

Slow and Stale Gradients Can Win the Race

Sanghamitra Dutta, Jianyu Wang and Gauri Joshi
Carnegie Mellon University

January 16, 2020

Abstract

Distributed Stochastic Gradient Descent (SGD) when run in a synchronous manner, suffers from delays in *runtime* as it waits for the slowest workers (stragglers). Asynchronous methods can alleviate stragglers, but cause gradient staleness that can adversely affect the convergence *error*. In this work we present a novel theoretical characterization of the speedup offered by asynchronous methods by analyzing the trade-off between the *error* in the trained model and the actual training *runtime* (wallclock time). The main novelty in our work is that our runtime analysis considers random straggling delays, which helps us design and compare distributed SGD algorithms that strike a balance between straggling and staleness. We also provide a new error convergence analysis of asynchronous SGD variants without bounded or exponential delay assumptions. Finally, based on our theoretical characterization of the error-runtime trade-off, we propose a method of gradually varying synchronicity in distributed SGD and demonstrate its performance on CIFAR10 dataset.

1 Introduction

Stochastic gradient descent (SGD) is the backbone of most state-of-the-art machine learning algorithms. Thus, improving the stability and convergence rate of SGD algorithms is critical for making machine learning algorithms fast and efficient. Classical SGD was designed to be run on a single computing node, and its error-convergence with respect to the number of iterations has been extensively analyzed and improved in optimization and learning theory literature. Due to the massive training data-sets and deep neural network architectures used today, running SGD at a single node can be prohibitively slow. This calls for distributed implementations of SGD, where gradient computation and aggregation is parallelized across multiple worker nodes. Although parallelism boosts the amount of data processed per iteration, it exposes SGD to unpredictable node slowdown and communication delays stemming from variability in the computing infrastructure. Thus, there is a critical need to make distributed SGD fast, and yet robust to system variability.

The convergence speed of distributed SGD is a product of two factors: 1) the error in the trained model versus the number of iterations, and 2) the number of iterations completed per second. Traditional single-node SGD analysis focuses on optimizing the first factor, because the second factor is generally a constant when SGD is run on a single dedicated server. In distributed SGD, which is often run on shared cloud infrastructure, the second factor depends on several aspects such as the number of worker nodes, their gradient computation delays, and the protocol (synchronous or asynchronous) used to aggregate their gradients. Hence, in order to achieve the fastest convergence speed we need: 1) optimization techniques to maximize the error-convergence rate with respect to iterations, and 2) scheduling techniques to maximize the number of iterations completed per second. These directions are inter-dependent and need to be explored together rather than in isolation. While many works have advanced the first direction, the second is less explored from a theoretical point

The authors are with the Department of Electrical and Computer Engineering, Carnegie Mellon University. Author Contacts: S. Dutta (sanghamd@andrew.cmu.edu), J. Wang (jianyuw1@andrew.cmu.edu), G. Joshi (gaurij@andrew.cmu.edu)

Some of the results have appeared in AISTATS 2018. This is an extended version with additional results, in particular, an adaptive synchronicity strategy called AdaSync.




Figure 1: (a) The parameter server framework (b) Synchronous SGD has lower error floor but higher runtime, while, asynchronous SGD converges faster but has a higher error floor. We want to achieve the the lower envelope between the two curves which characterizes the best error-runtime trade-off.

of view, and the juxtaposition of both is an unexplored problem. Our goal is to design SGD algorithms that easily lend themselves to distributed implementations, and are robust to fluctuations in computation and network delays as well as unpredictable node failures. This work improves the true convergence speed of distributed SGD with respect to wallclock time by jointly designing scheduling techniques to reduce per-iteration delay, and optimization algorithms to minimize error-versus-iterations.

A commonly used distributed SGD framework, which is first deployed at a large-scale in Google’s DistBelief [1], is the parameter server framework, which consists of a central parameter server (PS) that is used to aggregate gradients computed by worker nodes as shown in Figure 1 (a). In synchronous SGD, the PS waits for all workers to push gradients before it updates the model parameters. Random delays in computation (referred to as straggling) are common in today’s distributed systems as pointed out in the influential work of [2]. Waiting for slow and straggling workers can diminish the speedup offered by parallelizing the training. To alleviate the problem of stragglers, SGD can be run in an asynchronous manner, where the central parameters are updated without waiting for all workers. However, workers may return *stale* gradients that were evaluated at an older version of the model, and this can make the algorithm unstable. Synchronous SGD typically has better convergence error but has a higher wallclock runtime per iteration because it requires synchronization of straggling workers. On the other hand, asynchronous SGD has faster wallclock runtime per iteration but it also has higher convergence error due to the problem of gradient staleness.

Our goal is to achieve the lower envelope of the error-runtime trade-offs achieved by synchronous and asynchronous SGD (see Figure 1(b)), which characterizes the best error-runtime trade-off. Towards achieving this goal, in this work we present a systematic theoretical analysis of the trade-off between error and the actual runtime (instead of iterations), modelling wallclock runtimes as random variables with a general distribution. Based on our analysis, we propose AdaSync, which is a method of adaptively increasing the number of nodes whose gradients are aggregated synchronously by the central PS. Our theoretical results are also substantiated with experiments on CIFAR10 [3] dataset.

1.1 Related Works

Single Node SGD: Analysis of gradient descent dates back to classical works [4] in the optimization community. The problem of interest is the minimization of empirical risk of the form:

$$\min_{\mathbf{w}} \left\{ F(\mathbf{w}) \stackrel{\text{def}}{=} \frac{1}{N} \sum_{n=1}^N f(\mathbf{w}, \xi_n) \right\}. \quad (1)$$

Here, ξ_n denotes the n -th data point and its label where $n = 1, 2, \dots, N$, and $f(\mathbf{w}, \xi_n)$ denotes the composite loss function. Gradient descent is a way to iteratively minimize this objective function by updating the parameter \mathbf{w} in the opposite direction of the gradient of $F(\mathbf{w})$ at every iteration, as given by:

$$\mathbf{w}_{j+1} = \mathbf{w}_j - \eta \nabla F(\mathbf{w}_j) = \mathbf{w}_j - \frac{\eta}{N} \sum_{n=1}^N \nabla f(\mathbf{w}_j, \xi_n).$$

The computation of $\sum_{n=1}^N \nabla f(\mathbf{w}_j, \xi_n)$ over the entire dataset is expensive. Thus, stochastic gradient descent [5] with mini-batching is generally used in practice, where the gradient is evaluated over small, randomly chosen subsets of the data. Smaller mini-batches result in higher variance of the gradients, which affects convergence and error floor [6–8]. Algorithms such as AdaGrad [9] and Adam [10] gradually reduce learning rate to achieve a lower error floor. Another class of algorithms includes stochastic variation reduction techniques that include SVRG [11], SAGA [12] and their variants listed out in [13]. For a detailed survey of different SGD variants, refer to [14].

Synchronous SGD and Stragglers: To process large datasets, SGD is parallelized across multiple workers with a central PS. Each worker processes one mini-batch, and the PS aggregates all the gradients. The convergence of synchronous SGD is same as mini-batch SGD, with a P -fold larger mini-batch, where P is the number of workers. However, the time per iteration grows with the number of workers, because some straggling workers that slow down randomly [2]. Thus, it is important to juxtapose the error reduction per iteration with the runtime per iteration to understand the true convergence speed of distributed SGD.

To deal with stragglers and speed up machine learning, system designers have proposed several straggler mitigation techniques such as [15] that try to detect and avoid stragglers. An alternate direction of work is to use redundancy techniques, e.g., replication or erasure codes, as proposed in [16–37] to deal with the stragglers, as also discussed in Remark 1. See also [38–42] for other interesting related works in this direction.

Asynchronous SGD and Staleness: A complementary approach to deal with the issue of straggling is to use asynchronous SGD. In asynchronous SGD, any worker can evaluate the gradient and update the central PS without waiting for the other workers. Asynchronous variants of existing SGD algorithms have also been proposed and implemented in systems [1, 43–46]. In general, analyzing the convergence of asynchronous SGD with the number of iterations is difficult in itself because of the randomness of gradient staleness. There are only a few pioneering works such as [47–59] in this direction. In [47], a fully decentralized analysis was proposed that considers no central PS. In [50], a new asynchronous algorithm called Hogwild was proposed and analyzed under bounded gradient and bounded delay assumptions. This direction of research has been followed upon by several interesting works such as [48] which proposed novel theoretical analysis under bounded delay assumption for other asynchronous SGD variants. In [55–58], the framework of ARock was proposed for parallel co-ordinate descent and analyzed using Lyapunov functions, relaxing several existing assumptions such as bounded delay assumption and the independence of the delays and the index of the blocks being updated. In algorithms such as Hogwild, ARock etc. every worker only updates a part of the central parameter vector \mathbf{w} at every iteration and are thus essentially different in spirit from conventional asynchronous SGD settings [48, 51] where every worker updates the entire \mathbf{w} . In an alternate direction of work [52], asynchrony is modelled as a perturbation.

In this work, we present a new and simpler analysis of asynchronous SGD with number of iterations that relaxes some of the assumptions in previous literature, and helps us to characterize the error-runtime trade-off as well as easily derive adaptive update rule for gradually increasing synchrony.

