) the policy optimizes an upper bound on the error (Lemma 1) which is probabilistic and may be loose; (ii) the policy requires the knowledge of several system parameters including the optimal value of the loss function F^* which is typically unknown. Nevertheless, we use the insights provided by the theoretical analysis to design a practical algorithm for adaptive fastest- k SGD. This algorithm is based on a statistical heuristic and is oblivious to the system parameters as we explain in Section IV.

IV. ADAPTIVE FASTEST- k SGD ALGORITHM

In this section we present an algorithm for adaptive fastest- k SGD that is realizable in practice. As previously mentioned, SGD with fixed step size goes first through a transient phase where the error decreases exponentially fast, and then enters a stationary phase where the error oscillates around a constant term. Initially, the exponential decrease is fastest for $k = 1$. Then, as the stationary phase approaches, the error decrease becomes slower and slower until a point where the error starts oscillating around a constant term and does not decrease any further. At this point, increasing k allows the error to decrease further because the master would receive more partial gradients and hence would obtain a better estimate of the full gradient. The goal of the adaptive policy is to detect this phase transition in order to increase k and keep the error decreasing.

The adaptive policy we present in this section detects this phase transition by employing a statistical test based on a modified version of Pflug's procedure for stochastic approximation [35]. The main component of our policy is to monitor the signs of the products of consecutive gradients computed by the master based on (5). The underlying idea is that in the transient phase, due to the exponential decrease of the error, the gradients are likely to point in the same direction, hence their inner product is positive. Our policy consists of utilizing a counter that counts the difference between the number of times the product of consecutive gradients is negative (i.e., $\hat{\mathbf{g}}_j^T \hat{\mathbf{g}}_{j-1} < 0$) and the number of times this product is positive, throughout the iterations of the algorithm.

In the beginning of the algorithm, we expect the value of the counter to be negative and decrease because of the exponential decrease in the error. Then, as the error starts moving towards the stationary phase, negative gradient products will start accumulating until the value of the counter becomes larger than a certain positive threshold. At this point, we declare a phase transition and increase k . The complete algorithm is given in Algorithm 1.

Algorithm 1: Fastest- k SGD

```

input : starting point  $\mathbf{w}_0$ , data  $\{X, \mathbf{y}\}$ , number of workers  $n$ , step size  $\eta$ , maximum number of iterations  $J$ , adaptation
         parameters  $\text{step}$ ,  $\text{thresh}$ ,  $\text{burnin}$ 
output: weight vector  $\mathbf{w}_J$ 
 $j \leftarrow 1$ 
 $k \leftarrow 1$ 
 $\text{countNegative} \leftarrow 0$ 
 $\text{countIter} \leftarrow 1$ 
Distribute  $X$  to the  $n$  workers
while  $j \leq J$  do
    Send  $\mathbf{w}_{j-1}$  to all workers
    Collect the responses of the fastest  $k$  workers
     $\mathbf{w}_j \leftarrow \mathbf{w}_{j-1} - \eta \hat{\mathbf{g}}_{j-1}$ 
    if  $\hat{\mathbf{g}}_j^T \hat{\mathbf{g}}_{j-1} < 0$  then
         $\text{countNegative} \leftarrow \text{countNegative} + 1$ 
    else
         $\text{countNegative} \leftarrow \text{countNegative} - 1$ 
    if  $\text{countNegative} > \text{thresh}$  and  $\text{countIter} > \text{burnin}$  and  $k \leq n - \text{step}$  then
         $k \leftarrow k + \text{step}$ 
         $\text{countNegative} \leftarrow 0$ 
         $\text{countIter} \leftarrow 0$ 
     $\text{countIter} \leftarrow \text{countIter} + 1$ 
     $j \leftarrow j + 1$ 
return  $\mathbf{w}_J$ 

```

V. SIMULATIONS

A. Experimental setup

We simulated the performance of the fastest- k SGD (Algorithm 1) described earlier for n workers on synthetic data X . We generated X as follows: (i) we pick each row vector $\mathbf{x}_\ell$, $\ell = 1, \dots, m$, independently and uniformly at random from $\{1, 2, \dots, 10\}^d$; (ii) we pick a random vector $\bar{\mathbf{w}}$ with entries being integers chosen uniformly at random from $\{1, \dots, 100\}$; and (iii) we generate $\mathbf{y}_\ell \sim \mathcal{N}(\langle \mathbf{x}_\ell, \bar{\mathbf{w}} \rangle, 1)$ for all $\ell = 1, \dots, m$.