1.2 Main Contributions

Existing machine learning algorithms mostly try to optimize the trade-off of error with the number of iterations, epochs or “work complexity” [8], while assuming the time spent per iteration to be a constant. However, due to straggling and synchronization bottle-necks in the system, the same gradient computation task can often take different time to complete across different workers or iterations [2]. This work departs from the classic optimization theory view of analyzing error convergence with respect to the number of

iterations and takes the novel approach of minimizing the error with respect to the wallclock time. By taking a joint runtime and error optimization approach, we provide the first comprehensive runtime-per-iteration comparison of SGD variants and design adaptive synchronous SGD algorithms that can achieve a super-linear runtime speed-up over naive synchronous SGD, while still preserving a low error floor. The main contributions of this paper are summarized below.

- **Straggler-Resilient Variants of Synchronous and Asynchronous SGD.** In order to strike a balance between the two extremes: synchronous and asynchronous SGD, we propose partially synchronous SGD variants such as K -sync, K -batch-sync, K -async and K -batch-async SGD, where K is the number of workers (out of the total of P workers) that the parameter waits for when aggregating gradients. Although some of these distributed SGD variants have been proposed previously, to the best of our knowledge, this is the first work to provide a unified error convergence and runtime analysis of these variants.
- **Runtime Analysis of the Distributed SGD Variants.** We provide the first systematic analysis of the expected runtime per iteration of synchronous and asynchronous SGD and their variants. We do so by modelling the runtimes at each worker as random variables with an arbitrary general distribution. For commonly used delay distributions such as exponential, asynchronous SGD is $O(P \log P)$ times faster than synchronous SGD where P is the total number of workers.
- **More General Error Analysis of Asynchronous SGD Variants.** We propose a new error convergence analysis for asynchronous SGD and its variants for *strongly convex* objectives that can also be extended to provide relaxed guarantees for *non-convex* formulations. In this analysis we relax the bounded delay assumption in [48] and the bounded gradient assumption in [50]. We also remove the assumption of exponential computation time and the staleness process being independent of the parameter values [49] as we will elaborate in Section 4. Interestingly, our analysis also brings out the regimes where asynchronous SGD can be better or worse than synchronous SGD in terms of speed of convergence.
- **Insights from the Error-versus-wallclock Time Trade-off.** By combining our runtime and error analyses described above, we can theoretically characterize the error-versus-wallclock time trade-off for different SGD variants. Figure 2 illustrates the error at convergence (or error floor) versus the time to reach convergence of different SGD variants. Observe how the K -batch-async and K -async strategies can span different points on the trade-off as K varies. By choosing the right value of K we can achieve a desired error at convergence in minimum time. The theoretical results presented in this paper are corroborated by rigorous experiments on training deep neural networks for classification of the CIFAR10 dataset.
- **AdaSync strategy to Adapt Synchronicity during Training.** Instead of fixing K , we can achieve a win-win in the error-runtime trade-off by adapting K so as to gradually increasing the synchrony of the different SGD variants. We propose AdaSync, a method that uses the theoretical characterization of the error-runtime trade-off, to decide how to adapt K , as illustrated in Figure 3. This method is inspired from [60,61] which adapts the communication frequency for a different class of SGD methods known as periodic averaging SGD. Interestingly, similar to [60], our proposed method does not require knowledge of the algorithm parameters such as Lipschitz constant, variance of the stochastic gradient etc. as one would otherwise require if they choose to simply minimize the error-runtime trade-off with respect to parameter K . Experimental results on CIFAR 10 classification (see Figure 3b) show that AdaSync not only helps achieve the same training loss much faster but also gives smaller test error than fixed- K strategies.

The rest of the paper is organized as follows. Section 2 describes our problem formulation introducing the system model and assumptions. Section 3 provides the theoretical results on the analysis of true wallclock runtime per iteration for the different SGD variants and provides insights on quantifying the speedups that one variant provides over another. In Section 4, we discuss our analysis of error convergence with number of


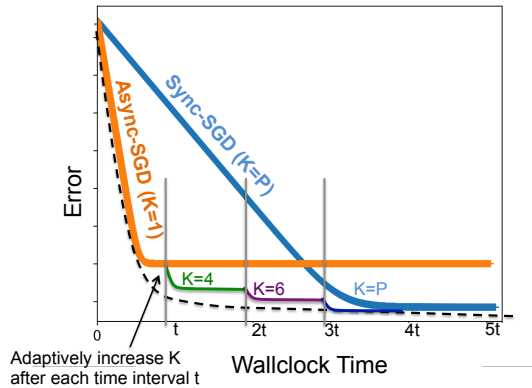



Figure 2: Distributed SGD variants span the error-runtime trade-off between fully Sync-SGD and fully Async-SGD. Here K is the number of workers or mini-batches the PS waits for before updating the model parameters, as we elaborate in Section 2.



(a) Idea behind the AdaSync strategy.



(b) Adaptive K -async for CIFAR 10 classification.

Figure 3: We propose AdaSync, a method of adaptively increasing K which is a measure of the synchronicity of the algorithm. AdaSync aggregates gradients from any K out the P total nodes. It helps achieve the best error-runtime trade-off by gradually increasing K .

iterations where we also include our new convergence analysis for asynchronous and K -async SGD. Proofs and detailed discussions are presented in the Appendix. In Section 5, we combine the runtime analysis with the error analysis to derive novel error-runtime trade-offs and demonstrate how our analysis could inform predicting the trend of the trade-off for distributed systems with different runtime distributions. Finally, in Section 6, we introduce our proposed method AdaSync that gradually varies synchronicity (parameter K) to achieve the desirable error-runtime trade-off, followed by experimental results in Section 6.1. We conclude with a brief discussion in Section 7.

2 Problem Formulation

Our objective is to minimize the risk function of the parameter vector \mathbf{w} as mentioned in (1) given N training samples. Let S denote the total set of N training samples, *i.e.*, a collection of some data points with their corresponding labels or values. We use the notation ξ to denote a random seed $\in S$ which consists of either a single data and its label or a single mini-batch (m samples) of data and their labels.




Figure 4: For $K = 2$ and $P = 3$, we illustrate the K -sync and K -batch-sync SGD in comparison with fully synchronous SGD. Lightly shaded arrows indicate straggling gradient computations that are cancelled.

2.1 System Model

We assume that there is a central parameter server (PS) with P parallel workers as shown in Section 1. The workers fetch the current parameter vector \mathbf{w}_j from the PS as and when instructed in the algorithm. Then they compute gradients using one mini-batch and push their gradients back to the PS as and when instructed in the algorithm. At each iteration, the PS aggregates the gradients computed by the workers and updates the parameter \mathbf{w} . Based on how these gradients are fetched and aggregated, we have different variants of synchronous or asynchronous SGD.

2.2 Variants of SGD

We now describe the SGD variants considered in this paper. We note that some of these variants have been proposed earlier under alternate names in different papers, as we will refer to during our descriptions. In this work, we give a unified runtime and error analysis to compare them with each other in terms of their true error-runtime trade-off, a problem that has not been considered in prior works. Please refer to Figure 4 and Figure 5 for a pictorial illustration of the SGD variants.

K -sync SGD: This is a generalized form of synchronous SGD, also suggested in [43, 62] to offer some resilience to straggling as the PS does not wait for all the workers to finish. The PS only waits for the first K out of P workers to push their gradients. Once it receives K gradients, it updates \mathbf{w}_j and cancels the remaining workers. The updated parameter vector \mathbf{w}_{j+1} is sent to all P workers for the next iteration. The update rule is given by:

$$\mathbf{w}_{j+1} = \mathbf{w}_j - \frac{\eta}{K} \sum_{l=1}^K g(\mathbf{w}_j, \xi_{l,j}). \quad (2)$$

Here $l = 1, 2, \dots, K$ denotes the index of the K workers that finish first, $\xi_{l,j}$ denotes the mini-batch of m samples used by the l -th worker at the j -th iteration and $g(\mathbf{w}_j, \xi_{l,j}) = \frac{1}{m} \sum_{\xi \in \xi_{l,j}} \nabla f(\mathbf{w}_j, \xi)$ denotes the average gradient of the loss function evaluated over the mini-batch $\xi_{l,j}$ of size m . For $K = P$, the algorithm is exactly equivalent to a fully synchronous SGD with P workers.

K -batch-sync SGD: In K -batch-sync, all the P workers start computing gradients with the same \mathbf{w}_j . Whenever any worker finishes, it pushes its update to the PS and evaluates the gradient on the next mini-batch at the same \mathbf{w}_j . The PS updates using the first K mini-batches that finish and cancels the remaining workers. Theoretically, the update rule is still the same as (2) but here l now denotes the index of the mini-batch (out of the K mini-batches that finished first) instead of the worker. However K -batch-sync will offer advantages over K -sync in runtime per iteration as no worker is idle.

K -async SGD: This is a generalized version of asynchronous SGD, also suggested in [43]. In K -async SGD, all the P workers compute their respective gradients on a single mini-batch. The PS waits for the first K out of P that finish first, but it does not cancel the remaining workers. As a result, for every update the gradients returned by each worker might be computed at a stale or older value of the parameter \mathbf{w} . The




Figure 5: For $K = 2$ and $P = 3$, we illustrate the K -async and K -batch-async algorithms in comparison with fully asynchronous SGD.

update rule is thus given by:

$$\mathbf{w}_{j+1} = \mathbf{w}_j - \frac{\eta}{K} \sum_{l=1}^K g(\mathbf{w}_{\tau(l,j)}, \xi_{l,j}). \quad (3)$$

Here $l = 1, 2, \dots, K$ denotes the index of the K workers that contribute to the update at the corresponding iteration, $\xi_{l,j}$ is one mini-batch of m samples used by the l -th worker at the j -th iteration and $\tau(l, j)$ denotes the iteration index when the l -th worker last read from the central PS where $\tau(l, j) \leq j$. Also, $g(\mathbf{w}_{\tau(l,j)}, \xi_{l,j}) = \frac{1}{m} \sum_{\xi \in \xi_{l,j}} \nabla f(\mathbf{w}_{\tau(l,j)}, \xi_{l,j})$ is the average gradient of the loss function evaluated over the mini-batch $\xi_{l,j}$ based on the stale value of the parameter $\mathbf{w}_{\tau(l,j)}$. For $K = 1$, the algorithm is exactly equivalent to fully asynchronous SGD, and the update rule can be simplified as:

$$\mathbf{w}_{j+1} = \mathbf{w}_j - \eta g(\mathbf{w}_{\tau(j)}, \xi_j). \quad (4)$$

Here ξ_j denotes the set of samples used by the worker that updates at the j -th iteration such that $|\xi_j| = m$ and $\tau(j)$ denotes the iteration index when that particular worker last read from the central PS. Note that $\tau(j) \leq j$.

K -batch-async SGD: Observe in Figure 5 that K -async also suffers from some workers being idle while others are still working on their gradients until any K finish. In K -batch-async (proposed in [48]), the PS waits for K mini-batches before updating itself but irrespective of which worker they come from. So wherever any worker finishes, it pushes its gradient to the PS, fetches current parameter at PS and starts computing gradient on the next mini-batch based on the current value of the PS. Surprisingly, the update rule is again similar to (3) theoretically except that now l denotes the indices of the K mini-batches that finish first instead of the workers and $\mathbf{w}_{\tau(l,j)}$ denotes the version of the parameter when the worker computing the l -th mini-batch last read from the PS. While the error convergence of K -batch-async is similar to K -async, it reduces the runtime per iteration as no worker is idle.

Remark 1. Recent works such as [21] propose erasure coding techniques to overcome straggling workers. Instead, the SGD variants considered in this paper such as K -sync and K -batch-sync SGD exploit the inherent redundancy in the data itself, and ignore the gradients returned by straggling workers. If the data is well-shuffled such that it can be assumed to be i.i.d. across workers, then for the same effective batch-size, ignoring straggling gradients will give equivalent error scaling as coded strategies, and at a lower computing cost. However, coding strategies may be useful in the non i.i.d. case, when the gradients supplied by each worker provide diverse information that is important to capture in the trained model.

2.3 Performance Metrics and Goal

There are two metrics of interest: Expected Runtime and Error.

Definition 1 (Expected Runtime per iteration). *The expected runtime per iteration is the expected time (average time) taken to perform each iteration, i.e., the expected time between two consecutive updates of the parameter \mathbf{w} at the central PS.*

CONSTANTS		RANDOM VARIABLES	
Mini-batch Size	m	Runtime of a worker for one mini-batch	X_i
Total Iterations	J	Runtime per iteration	T
Number of workers (Processors)	P		
Number of workers to wait for	K		
Learning rate	η		
Lipschitz Constant	L		
Strong-convexity parameter	c		

Table 1: LIST OF NOTATIONS

Definition 2 (Expected Error after J iterations). *The expected error after J iterations is defined as $\mathbb{E}[F(\mathbf{w}_J) - F^*]$, i.e., the expected gap of the risk function from its optimal value.*

Goal: Our goal is to determine the trade-off between the expected error (measures the accuracy of the algorithm) and the expected runtime **after a total of J iterations** for the different SGD variants. Based on our characterization, we would like to derive a method of gradually varying synchronicity (parameter K) in the SGD variants to achieve a desirable error-runtime trade-off.

In Table 1, we provide a list of the notations used in this paper for referencing, that we again revisit in the respective sections where they appear.

3 Runtime Analysis: Insights on Quantifying Speedup

Our runtime analysis provides useful insights in quantifying the speedup offered by different SGD variants. We first state our key modeling assumptions in Section 3.1, followed by our main theoretical results on quantifying speedup in Section 3.2. Next, we include our detailed runtime analysis for the four SGD variants considered in this paper in Section 3.3, some of which are useful in the proofs of the main results on speedups. For a summary of the expected runtime of the different SGD variants, we refer to Table 2.

Table 2: Expected Runtime for the different variants of SGD

SGD Variant	Expected Runtime per iteration $\mathbb{E}[T]$
K -sync	$\mathbb{E}[T] = \mathbb{E}[X_{K:P}]$ for all distributions
K -batch-sync	$\mathbb{E}[T] \leq K\mathbb{E}[X_{1:P}]$ for new-longer-than-used distributions
K -async	$\mathbb{E}[T] \leq \mathbb{E}[X_{K:P}]$ for new-longer-than-used distributions
K -batch-async	$\mathbb{E}[T] = \frac{K}{P}\mathbb{E}[X]$ for all distributions

3.1 Modeling Assumptions:

The time taken by a worker to compute gradient of one mini-batch is denoted by random variable X_i for $i = 1, 2, \dots, P$. We assume that these X_i 's are i.i.d. across mini-batches and workers.

3.2 Main Results on Quantifying Speedups

Our first result (Theorem 1) analytically characterizes the speedup offered by asynchronous SGD for *any general distribution on the wallclock time of each worker*.




Figure 6: Simulation demonstrating the speedup using asynchronous over synchronous SGD for three different types of distributions, namely, $Exp(1)$, $1 + Exp(1)$ and $Pareto(2, 1)$. The speedup is shown in log-scale, *i.e.*, $\log \frac{\mathbb{E}[T_{Sync}]}{\mathbb{E}[T_{Async}]}$ is plotted with total number of nodes P .

Theorem 1. *Let the wallclock time of each worker to process a single mini-batch be i.i.d. random variables X_1, X_2, \dots, X_P . Then the ratio of the expected runtimes per iteration for synchronous and asynchronous SGD is*

$$\frac{\mathbb{E}[T_{Sync}]}{\mathbb{E}[T_{Async}]} = P \frac{\mathbb{E}[X_{P:P}]}{\mathbb{E}[X]}$$

where $X_{(P:P)}$ is the P^{th} order statistic of P i.i.d. random variables X_1, X_2, \dots, X_P .

Proof of Theorem 1. Note that fully synchronous SGD is actually K -sync SGD with $K = P$, *i.e.*, waiting for all the P workers to finish. On the other hand, fully asynchronous SGD is actually K -batch-async with $K = 1$. By taking the ratio of the expected runtimes per iteration for K -sync SGD (see Lemma 1 in Section 3.3) with $K = P$ and K -batch-async (see Lemma 5 in Section 3.3) with $K = 1$, we get the result in Theorem 1. \square

In the following corollary, we highlight this speedup for the special case of exponential computation time.

Corollary 1. *Let the wallclock time of each worker to process a single mini-batch be i.i.d. exponential random variables $X_1, X_2, \dots, X_P \sim Exp(\mu)$. Then the ratio of the expected runtimes per iteration for synchronous and asynchronous SGD is $\Theta(P \log P)$.*

Thus, the speedup scales with P and can diverge to infinity for large P . We illustrate the speedup for different distributions in Figure 6. It might be noted that a similar speedup as Corollary 1 has also been obtained in a recent work [56] under exponential assumptions.

Proof of Corollary 1. The expectation of the maximum of P i.i.d. $X_i \sim Exp(\mu)$ is $\mathbb{E}[X_{P:P}] = \sum_{i=1}^P \frac{1}{i\mu} \approx \frac{\log P}{\mu}$ [63]. This can be substituted in Theorem 1 to get Corollary 1. \square

The next result illustrates the advantages offered by K -batch-sync and K -batch-async over their corresponding counterparts K -sync and K -async respectively.

Theorem 2. *Let the wallclock time of each worker to process a single mini-batch be i.i.d. exponential random variables $X_1, X_2, \dots, X_P \sim Exp(\mu)$. Then the ratio of the expected runtimes per iteration for K -async (or sync) SGD and K -batch-async (or sync) SGD is*

$$\frac{\mathbb{E}[T_{K-async}]}{\mathbb{E}[T_{K-batch-async}]} = \frac{P\mathbb{E}[X_{K:P}]}{K\mathbb{E}[X]} \approx \frac{P \log \frac{P}{P-K}}{K}$$

where $X_{K:P}$ is the K^{th} order statistic of i.i.d. random variables X_1, X_2, \dots, X_P .