We run linear regression using the ℓ_2 loss function, i.e., $F(\mathbf{w}) = \sum_{\ell=1}^m \frac{1}{2} \|\langle \mathbf{x}_\ell, \mathbf{w} \rangle - \mathbf{y}_\ell\|_2^2$. At each iteration, we generate n independent exponential random variables with rate $\mu = 1$.

B. Adaptive k -sync SGD vs Non-adaptive k -sync SGD

Figure 3 compares the performance of the adaptive fastest- k SGD to non-adaptive for $n = 50$ workers. In the adaptive version we start with $k = 10$ and then increase k by 10 until reaching $k = 40$, where the switching times are given by Algorithm 1. Whereas for the non-adaptive version, k is fixed throughout the runtime of the algorithm. The comparison shows that the adaptive version is able to achieve a better error-runtime trade-off than the non-adaptive one. Namely, notice that the adaptive k -sync reaches its lowest error at approximately $t = 2000$, whereas the non-adaptive version reaches the same error only for $k = 40$ at approximately $t = 6000$. These results confirm our intuition and previous theoretical results.




Fig. 3: Error as a function of the wall-clock time for non-adaptive fastest- k SGD with fixed $k = 10, 20, 30, 40$; and adaptive fastest- k SGD (Algorithm 1). The experimental setup is the following: $d = 100$, $m = 2000$, $n = 50$, $\eta = 0.0005$. The adaptation parameters chosen here are $step = 10$, $thresh = 10$, and $burnin = 0.1 \times (\text{number of data points}) = 200$. We start the adaptive fastest- k SGD with $k = 10$ and increase k by 10 until reaching $k = 40$.

C. Comparison to Asynchronous SGD

Figure 4 compares the adaptive fastest- k SGD to the fully asynchronous version of distributed SGD. Recall that in asynchronous SGD, when a worker finishes the computation of its partial gradient, it sends it to the master who immediately updates the model and sends the new model only to this worker. In the meantime, the other workers continue their ongoing computations. The comparison shows that in the beginning, the decrease in the error of fastest- k SGD is as fast as that of asynchronous SGD. However, the error floor for asynchronous SGD is higher than that of fastest- k SGD. Therefore, fastest- k SGD can achieve a higher model accuracy by adapting the parameter k throughout the runtime of the algorithm.




Fig. 4: Error as a function of time for adaptive fastest- k SGD (Algorithm 1) and asynchronous SGD. The experimental setup is the following: $d = 100$, $m = 2000$, $n = 50$, $\eta = 0.0002$. The adaptation parameters chosen here are $step = 5$, $thresh = 10$, and $burnin = 0.1 \times (\text{number of data points}) = 200$. We start the adaptive fastest- k SGD with $k = 1$ and increase k by 5 until reaching $k = 36$.

VI. PROOFS

A. Proof of Lemma 1

Proposition 1 gives an upper bound on the error of fastest- k SGD as a function of the number of iterations j , given the conditions in the Appendix. Let $J(t)$ be the arrival process representing the number of iterations completed in time t . The upper bound as a function of time is given by

$$\mathbb{E}[F(\mathbf{w}_t) - F^* | J(t)] \leq \frac{\eta L \sigma^2}{2cks} + (1 - \eta c)^{J(t)} \left(F(\mathbf{w}_0) - F^* - \frac{\eta L \sigma^2}{2cks} \right). \quad (8)$$