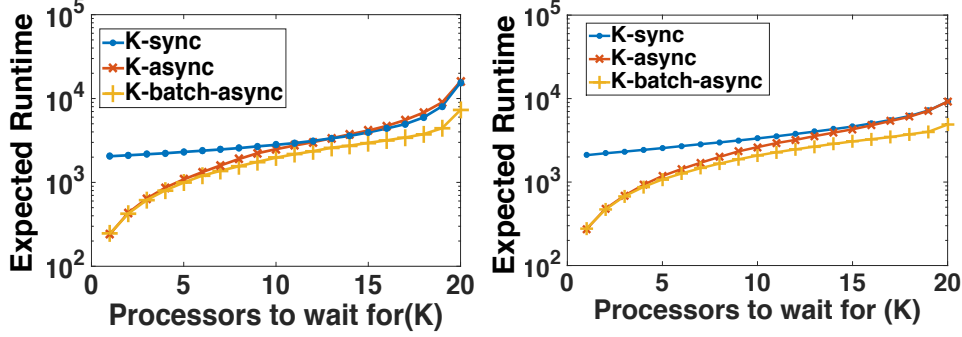


Figure 7: Simulation of expected runtime for 2000 iterations for K -sync, K -async and K -batch-async SGD: (Left) Pareto distribution $Pareto(2, 1)$ and (Right) Shifted exponential distribution $1 + Exp(1)$.

Proof of Theorem 2. For the exponential X_i , equality holds in (9) in Lemma 3, as we justify in Appendix B.3.1. The expectation can be derived as $\mathbb{E}[X_{K:P}] = \sum_{i=P-K+1}^P \frac{1}{i\mu} \approx \frac{\log(P/P-K)}{\mu}$. For exponential X_i , the expected runtime per iteration for K -batch-async is given by $\mathbb{E}[T] = \frac{K\mathbb{E}[X]}{P} = \frac{K}{\mu P}$ from Lemma 5. \square

Theorem 2 shows that as $\frac{K}{P}$ increases, the speedup using K -batch-async increases and can be up to $\log P$ times higher. For non-exponential distributions, we simulate the behaviour of expected runtime in Figure 7 for K -sync, K -async and K -batch-async respectively for Pareto and Shifted Exponential.

3.3 Runtime Analysis for four different SGD variants

Here, we rigorously analyze the theoretical wallclock runtime of the four different SGD variants. These results have also been used in providing insights on the speedup offered by different asynchronous and batch variants in Theorem 1 and Theorem 2.

3.3.1 Runtime of K -sync SGD

Lemma 1 (Runtime of K -sync SGD). *The expected runtime per iteration for K -sync SGD is,*

$$\mathbb{E}[T] = \mathbb{E}[X_{K:P}] \quad (5)$$

where $X_{K:P}$ is the K^{th} order statistic of P i.i.d. random variables X_1, X_2, \dots, X_P .

Proof of Lemma 1. When all the workers start together, and we wait for the first K out of P i.i.d. random variables to finish, the expected computation time for that iteration is $\mathbb{E}[X_{K:P}]$, where $X_{K:P}$ denotes the K -th statistic of P i.i.d. random variables X_1, X_2, \dots, X_P . \square

Thus, for a total of J iterations, the expected runtime is given by $J\mathbb{E}[X_{K:P}]$.

Remark 2. For $X_i \sim Exp(\mu)$, the expected runtime per iteration is given by,

$$\mathbb{E}[T] = \frac{1}{\mu} \sum_{i=P-K+1}^P \frac{1}{i} \approx \frac{1}{\mu} \left(\frac{\log \frac{P}{P-K}}{\mu} \right)$$

where the last step uses an approximation from [63]. For justification, the reader is referred to Appendix B.1.

3.3.2 Runtime of K-batch-sync SGD

The expected runtime of K -batch-sync SGD is not analytically tractable in general, but we obtain an upper bound on it for a class of distributions called the “new-longer-than-used” distributions, as defined below.

Definition 3 (New-longer-than-used). *A random variable is said to have a new-longer-than-used distribution if the following holds for all $t, u \geq 0$:*

$$\Pr(U > u + t | U > t) \leq \Pr(U > u). \quad (6)$$

Most of the continuous distributions we encounter like normal, shifted-exponential, gamma, beta are new-longer-than-used. Alternately, the hyper-exponential distribution is new-shorter-than-used and it satisfies

$$\Pr(U > u + t | U > t) \geq \Pr(U > u) \quad \forall t, u \geq 0. \quad (7)$$

For the exponential distribution, the inequality (6) holds with equality due to the memoryless property, and thus it can be thought of as both new-longer-than-used and new-shorter-than-used.

Lemma 2 (Runtime of K -batch-sync SGD). *Suppose that each X_i has a new-longer-than-used distribution. Then, the expected runtime per iteration for K -batch-sync is upper-bounded as*

$$\mathbb{E}[T] \leq K \mathbb{E}[X_{1:P}] \quad (8)$$

where $X_{1:P}$ is the minimum of P i.i.d. random variables X_1, X_2, \dots, X_P .

The proof is provided in Appendix B.2. For the special case of $X_i \sim \text{Exp}(\mu)$, the runtime per iteration is distributed as $\text{Erlang}(K, P\mu)$ (see Appendix B.2). Thus, for K -batch-sync SGD, the expected time per iteration is given by,

$$\mathbb{E}[T] = \frac{K}{P\mu},$$

which is precisely what we obtain when (8) holds with equality.

3.3.3 Runtime of K-Async SGD

The expected runtime per iteration of K -async SGD is also not analytically tractable for non-exponential X_i , but we again obtain an upper bound on it for “new-longer-than-used” distributions.

Lemma 3 (Runtime of K -async SGD). *Suppose that each X_i has a new-longer-than-used distribution. Then, the expected runtime per iteration for K -async is upper-bounded as*

$$\mathbb{E}[T] \leq \mathbb{E}[X_{K:P}] \quad (9)$$

where $X_{K:P}$ is the K^{th} order statistic of P i.i.d. random variables X_1, X_2, \dots, X_P .

The proof of this lemma is provided in Appendix B.3.

Remark 3. *Recall that the runtime of K -sync SGD is $\mathbb{E}[X_{K:P}]$. Therefore, Lemma 3 essentially implies that for new-longer-than-used distributions, the runtime of K -async SGD is upper-bounded by the runtime of K -sync SGD. For the special case of exponential runtimes, (6) holds with equality, and the expected runtime of both K -sync SGD and K -async SGD match theoretically. However, for other classes of distributions where (6) holds with strict inequality, the upper bound of Lemma 3 also holds with strict inequality. In general, intuitively the “more” is the new-longer-than-used property, the lower is the runtime of K -async SGD as compared to K -sync SGD. We show this by explicitly deriving an alternate upper bound for the shifted-exponential distribution (a special case of the new-longer-than-used distributions) that is lower than the expected runtime of K -sync SGD when the shift of the distribution is large.*

Lemma 4 (Alternate Upper Bound for Shifted Exponential). *Let $P = nK$ where n is an integer greater than 1, and each X_i follow a shifted exponential distribution with shift Δ , i.e., $X_i \sim \Delta + \text{Exp}(\mu)$. Then, the expected runtime for any n consecutive iterations of K -async SGD is upper-bounded as*

$$\mathbb{E}[T_1 + T_2 + \dots + T_n] \leq \Delta + \sum_{i=0}^{n-1} \mathbb{E} \left[\tilde{X}_{K:(P-iK)} \right], \quad (10)$$

where $\tilde{X} \sim \text{Exp}(\mu)$ and $\tilde{X}_{K:(P-iK)}$ denotes the K^{th} order statistic out of $P - iK$ i.i.d. random variables.

The proof of this lemma is also provided in Appendix B.3. Based on this lemma, the upper-bound on per-iteration runtime can be approximated as:

$$\begin{aligned} & \frac{1}{n} \left(\Delta + \sum_{i=0}^{n-1} \mathbb{E} \left[\tilde{X}_{K:(P-iK)} \right] \right) \\ & \approx \frac{\Delta}{n} + \frac{1}{n\mu} \left(\log \frac{P}{P-K} + \log \frac{P-K}{P-2K} + \dots + \log \frac{P-(n-2)K}{P-(n-1)K} \right) + \frac{1}{n\mu} \log(K) \\ & \approx \frac{\Delta}{n} + \frac{\log P}{n\mu} \\ & = \frac{K\Delta}{P} + \frac{K \log P}{P\mu}. \end{aligned} \quad (11)$$

In comparison, the runtime per iteration for K -sync SGD is $\left(\Delta + \frac{\log P/(P-K)}{\mu} \right)$. Thus, a high value of Δ implies that the runtime of K -async SGD is strictly lower than K -sync SGD.

3.3.4 Runtime of K-batch-async SGD

For this variant, we derive an expression that holds for any general distribution on X_i .

Lemma 5 (Runtime of K -batch-async SGD). *The expected runtime per iteration for K -batch-async SGD in the limit of large number of iterations is given by:*

$$\mathbb{E}[T] = \frac{K\mathbb{E}[X]}{P}. \quad (12)$$

Unlike the results for the synchronous variants, this result on average runtime per iteration holds only in the limit of large number of iterations. To prove the result we use ideas from renewal theory. For a brief background on renewal theory, the reader is referred to Appendix B.4.

Proof of Lemma 5. For the i -th worker, let $\{N_i(t), t > 0\}$ be the number of times the i -th worker pushes its gradient to the PS over in time t . The time between two pushes is an independent realization of X_i . Thus, the inter-arrival times $X_i^{(1)}, X_i^{(2)}, \dots$ are i.i.d. with mean inter-arrival time $\mathbb{E}[X_i]$. Using the elementary renewal theorem [64, Chapter 5] we have,

$$\lim_{t \rightarrow \infty} \frac{\mathbb{E}[N_i(t)]}{t} = \frac{1}{\mathbb{E}[X_i]}. \quad (13)$$

Thus, the rate of gradient pushes by the i -th worker is $1/\mathbb{E}[X_i]$. As there are P workers, we have a superposition of P renewal processes and thus the average rate of gradient pushes to the PS is

$$\lim_{t \rightarrow \infty} \sum_{i=1}^P \frac{\mathbb{E}[N_i(t)]}{t} = \sum_{i=1}^P \frac{1}{\mathbb{E}[X_i]} = \frac{P}{\mathbb{E}[X]}. \quad (14)$$

Every K pushes are one iteration. Thus, the expected runtime per iteration or effectively the expected time for K pushes is given by $\mathbb{E}[T] = \frac{K\mathbb{E}[X]}{P}$. \square

4 Error Analysis: New Convergence Analysis for Asynchronous SGD

In this section, we discuss our analysis of error convergence with the number of iterations. We first state some assumptions on the objective function in Section 4.1, followed by the main theoretical results on the error analysis including our novel analysis for asynchronous SGD and its variants in Section 4.2.

4.1 Assumptions on the Objective Function

Closely following [8], we make the following assumptions:

1. $F(\mathbf{w})$ is an L -smooth function. Thus,

$$\|\nabla F(\mathbf{w}) - \nabla F(\tilde{\mathbf{w}})\|_2 \leq L\|\mathbf{w} - \tilde{\mathbf{w}}\|_2 \quad \forall \mathbf{w}, \tilde{\mathbf{w}}. \quad (15)$$

2. $F(\mathbf{w})$ is strongly convex with parameter c . Thus,

$$2c(F(\mathbf{w}) - F^*) \leq \|\nabla F(\mathbf{w})\|_2^2 \quad \forall \mathbf{w}. \quad (16)$$

Refer to Appendix A for discussion on strong convexity. Our results also extend to non-convex objectives, as discussed in Theorem 4.

3. The stochastic gradient is an unbiased estimate of the true gradient:

$$\mathbb{E}_{\xi_j|\mathbf{w}_k} [g(\mathbf{w}_k, \xi_j)] = \nabla F(\mathbf{w}_k) \quad \forall k \leq j. \quad (17)$$

Observe that this is slightly different from the common assumption that says $\mathbb{E}_{\xi_j} [g(\mathbf{w}, \xi_j)] = \nabla F(\mathbf{w})$ for all \mathbf{w} . Observe that all \mathbf{w}_j for $j > k$ is actually not independent of the data ξ_j . We thus make the assumption more rigorous by conditioning on \mathbf{w}_k for $k \leq j$. Our requirement $k \leq j$ means that \mathbf{w}_k is the value of the parameter at the PS before the data ξ_j was accessed and can thus be assumed to be independent of the data ξ_j .

4. Similar to the previous assumption, we also assume that the variance of the stochastic update given \mathbf{w}_k at iteration k before the data point was accessed is also bounded as follows:

$$\mathbb{E}_{\xi_j|\mathbf{w}_k} [\|g(\mathbf{w}_k, \xi_j) - \nabla F(\mathbf{w}_k)\|_2^2] \leq \frac{\sigma^2}{m} + \frac{M_G}{m} \|\nabla F(\mathbf{w}_k)\|_2^2 \quad \forall k \leq j. \quad (18)$$

5. In the analysis of K -async and K -batch-async SGD, we replace some assumptions in existing literature that we discuss in Section 4.2, and instead use an alternate staleness bound that allows for large, but rare delays. We assume that for some $\gamma \leq 1$,

$$\mathbb{E} [\|\nabla F(\mathbf{w}_j) - \nabla F(\mathbf{w}_{\tau(l,j)})\|_2^2] \leq \gamma \mathbb{E} [\|\nabla F(\mathbf{w}_j)\|_2^2]. \quad (19)$$

4.2 Main Theoretical Results

In this work, we provide a novel convergence analysis of K -async SGD for fixed η , relaxing the following assumptions in existing literature.

- In several prior works such as [20, 22, 49, 56], it is often assumed, for the ease of analysis, that runtimes are exponentially distributed. In this paper, we extend our analysis for any general service time X_i .
- In [49], it is also assumed that the staleness process is independent of \mathbf{w} . While this assumption simplifies the analysis greatly, it is not true in practice. For instance, for a two worker case, the parameter \mathbf{w}_2 after 2 iterations depends on whether the update from \mathbf{w}_1 to \mathbf{w}_2 was based on a stale gradient at \mathbf{w}_0 or the current gradient at \mathbf{w}_1 , depending on which worker finished first. In this work, we remove this independence assumption.

- Instead of the bounded delay assumption in [48], we use a general staleness bound

$$\mathbb{E} [\|\nabla F(\mathbf{w}_j) - \nabla F(\mathbf{w}_{\tau(l,j)})\|_2^2] \leq \gamma \mathbb{E} [\|\nabla F(\mathbf{w}_j)\|_2^2]$$

which allows for large, but rare delays.

- In [50], the norm of the gradient is assumed to be bounded. However, if we assume that $\|\nabla F(\mathbf{w})\|_2^2 \leq M$ for some constant M , then using (16) we obtain $\|\mathbf{w} - \mathbf{w}^*\|_2^2 \leq \frac{2}{c}(F(\mathbf{w}) - F^*) \leq \frac{M}{c^2}$ implying that \mathbf{w} itself is bounded which is a very strong and restrictive assumption, that we relax in this result.

Some of these assumptions have been addressed in the context of alternative asynchronous SGD variants in the recent works of [56–59].

4.2.1 Convex Loss Function

Theorem 3. *Suppose the objective $F(\mathbf{w})$ is c -strongly convex and the learning rate $\eta \leq \frac{1}{2L(\frac{M_G}{Km} + \frac{1}{K})}$. Also assume that for some $\gamma \leq 1$,*

$$\mathbb{E} [\|\nabla F(\mathbf{w}_j) - \nabla F(\mathbf{w}_{\tau(l,j)})\|_2^2] \leq \gamma \mathbb{E} [\|\nabla F(\mathbf{w}_j)\|_2^2].$$

Then, the error of K -async SGD after J iterations is,

$$\mathbb{E}[F(\mathbf{w}_J)] - F^* \leq \frac{\eta L \sigma^2}{2c\gamma' Km} + (1 - \eta c\gamma')^J \left(F(\mathbf{w}_0) - F^* - \frac{\eta L \sigma^2}{2c\gamma' Km} \right) \quad (20)$$

where $\gamma' = 1 - \gamma + \frac{p_0}{2}$ and p_0 is a lower bound on the conditional probability that $\tau(l, j) = j$, given all the past delays and parameters.

Here, γ is a measure of staleness of the gradients returned by workers; smaller γ indicates less staleness. The full proof is provided in Appendix C. We first prove the result for $K = 1$ in Appendix C.2 for ease of understanding, and then provide the more general proof for any K in Appendix C.3.

Lemma 6 below provides bounds on p_0 for different classes of distributions.

Lemma 6 (Bounds on p_0). *Define $p_0 = \inf_j p_0^{(j)}$. Then the following holds:*

- For exponential computation times, $p_0^{(j)} = \frac{1}{\bar{p}}$ for all j (invariant of j) and $p_0 = \frac{1}{\bar{p}}$.
- For new-longer-than-used (See Definition 3) computation times, $p_0^{(j)} \leq \frac{1}{\bar{p}}$ and thus $p_0 \leq \frac{1}{\bar{p}}$.
- For new-shorter-than-used computation times, $p_0^{(j)} \geq \frac{1}{\bar{p}}$ and thus $p_0 \geq \frac{1}{\bar{p}}$.

The proof is provided in Appendix C.1.

Remark 4. *For K -batch-async, the update rule is same as K -async except that the index l denotes the index of the mini-batch. Thus, the error analysis will be exactly similar.*

Now let us compare with K -sync SGD. We observe that the analysis of K -sync SGD is same as serial SGD with mini-batch size Km . Thus,

Lemma 7 (Error of K -sync [8]). *Suppose that the objective $F(\mathbf{w})$ is c -strongly convex and learning rate $\eta \leq \frac{1}{2L(\frac{M_G}{Km} + 1)}$. Then, the error after J iterations of K -sync SGD is*

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

Can stale gradients win the race? For the same η , observe that the error given by Theorem 3 decays at the rate $(1 - \eta c(1 - \gamma + \frac{p_0}{2}))$ for K -async or K -batch-async SGD while for K -sync, the decay rate with number of iterations is $(1 - \eta c)$. Thus, depending on the values of γ and p_0 , the decay rate of K -async or K -batch-async SGD can be faster or slower than K -sync SGD. The decay rate of K -async or K -batch-async SGD is faster if $\frac{p_0}{2} > \gamma$. As an example, one might consider an exponential or new-shorter-than-used service time where $p_0 \geq \frac{1}{P}$ and γ can be made smaller by increasing K . It might be noted that asynchronous SGD can still be faster than synchronous SGD with respect to wallclock time even if its decay rate with respect to number of iterations is lower as every iteration is much faster in asynchronous SGD (Roughly $P \log P$ times faster for exponential service times).