The time between any two consecutive iterations in fastest- k SGD is given by $X_{(k)}$, where $X_{(k)}$ is the k^{th} order statistic of the random variables $X_1, \dots, X_n$. Therefore, $J(t)$ has *iid* inter-arrival times. Let $(T_z \in \mathbb{R}_+ : z \in \mathbb{N})$ be the *iid* sequence of inter-arrival times, where $\mathbb{E}[T_z] = \mu_k$ and $\text{Var}[T_z] = \sigma_k^2$. We define $S_j = \sum_{z=1}^j T_z$, such that $S_0 = 0$ and S_j is the j th renewal instant for *iid* sequence $T = (T_z : z \in \mathbb{N})$ with mean μ_k , variance σ_k^2 , and the moment generating function $M_T(\theta) = e^{\Phi_T(\theta)}$. Then, we can write following concentration inequalities for S_j . First, from the Chebyshev inequality we have

$$P \left\{ \left| \frac{S_j - j\mu_k}{\sqrt{j}\sigma_k} \right| > t \right\} \leq \frac{1}{t^2}. \quad (9)$$

Corresponding to the renewal process, we obtain the counting process

$$J(t) = \sum_{z \in \mathbb{N}} 1_{\{T_z \leq t\}}. \quad (10)$$

Defining $x_l \triangleq y_l \frac{\sigma_k}{\mu_k} \sqrt{\frac{t}{\mu_k}}$ and $x_u \triangleq y_u \frac{\sigma_k}{\mu_k} \sqrt{\frac{t}{\mu_k}}$, we can upper bound the union of the events $\left\{ J(t) > \frac{t}{\mu_k} + x_u \right\}$ and $\left\{ J(t) < \frac{t}{\mu_k} - x_l \right\}$. Since $J(t)$ is integer valued, taking the floor $n_u = \lfloor \frac{t}{\mu_k} + x_u \rfloor$ and the ceil $n_l = \lceil \frac{t}{\mu_k} - x_l \rceil$, we can write using the inverse relationship between the counting process $(J(t) : t \geq 0)$ and the renewal process $(S_n : n \in \mathbb{N})$ as

$$P \left\{ J(t) > \frac{t}{\mu_k} + x_u \right\} = \{S_{n_u} - \mu_k n_u \leq t - \mu_k n_u\} = P \left\{ \frac{S_{n_u} - \mu_k n_u}{\sqrt{n_u} \sigma_k} \leq -\frac{\mu_k x_u}{\sqrt{n_u} \sigma_k} \right\} \leq \frac{\sigma_k^2 n_u}{\mu_k^2 x_u^2}, \quad (11)$$

$$P \left\{ J(t) < \frac{t}{\mu_k} - x_l \right\} = \{S_{n_l} - \mu_k n_l \geq t - \mu_k n_l\} = P \left\{ \frac{S_{n_l} - \mu_k n_l}{\sqrt{n_l} \sigma_k} \geq \frac{\mu_k x_l}{\sqrt{n_l} \sigma_k} \right\} \leq \frac{\sigma_k^2 n_l}{\mu_k^2 x_l^2}. \quad (12)$$

Substituting the value of x_u, x_l in terms of y_u, y_l , we get

$$P \left\{ -y_l \leq \frac{J(t) - \frac{t}{\mu_k}}{\sigma_k \sqrt{\frac{t}{\mu_k^3}}} \leq y_u \right\} \geq 1 - \left(\frac{1}{y_l^2} + \frac{1}{y_u^2} + \frac{\mu_k}{t y_l^2} \right) - \frac{\sigma_k}{\sqrt{t \mu_k}} \left(\frac{1}{y_u} - \frac{1}{y_l} \right). \quad (13)$$

Let $\theta = -\ln(1 - \eta c) > 0$, We can write the following bound for $(1 - \eta c)^{J(t)} = e^{-\theta J(t)}$.

$$P \left\{ e^{-\theta \frac{t}{\mu_k} (1 + \frac{y \sigma_k}{\sqrt{t \mu_k}})} \leq e^{-\theta J(t)} \leq e^{-\theta \frac{t}{\mu_k} (1 - \frac{y \sigma_k}{\sqrt{t \mu_k}})} \right\} \geq 1 - \left(\frac{\frac{\mu_k}{t} + 2}{y^2} \right). \quad (14)$$