The maximum allowable learning rate for synchronous SGD is $\max\{\frac{1}{c}, \frac{1}{2L(\frac{MG}{Pm} + 1)}\}$ which can be much higher than that for asynchronous SGD, *i.e.*, $\max\{\frac{1}{c(1-\gamma+\frac{p_0}{2})}, \frac{1}{2L(\frac{MG}{m} + 1)}\}$. Similarly the error floor for synchronous is $\frac{\eta L \sigma^2}{2cPm}$ as compared to asynchronous whose error floor is $\frac{\eta L \sigma^2}{2c(1-\gamma+\frac{p_0}{2})m}$.

4.2.2 Extension to Non-Convex Loss Function

The analysis can be extended to provide weaker guarantees for non-convex objectives. Let $\gamma' = 1 - \gamma + \frac{p_0}{2}$. For non-convex objectives, we have the following result.

Theorem 4. *For non-convex objective function $F(\cdot)$, where $F^* = \min_{\mathbf{w}} F(\mathbf{w})$, we have the following ergodic convergence result for K -async SGD:*

$$\frac{1}{J} \sum_{j=0}^{J-1} \mathbb{E} [\|\nabla F(\mathbf{w}_j)\|_2^2] \leq \frac{2(F(\mathbf{w}_0) - F^*)}{J\eta\gamma'} + \frac{L\eta\sigma^2}{Km\gamma'}. \quad (21)$$

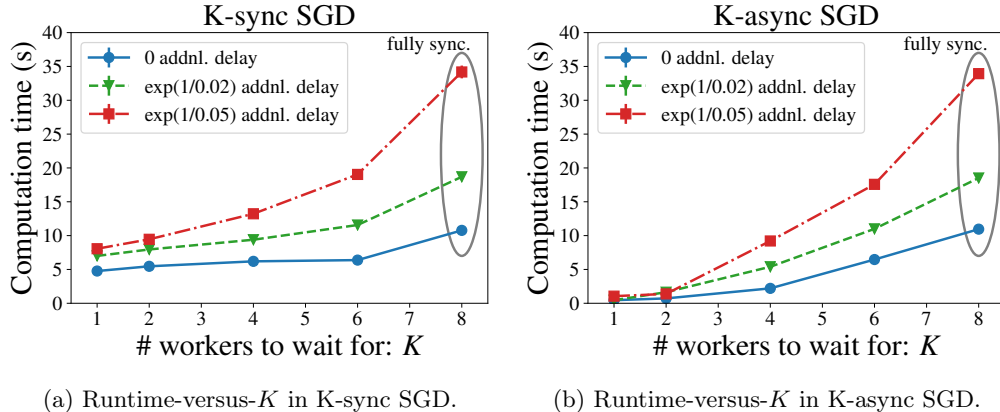
The proof is provided in Appendix C.4. As before, the same analysis also holds for K -batch-async SGD. For K -sync and K -batch-sync, we can also obtain a similar result, substituting $\gamma' = 1$ in (21) (see [8]). Next, we combine our runtime analysis with the error analysis to characterize the error-runtime trade-off.

5 Experiments and Insights on the Error-Runtime Trade-off

We can combine our expressions for runtime per iteration with the error convergence per iteration to derive the error-runtime trade-off. In Figure 1(b), we compare the theoretical trade-offs between synchronous ($K = P$ in Lemma 7 and Lemma 1) and asynchronous SGD ($K = 1$ in Theorem 3 and Lemma 3) under the strongly-convex assumption. Asynchronous SGD converges very quickly, but to a higher floor. On the other hand, synchronous SGD converges slowly with respect to time, but reaches a much lower error-floor. To validate the trend observed in theory, we conduct experiments on training neural networks to perform image classification on CIFAR 10 dataset. These experiments give insights on how choosing the right K helps us strike the best error-runtime trade-off, depending upon the distribution of the gradient computation delays.

5.1 Experimental Setting

The algorithms discussed in this paper are implemented in Pytorch (v1.0) using multiple nodes. Ray (v0.7) is used for supporting the distributed execution. We use the CIFAR-10 [3] dataset. This dataset consists of 60,000 32×32 color images in 10 classes. There are 50,000 training images and 10,000 validation images. The neural network used to classify this dataset has two convolutional layers and three fully connected layers. Experiments were conducted on a local cluster with 8 worker machines, each of which has an NVIDIA TitanX GPU. Machines are connected via a 40 Gbps (5000 Mb/s) Ethernet interface. Mini-batch size per worker machine is 32 and learning rate is $\eta = 0.12$.



(a) Runtime-versus- K in K -sync SGD.

(b) Runtime-versus- K in K -async SGD.

Figure 8: Runtime per iteration changes along with the parameter K for CIFAR10 dataset: K -async SGD always has lower runtime than K -sync SGD as one might expect from our theoretical analysis.

5.2 Speedup in Runtime

In Figure 8(a) and Figure 8(b) we compare the average runtime per epoch of K -sync and K -async SGD for different values of K . When $K = P = 8$, both the variants become identical to fully synchronous SGD. As we decrease K from $P = 8$ to 1, the computation time drastically reduces since we do not have to wait for straggling nodes. K -async SGD gives a larger delay reduction than K -sync SGD because we do not cancel partially completed gradient computation tasks. Each plot shows three cases: 1) with no artificial delays added to induce straggling, 2) with an additional exponential delay with mean 0.02sec, and 3) with an additional exponential delay with mean 0.05sec. The purpose of these curves is to demonstrate how variability in gradient computation time affects the runtime per iteration. Higher variability means that the system is more susceptible to straggling workers. Thus, as the delay variability increases (as we add a higher mean exponential delay per worker), setting a smaller K gives sharper delay reduction as compared to the $K = 8$ (fully synchronous SGD) case.

5.3 Accuracy-Runtime Trade-off in K -sync and K -async SGD

In this subsection, we examine the accuracy-runtime trade-offs for K -sync and K -async SGD variants for CIFAR10 dataset. We plot the test error against wallclock runtime for three cases: (a) No artificial delay; (b) Exponential delay with mean 0.02sec; and (c) Exponential delay with mean 0.05sec. K -sync SGD in Figure 9 shows the test error for K -sync SGD and Figure 10 shows K -async SGD. For brevity, we present test error plots here and include the training loss plots in Appendix E. The training loss follows the same trend as test error.

As predicted from our theoretical analysis, increasing K improves the final error floor for all the SGD variants. But increasing K also increases the runtime per iteration. Hence, when the error is plotted against wallclock run-time, we begin to observe interesting trends in the error-runtime trade-offs – the highest K does not always achieve the best error-runtime trade-off. An intermediate value of K often achieves a better trade-off than fully-asynchronous or fully-synchronous SGD. For example, in Figure 9a, since there is little delay variability in the gradient computation time, $K = 8$ (fully synchronous SGD) is the best choice of K , but as the delay variability increases in Figure 9b and Figure 9c, $K = 4$ becomes the case that gives the fastest error-versus-wallclock time convergence. A similar trend can be observed in Figure 10 for K -async SGD.

Remark 5. Note that the problem is choosing K is similar in spirit to that of selecting the best mini-batch size in standard synchronous SGD. The main difference is that we consider the error-runtime trade-off instead of the error-iterations trade-off when making the choice.

So far we have considered partially synchronous/asynchronous SGD variants with a fixed value of K . In Section 6 we propose the AdaSync strategy that gradually adapts K during the training process in order to get the best error-runtime trade-off.




Figure 9: Test error (%) of **K-sync SGD** on CIFAR-10 with 8 worker nodes. We now demonstrate the error-runtime trade-off for the case with no artificial delay, and then also plot the trade-off as we add an exponential delay on each worker. As the mean of the additional delay increases, using an intermediate value of K achieves a better error-runtime trade-off. The trend of training loss is similar (see Figure 12 in Appendix E).




Figure 10: Test error (%) of **K-async SGD** on CIFAR-10 with 8 worker nodes. We demonstrate the error-runtime trade-off for the case with no artificial delay, and then also plot the trade-off as we add an exponential delay on each worker. As the mean of the additional delay increases, using an intermediate value of K achieves a better error-runtime trade-off. The trend of training loss is again similar (see Figure 13 in Appendix E).

6 AdaSync: Adaptive Synchronicity for Achieving the Best Error- Runtime Trade-Off

As we observed both theoretically and empirically above, asynchronous SGD converges faster (with respect to wallclock time) but has a higher error floor. On the other hand, synchronous SGD converges slower (again with respect to wallclock time) but achieves a lower error floor. In this section, our goal is to try to achieve the best of both worlds, *i.e.*, attain the most desirable error-runtime trade-off by gradually varying the level of synchronicity (parameter K) for the different SGD variants.

Let us partition the training time into intervals of time t each, such that, after every slot of time t we vary K . The number of iterations performed within time t is assumed to be approximately $N(t) \approx t/\mathbb{E}[T]$

where $\mathbb{E}[T]$ is the expected runtime per-iteration for the chosen SGD variant¹. From Theorem 4, we can write a (heuristic) upper bound on the average of $\mathbb{E}[\|\nabla F(\mathbf{w})\|_2^2]$ within each time interval t as follows:

$$\text{Upper Bound as a function of } K = u(K) = \frac{2(F(\mathbf{w}_{start}) - F^*)\mathbb{E}[T]}{t\eta\gamma'} + \frac{L\eta\sigma^2}{Km\gamma'}, \quad (22)$$

where \mathbf{w}_{start} denotes the value of the model \mathbf{w} at the beginning of that time interval². Our goal is to minimize $u(K)$ with respect to K for each time interval.

Observe that,

$$\frac{du(K)}{dK} = \frac{2(F(\mathbf{w}_{start}))}{t\eta\gamma'} \frac{d\mathbb{E}[T]}{dK} - \frac{L\eta\sigma^2}{K^2m\gamma'}. \quad (23)$$

Setting $\frac{du(K)}{dK}$ to 0 therefore provides a rough heuristic on how to choose parameter K for each time interval, as long as, $\frac{d^2u(K)}{dK^2}$ is positive. We derive the rule for adaptively varying K for each of the SGD variants in Appendix D. Here, we include the method for one variant K -async to demonstrate the key idea.

For general distributions, the runtime of K -async is upper bounded by that of K -sync (see Lemma 3). For exponential distributions, the two become equal and an algorithm similar to K -sync works. Here, we examine the interesting case of shifted-exponential distribution. We approximate $\mathbb{E}[T] \approx \frac{K\Delta}{P} + \frac{K \log P}{P\mu}$ (see Lemma 4). This leads to

$$K^2 = \frac{L\eta\sigma^2 t\eta P\mu}{2m(F(\mathbf{w}_{start}))(\Delta\mu + \log P)}.$$

To actually solve for this equation, we would need the values of the Lipschitz constant, variance of the gradient etc. which are not always available. We sidestep this issue by proposing a heuristic here that relies on the ratio of the parameter K at different instants.

Observe that, the larger is $F(\mathbf{w}_{start})$, the smaller is the value of K required to minimize $u(K)$. We assume that $F(\mathbf{w}_{start})$ is maximum at the beginning of training, *i.e.*, when $\mathbf{w}_{start} = \mathbf{w}_0$. Hence we start with the smallest initial K , *e.g.*, $K_0 = 1$. Thus, we could start with a small K_0 and after each time interval t , we can update K by solving for

$$K^2 = K_0^2 \frac{F(\mathbf{w}_0)}{F(\mathbf{w}_{start})}.$$

We can also verify that the second derivative is positive, *i.e.*,

$$\frac{d^2u(K)}{dK^2} = \frac{2L\eta\sigma^2}{K^3m} > 0.$$

The detailed algorithm for AdaSync for K -async SGD is described in Algorithm 1.

Algorithm 1 AdaSync for K -async SGD

- 1: Start with $K = K_0$ (typically $K_0 = 1$).
 - 2: **While** Wallclock Time \leq Total Time Budget **do**:
 - 3: Perform an iteration of K -async SGD
 - 4: **If** (Wallclock Time % $t = 0$) and ($K < P$) **do**:
 - 5: Update K as follows: $K = K_0 \sqrt{\frac{F(\mathbf{w}_0)}{F(\mathbf{w}_{start})}}$
-

For K -sync SGD under exponential assumption, the update rule for K is derived by solving for K in the quadratic equation: $\frac{K^2}{P-K} = \frac{K_0^2}{P-K_0} \frac{F(\mathbf{w}_0)}{F(\mathbf{w}_{start})}$, as discussed in Appendix D. For the two other variants of distributed SGD, the adaptive update rule for K is $K = K_0 \sqrt{\frac{F(\mathbf{w}_0)}{F(\mathbf{w}_{start})}}$, under certain assumptions on the runtime distribution, as also discussed in Appendix D.

¹Note that, for exponential inter-arrival times, this approximation holds in the limit of large t .

²Note that γ' can be set as 1 for synchronous variants. Though, this does not matter much here as the minimizing K does not depend on it.

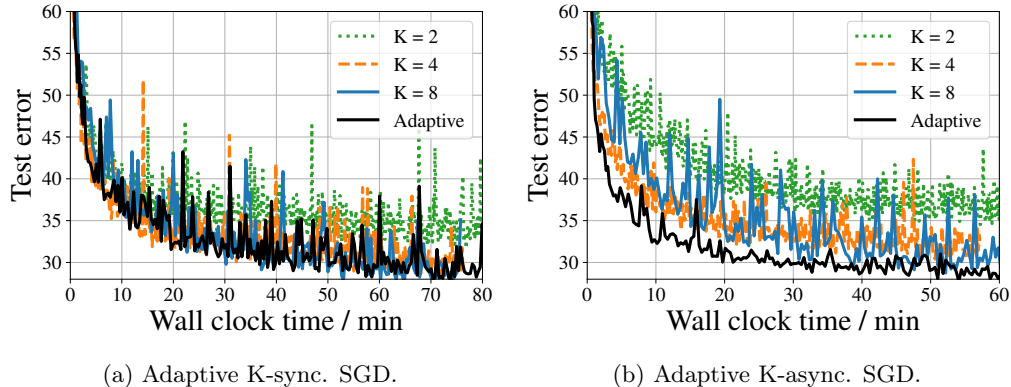


Figure 11: Test error of **AdaSync SGD** on CIFAR-10 with 8 worker nodes. We add an exponential delay with mean 0.02s on each worker. The value of K is changed after every 60 seconds.

6.1 AdaSync Experimental Results

In this subsection, we evaluate the effectiveness of AdaSync for both K-sync SGD and K-async SGD algorithms. An exponential delay with mean 0.02s is added to each worker node independently. We fix K for every $t = 60$ seconds (about 10 epochs). The initial values of K are fine-tuned and set to 2 and 4 for K-sync SGD and K-async SGD, respectively. As shown in Figure 11, the adaptive strategy achieves the fastest convergence in terms of error-versus-time. The adaptive K-async algorithm can even achieve a better error-runtime trade-off than the $K = 8$ case (i.e., fully synchronous case).

7 Concluding Remarks

This work introduces a novel analysis of error-runtime trade-off of distributed SGD, accounting for both error reduction per iteration as well as the wallclock runtime in a delay-prone computing environment. Furthermore, we also give a new analysis of asynchronous SGD by relaxing some commonly made assumptions in existing literature. Lastly, we also propose a novel strategy called AdaSync that adaptively increases synchronicity during distributed machine learning to achieve the best error-runtime trade-off. Our results provide valuable insights into distributed machine learning that could inform choice of workers and preferred method of parallelization for a particular distributed SGD algorithm in a chosen distributed computing environment.

As future work, we plan to explore methods of gradually increasing synchrony in other distributed optimization frameworks, e.g., federated learning, decentralized SGD, elastic averaging etc., that is closely related to [60, 61, 65–72]. Our proposed techniques can also inform hyperparameter tuning. Given some knowledge of the computing environment, our technique could allow one to simulate the expected error-runtime trade-off in advance, and possibly choose training parameters such as parameter K , mini-batch size m etc. It is also an interesting future direction to extend our current theoretical analysis for non iid scenarios, i.e., when the runtime or the dataset of different workers are not independent and identically distributed.

Acknowledgements

The authors thank Soumyadip Ghosh, Parijat Dube, Pulkit Grover, Souptik Sen, Li Zhang, Wei Zhang, and Priya Nagpurkar for helpful discussions. This work was supported in part by an IBM Faculty Award, and NSF CRII Award (CCF-1717314), and the Qualcomm Innovation fellowship.

References

- [1] Jeffrey Dean et al. Large scale distributed deep networks. In *Advances in Neural Information Processing Systems*, pages 1223–1231, 2012.
- [2] Jeffrey Dean and Luiz Andre Barroso. The tail at scale. *Communications of the ACM*, 56(2):74–80, 2013.
- [3] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. *Technical report, University of Toronto*, 2009.
- [4] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [5] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.
- [6] Ofer Dekel, Ran Gilad-Bachrach, Ohad Shamir, and Lin Xiao. Optimal distributed online prediction using mini-batches. *Journal of Machine Learning Research*, 13(1):165–202, 2012.
- [7] Mu Li, Tong Zhang, Yuqiang Chen, and Alexander J. Smola. Efficient mini-batch training for stochastic optimization. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 661–670, 2014.
- [8] Leon Bottou, Frank E. Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *arXiv:1606.04838*, 2016.
- [9] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 2, July 2011.
- [10] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations (ICLR)*, 2015.
- [11] Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in Neural Information Processing Systems*, pages 315–323, 2013.
- [12] Nicolas L Roux, Mark Schmidt, and Francis R Bach. A stochastic gradient method with an exponential convergence rate for finite training sets. In *Advances in Neural Information Processing Systems*, pages 2663–2671, 2012.
- [13] Lam M Nguyen, Jie Liu, Katya Scheinberg, and Martin Takac. Sarah: A novel method for machine learning problems using stochastic recursive gradient. In *International Conference on Machine Learning*, pages 2613–2621, 2017.
- [14] Sebastian Ruder. An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*, 2016.
- [15] Aaron Harlap, Henggang Cui, Wei Dai, Jinliang Wei, Gregory R. Ganger, Phillip B. Gibbons, Garth A. Gibson, and Eric P. Xing. Addressing the straggler problem for iterative convergent parallel ml. In *ACM Symposium on Cloud Computing (SoCC)*, pages 98–111, 2016.
- [16] Gauri Joshi, Yanpei Liu, and Emina Soljanin. On the delay-storage trade-off in content download from coded distributed storage systems. *IEEE Journal on Selected Areas in Communications*, 32(5):989–997, 2014.
- [17] Da Wang, Gauri Joshi, and Gregory Wornell. Using straggler replication to reduce latency in large-scale parallel computing. *ACM SIGMETRICS Performance Evaluation Review*, 43(3):7–11, 2015.