By substituting $y = \epsilon \frac{\sqrt{t \mu_k}}{\sigma_k}$ in (14), where $0 < \epsilon \ll 1$ is a constant, we obtain

$$Pr \left\{ (1 - \eta c)^{J(t)} \leq (1 - \eta c)^{\frac{t}{\mu_k} (1 - \epsilon)} \right\} \geq 1 - \frac{\sigma_k^2}{\epsilon^2} \left(\frac{2}{t \mu_k} + \frac{1}{t^2} \right). \quad (15)$$

From (8) and (15), we can see that the bound (6) in Lemma (1) holds with high probability for large t .

B. Proof of Theorem 1

To minimize the error achieved within any time interval, we maximize the convergence rate at every time t , based on the result in Lemma 1 (the error term ϵ is omitted). Substituting $\alpha \triangleq -\ln(1 - \eta c) > 0$ in (6), we define

$$e(t) \triangleq \frac{\eta L \sigma^2}{2ck s} + e^{\frac{-\alpha t}{\mu_k}} \left(F(\mathbf{w}_0) - F^* - \frac{\eta L \sigma^2}{2ck s} \right), \quad (16)$$

where $e(0) = F(\mathbf{w}_0) - F^*$. To evaluate the rate of convergence, we analyze the magnitude of the first order derivative of $e(t)$ with respect to t given by

$$\left| \frac{de(t)}{dt} \right| = \frac{\alpha}{\mu_k} e^{\frac{-\alpha t}{\mu_k}} \left(e(0) - \frac{\eta L \sigma^2}{2ck s} \right). \quad (17)$$

For a fixed $t' \leq t$, we can write the following equivalent definitions of (16) and (17).

$$e(t) = \frac{\eta L \sigma^2}{2ck s} + e^{\frac{-\alpha(t-t')}{\mu_k}} \left(e(t') - \frac{\eta L \sigma^2}{2ck s} \right). \quad (18)$$

$$\left| \frac{de(t)}{dt} \right| = \frac{\alpha}{\mu_k} e^{\frac{-\alpha(t-t')}{\mu_k}} \left(e(t') - \frac{\eta L \sigma^2}{2ck s} \right). \quad (19)$$

The value of k in the expressions above is considered to be fixed. We are interested in an adaptive policy for choosing k that maximizes this magnitude at every time t , and hence we look for a scheduling function $k(t)$ that varies k with time. We observe that there is trade-off between the two terms of $e(t)$ in (16). Waiting for larger number of workers k increases the mean iteration time μ_k , whereas it decreases the stationary phase error (i.e., error floor) $\eta L \sigma^2 / 2ck s$. Assuming that the initial error $e(0)$ is much larger than the stationary phase error, it is easy to see that the rate of decrease in error at $t = 0$ is highest for $k = 1$. Hence, the master starts at $t = 0$ by waiting for the fastest single worker. Subsequently, let $t_k, k = 0, \dots, n-1$, be the time when the master switches from waiting for the fastest k workers to waiting for the fastest $k+1$ workers, with $t_0 = 0$. For $t_{k-1} \leq t \leq t_k$, we have

$$\left| \frac{de(t)}{dt} \right| = \frac{\alpha}{\mu_k} e^{\frac{-\alpha(t-t_{k-1})}{\mu_k}} \left(e(t_{k-1}) - \frac{\eta L \sigma^2}{2ck s} \right). \quad (20)$$

For $t_k \leq t \leq t_{k+1}$, we have

$$\left| \frac{de(t)}{dt} \right| = \frac{\alpha}{\mu_{k+1}} e^{\frac{-\alpha(t-t_k)}{\mu_{k+1}}} \left(e(t_k) - \frac{\eta L \sigma^2}{2c(k+1)s} \right). \quad (21)$$

The bound-optimal switching time t_k is given by the smallest t for which the rate in (21) becomes larger than the rate in (20). Hence t_k can be determined by evaluating the following,

$$\frac{\alpha}{\mu_{k+1}} \left(e(t_k) - \frac{\eta L \sigma^2}{2c(k+1)s} \right) \geq \frac{\alpha}{\mu_k} e^{\frac{-\alpha(t_k-t_{k-1})}{\mu_k}} \left(e(t_{k-1}) - \frac{\eta L \sigma^2}{2ck s} \right), \quad (22)$$

where the LHS in (22) is obtained by evaluating (21) at $t = t_k$. The expression of $e(t_k)$ is given by

$$e(t_k) = \frac{\eta L \sigma^2}{2ck_s} + e^{\frac{-\alpha(t_k - t_{k-1})}{\mu_k}} \left(e(t_{k-1}) - \frac{\eta L \sigma^2}{2ck_s} \right). \quad (23)$$

Substituting (23) in (22) and simplifying we get

$$e^{\frac{-\alpha(t_k - t_{k-1})}{\mu_k}} \leq \frac{(\mu_{k+1} - \mu_k)(2ck_s(k+1)e(t_{k-1}) - \eta L(k+1)\sigma^2)}{\eta L \sigma^2 \mu_k}. \quad (24)$$

Taking logarithms on both sides of (24) we obtain

$$t_k = t_{k-1} + \frac{\mu_k}{\alpha} \left[\ln(\mu_{k+1} - \mu_k) + \ln(2ck_s(k+1)e(t_{k-1}) - \eta L(k+1)\sigma^2) - \ln(\eta L \sigma^2 \mu_k) \right]. \quad (25)$$

where $t_0 = 0$ and $e(t_0) = F(\mathbf{w}_0) - F^*$.

VII. ACKNOWLEDGMENT

This work was supported in part by the Army Research Lab (ARL) under Grant W911NF-1820181.

REFERENCES

- [1] J. Dean and L. A. Barroso, "The tail at scale," *Communications of the ACM*, vol. 56, no. 2, pp. 74–80, 2013.
- [2] S. Dutta, G. Joshi, S. Ghosh, P. Dube, and P. Nagpurkar, "Slow and stale gradients can win the race: Error-runtime trade-offs in distributed sgd," *arXiv preprint arXiv:1803.01113*, 2018.
- [3] J. Chen, X. Pan, R. Monga, S. Bengio, and R. Jozefowicz, "Revisiting distributed synchronous sgd," *arXiv preprint arXiv:1604.00981*, 2016.
- [4] H. Robbins and S. Monro, "A stochastic approximation method," *The Annals of Mathematical Statistics*, vol. 22, no. 3, pp. 400–407, 1951.
- [5] L. Bottou, F. E. Curtis, and J. Nocedal, "Optimization methods for large-scale machine learning," *SIAM Review*, vol. 60, no. 2, pp. 223–311, 2018.
- [6] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and distributed computation: numerical methods*, vol. 23. Prentice hall Englewood Cliffs, NJ, 1989.
- [7] A. Cotter, O. Shamir, N. Srebro, and K. Sridharan, "Better mini-batch algorithms via accelerated gradient methods," in *Advances in neural information processing systems*, pp. 1647–1655, 2011.
- [8] A. Agarwal and J. C. Duchi, "Distributed delayed stochastic optimization," in *Advances in Neural Information Processing Systems*, pp. 873–881, 2011.
- [9] O. Dekel, R. Gilad-Bachrach, O. Shamir, and L. Xiao, "Optimal distributed online prediction using mini-batches," *Journal of Machine Learning Research*, vol. 13, no. Jan, pp. 165–202, 2012.
- [10] O. Shamir and N. Srebro, "Distributed stochastic optimization and learning," in *Communication, Control, and Computing (Allerton), 2014 52nd Annual Allerton Conference on*, pp. 850–857, IEEE, 2014.
- [11] R. Tandon, Q. Lei, A. G. Dimakis, and N. Karampatziakis, "Gradient coding: Avoiding stragglers in distributed learning," in *International Conference on Machine Learning*, pp. 3368–3376, 2017.
- [12] M. Ye and E. Abbe, "Communication-computation efficient gradient coding," *arXiv preprint arXiv:1802.03475*, 2018.
- [13] N. Raviv, I. Tamo, R. Tandon, and A. G. Dimakis, "Gradient coding from cyclic mds codes and expander graphs," *arXiv preprint arXiv:1707.03858*, 2017.
- [14] K. Lee, M. Lam, R. Pedarsani, D. Papailiopoulos, and K. Ramchandran, "Speeding up distributed machine learning using codes," *IEEE Transactions on Information Theory*, vol. 64, no. 3, pp. 1514–1529, 2018.
- [15] N. Ferdinand and S. Draper, "Anytime stochastic gradient descent: A time to hear from all the workers," *arXiv preprint arXiv:1810.02976*, 2018.
- [16] Q. Yu, N. Raviv, J. So, and A. S. Avestimehr, "Lagrange coded computing: Optimal design for resiliency, security and privacy," *arXiv preprint arXiv:1806.00939*, 2018.
- [17] S. Kiani, N. Ferdinand, and S. C. Draper, "Exploitation of stragglers in coded computation," in *2018 IEEE International Symposium on Information Theory (ISIT)*, pp. 1988–1992, IEEE, 2018.
- [18] L. Chen, Z. Charles, D. Papailiopoulos, *et al.*, "Draco: Robust distributed training via redundant gradients," *arXiv preprint arXiv:1803.09877*, 2018.
- [19] C. Karakus, Y. Sun, S. Diggavi, and W. Yin, "Straggler mitigation in distributed optimization through data encoding," in *Advances in Neural Information Processing Systems*, pp. 5434–5442, 2017.
- [20] W. Halbawi, N. Azizan-Ruhi, F. Salehi, and B. Hassibi, "Improving distributed gradient descent using reed-solomon codes," *arXiv preprint arXiv:1706.05436*, 2017.
- [21] S. Dutta, V. Cadambe, and P. Grover, "Short-dot: Computing large linear transforms distributedly using coded short dot products," in *29th Annual Conference on Neural Information Processing Systems (NIPS)*, pp. 2092–2100, 2016.
- [22] Y. Keshtkarjahreni and H. Seferoglu, "Coded cooperative computation for internet of things," *CoRR*, vol. abs/1801.04357, 2018.

- [23] M. Fahim, H. Jeong, F. Haddadpour, S. Dutta, V. Cadambe, and P. Grover, “On the optimal recovery threshold of coded matrix multiplication,” in *Communication, Control, and Computing (Allerton), 2017 55th Annual Allerton Conference on*, pp. 1264–1270, IEEE, 2017.
- [24] R. Bitar, P. Parag, and S. El Rouayheb, “Minimizing latency for secure distributed computing,” in *International Symposium on Information Theory (ISIT)*, pp. 2900–2904, June 2017.
- [25] Z. Charles, D. Papailiopoulos, and J. Ellenberg, “Approximate gradient coding via sparse random graphs,” *arXiv preprint arXiv:1711.06771*, 2017.
- [26] Z. Charles, D. Papailiopoulos, and J. Ellenberg, “Approximate gradient coding via sparse random graphs,” *arXiv preprint arXiv:1711.06771*, 2017.
- [27] R. K. Maity, A. S. Rawat, and A. Mazumdar, “Robust gradient descent via moment encoding with ldpc codes,” *arXiv preprint arXiv:1805.08327*, 2018.
- [28] 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*, pp. 693–701, 2011.
- [29] J. Liu, S. J. Wright, C. Ré, V. Bittorf, and S. Sridhar, “An asynchronous parallel stochastic coordinate descent algorithm,” *The Journal of Machine Learning Research*, vol. 16, no. 1, pp. 285–322, 2015.
- [30] S. Shalev-Shwartz and T. Zhang, “Accelerated mini-batch stochastic dual coordinate ascent,” in *Advances in Neural Information Processing Systems*, pp. 378–385, 2013.
- [31] S. J. Reddi, A. Hefny, S. Sra, B. Póczos, and A. J. Smola, “On variance reduction in stochastic gradient descent and its asynchronous variants,” in *Advances in Neural Information Processing Systems*, pp. 2647–2655, 2015.
- [32] X. Pan, M. Lam, S. Tu, D. Papailiopoulos, C. Zhang, M. I. Jordan, K. Ramchandran, and C. Ré, “Cyclades: Conflict-free asynchronous machine learning,” in *Advances in Neural Information Processing Systems*, pp. 2568–2576, 2016.
- [33] N. Murata, “A statistical study of on-line learning,” *Online Learning and Neural Networks. Cambridge University Press, Cambridge, UK*, pp. 63–92, 1998.
- [34] J. Chee and P. Toulis, “Convergence diagnostics for stochastic gradient descent with constant learning rate,” in *International Conference on Artificial Intelligence and Statistics*, pp. 1476–1485, 2018.
- [35] G. C. Pflug, “Non-asymptotic confidence bounds for stochastic approximation algorithms with constant step size,” *Monatshefte für Mathematik*, vol. 110, no. 3-4, pp. 297–314, 1990.
- [36] S. L. Smith, P.-J. Kindermans, C. Ying, and Q. V. Le, “Don’t decay the learning rate, increase the batch size,” *arXiv preprint arXiv:1711.00489*, 2017.
- [37] D. Needell and R. Ward, “Batched stochastic gradient descent with weighted sampling,” in *International Conference Approximation Theory*, pp. 279–306, Springer, 2016.

APPENDIX

The result in Proposition 1 holds under the following conditions stated in [2], [5]:

- 1) The loss function $F(\mathbf{w})$ is continuously differentiable and the gradient function of F , given by ∇F , is Lipschitz continuous with Lipschitz constant $L > 0$, i.e.,

$$\|\nabla F(\mathbf{w}_1) - \nabla F(\mathbf{w}_2)\|_2 \leq L \|\mathbf{w}_1 - \mathbf{w}_2\|_2.$$

- 2) The loss function $F(\mathbf{w})$ is strongly convex, i.e., there exists a constant $c > 0$ such that

$$2c(F(\mathbf{w}) - F^*) \leq \|\nabla F(\mathbf{w})\|_2^2 \quad \forall \mathbf{w},$$

where $F^* = F(\mathbf{w}^*)$ is the loss function evaluated at the optimal point $\mathbf{w}^*$.

- 3) The stochastic gradient is an unbiased estimate of the actual gradient and has bounded variance, i.e., for all iterations indexed by $j \in \mathbb{N}$,

$$\mathbb{E}[\hat{\mathbf{g}}_j] = \nabla F(\mathbf{w}_j), \quad \text{Var}[\hat{\mathbf{g}}_j] \leq \frac{\sigma^2}{s} + \frac{M_G}{s} \|\nabla F(\mathbf{w}_j)\|_2^2,$$

where the expectation is taken over the observed rows $\mathbf{a}_\ell$ at the master, and σ and M_G are positive scalars.

The formal statement of the proposition is the following.

Proposition 1 (formal). *Suppose that the loss function $F(\mathbf{w})$ is c -strongly convex and L -lipchitz, and consider a fixed learning rate η such that $\eta \leq 1/[2L(\frac{M_G}{ks} + 1)]$. Then, the error of fastest- k SGD after j iterations satisfies*

$$\mathbb{E}[F(\mathbf{w}_j) - F^*] \leq \frac{\eta L \sigma^2}{2cks} + (1 - \eta c)^j \left(F(\mathbf{w}_0) - F^* - \frac{\eta L \sigma^2}{2cks} \right).$$

Remark 1. *One can easily verify that the gradient estimator given in (5) satisfies the unbiasedness condition. Moreover, the bounded variance condition can be also easily shown when the optimization is done over a convex compact set $\mathcal{W} \subset \mathbb{R}^d$. Furthermore, the least square errors (L2) is an example of a loss function that satisfies the first and second conditions.*