- [18] Gauri Joshi, Emina Soljanin, and Gregory Wornell. Queues with redundancy: Latency-cost analysis. *ACM SIGMETRICS Performance Evaluation Review*, 43(2):54–56, 2015.
- [19] Gauri Joshi, Emina Soljanin, and Gregory Wornell. Efficient redundancy techniques for latency reduction in cloud systems. *ACM Transactions on Modeling and Performance Evaluation of Computing Systems*, 2(12), may 2017.
- [20] Kangwook Lee, Maximilian Lam, Ramtin Pedarsani, Dimitris Papailiopoulos, and Kannan Ramchandran. Speeding up distributed machine learning using codes. *IEEE Transactions on Information Theory*, 64(3):1514–1529, 2018.
- [21] Rashish Tandon, Qi Lei, Alexandros G. Dimakis, and Nikos Karampatziakis. Gradient coding: Avoiding stragglers in distributed learning. In *International Conference on Machine Learning*, pages 3368–3376, 2017.
- [22] Sanghamitra Dutta, Viveck Cadambe, and Pulkit Grover. Short-dot: Computing large linear transforms distributedly using coded short dot products. In *Advances In Neural Information Processing Systems*, pages 2100–2108, 2016.
- [23] Wael Halbawi, Navid Azizan-Ruhi, Fariborz Salehi, and Babak Hassibi. Improving distributed gradient descent using reed-solomon codes. *arXiv preprint arXiv:1706.05436*, 2017.
- [24] Yaoqing Yang, Pulkit Grover, and Soumya Kar. Coded distributed computing for inverse problems. In *Advances in Neural Information Processing Systems*, pages 709–719, 2017.
- [25] Yaoqing Yang, Pulkit Grover, and Soumya Kar. Fault-tolerant distributed logistic regression using unreliable components. In *Communication, Control, and Computing (Allerton)*, pages 940–947. IEEE, 2016.
- [26] Can Karakus, Yifan Sun, and Suhas Diggavi. Encoded distributed optimization. In *IEEE International Symposium on Information Theory (ISIT)*, pages 2890–2894, 2017.
- [27] Can Karakus, Yifan Sun, Suhas Diggavi, and Wotao Yin. Straggler mitigation in distributed optimization through data encoding. In *Advances in Neural Information Processing Systems*, pages 5440–5448, 2017.
- [28] Zachary Charles, Dimitris Papailiopoulos, and Jordan Ellenberg. Approximate gradient coding via sparse random graphs. *arXiv preprint arXiv:1711.06771*, 2017.
- [29] Songze Li, Sucha Supittayapornpong, Mohammad Ali Maddah-Ali, and Salman Avestimehr. Coded terasort. In *IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 389–398, 2017.
- [30] Mohammad Fahim, Haewon Jeong, Farzin Haddadpour, Sanghamitra Dutta, Viveck Cadambe, and Pulkit Grover. On the optimal recovery threshold of coded matrix multiplication. In *Communication, Control, and Computing (Allerton)*, pages 1264–1270. IEEE, 2017.
- [31] Min Ye and Emmanuel Abbe. Communication-computation efficient gradient coding. *arXiv preprint arXiv:1802.03475*, 2018.
- [32] Songze Li, Mohammad Ali Maddah-Ali, Qian Yu, and A Salman Avestimehr. A fundamental tradeoff between computation and communication in distributed computing. *IEEE Transactions on Information Theory*, 64(1):109–128, 2018.
- [33] Viveck Cadambe and Pulkit Grover. Codes for Distributed Computing: A Tutorial. *IEEE Information Theory Society Newsletter*, 67(4):3–15, December 2017.

- [34] Sanghamitra Dutta, Ziqian Bai, Haewon Jeong, Tze Meng Low, and Pulkit Grover. A Unified Coded Deep Neural Network Training Strategy based on Generalized PolyDot codes. In *IEEE International Symposium on Information Theory (ISIT)*, 2018.
- [35] Ankur Mallick, Malhar Chaudhari, and Gauri Joshi. Rateless codes for near-perfect load balancing in distributed matrix-vector multiplication. *arXiv preprint arXiv:1804.10331*, 2018.
- [36] Sanghamitra Dutta, Viveck Cadambe, and Pulkit Grover. Coded convolution for parallel and distributed computing within a deadline. In *2017 IEEE International Symposium on Information Theory (ISIT)*, pages 2403–2407. IEEE, 2017.
- [37] Utsav Sheth, Sanghamitra Dutta, Malhar Chaudhari, Haewon Jeong, Yaoqing Yang, Jukka Kohonen, Teemu Roos, and Pulkit Grover. An application of storage-optimal matdot codes for coded matrix multiplication: Fast k-nearest neighbors estimation. In *2018 IEEE International Conference on Big Data (Big Data)*, pages 1113–1120. IEEE, 2018.
- [38] Emre Ozfatura, Deniz Gündüz, and Sennur Ulukus. Speeding up distributed gradient descent by utilizing non-persistent stragglers. In *2019 IEEE International Symposium on Information Theory (ISIT)*, pages 2729–2733. IEEE, 2019.
- [39] Haider Al-Lawati, Nuwan Ferdinand, and Stark C Draper. Anytime minibatch with stale gradients. In *2019 53rd Annual Conference on Information Sciences and Systems (CISS)*, pages 1–5. IEEE, 2019.
- [40] Raj Kumar Maity, Ankit Singh Rawa, and Arya Mazumdar. Robust gradient descent via moment encoding and ldpc codes. In *2019 IEEE International Symposium on Information Theory (ISIT)*, pages 2734–2738. IEEE, 2019.
- [41] Amirhossein Reiszadeh, Hossein Taheri, Aryan Mokhtari, Hamed Hassani, and Ramtin Pedarsani. Robust and communication-efficient collaborative learning. In *Advances in Neural Information Processing Systems*, pages 8386–8397, 2019.
- [42] Mohammad Mohammadi Amiri and Deniz Gündüz. Computation scheduling for distributed machine learning with straggling workers. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8177–8181. IEEE, 2019.
- [43] Suyog Gupta, Wei Zhang, and Fei Wang. Model accuracy and runtime tradeoff in distributed deep learning: A systematic study. In *International Conference on Data Mining*, pages 171–180, 2016.
- [44] James Cipar, Qirong Ho, Jin Kyu Kim, Seunghak Lee, Gregory R. Ganger, Garth Gibson, Kimberly Keeton, and Eric Xing. Solving the straggler problem with bounded staleness. In *Workshop on Hot Topics in Operating Systems*, 2013.
- [45] Henggang Cui, James Cipar, Qirong Ho, Jin Kyu Kim, Seunghak Lee, Abhimanu Kumar, Jinliang Wei, Wei Dai, Gregory R Ganger, Phillip B Gibbons, et al. Exploiting bounded staleness to speed up big data analytics. In *USENIX Annual Technical Conference (ATC)*, pages 37–48, 2014.
- [46] Qirong Ho, James Cipar, Henggang Cui, Seunghak Lee, Jin Kyu Kim, Phillip B Gibbons, Garth A Gibson, Greg Ganger, and Eric P Xing. More effective distributed ml via a stale synchronous parallel parameter server. In *Advances in Neural Information Processing Systems*, pages 1223–1231, 2013.
- [47] John Tsitsiklis, Dimitri Bertsekas, and Michael Athans. Distributed asynchronous deterministic and stochastic gradient optimization algorithms. *IEEE Transactions on Automatic Control*, 31(9):803–812, 1986.
- [48] Xiangru Lian, Yijun Huang, Yuncheng Li, and Ji Liu. Asynchronous parallel stochastic gradient for nonconvex optimization. In *Advances in Neural Information Processing Systems*, pages 2737–2745, 2015.

- [49] Ioannis Mitliagkas, Ce Zhang, Stefan Hadjis, and Christopher Re. Asynchrony begets momentum, with an application to deep learning. In *Allerton Conference on Communication, Control, and Computing*, pages 997–1004. IEEE, 2016.
- [50] Benjamin Recht, Christopher Re, Stephen Wright, and Feng Niu. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *Advances in Neural Information Processing Systems*, pages 693–701, 2011.
- [51] Alekh Agarwal and John C Duchi. Distributed delayed stochastic optimization. In *Advances in Neural Information Processing Systems*, pages 873–881, 2011.
- [52] Horia Mania, Xinghao Pan, Dimitris Papailiopoulos, Benjamin Recht, Kannan Ramchandran, and Michael I Jordan. Perturbed iterate analysis for asynchronous stochastic optimization. *SIAM Journal on Optimization*, 27(4):2202–2229, 2017.
- [53] Sorathan Chaturapruek, John C Duchi, and Christopher Re. Asynchronous stochastic convex optimization: the noise is in the noise and sgd don’t care. In *Advances in Neural Information Processing Systems*, pages 1531–1539, 2015.
- [54] Wei Zhang, Suyog Gupta, Xiangru Lian, and Ji Liu. Staleness-aware async-sgd for distributed deep learning. In *International Joint Conference on Artificial Intelligence*, pages 2350–2356. AAAI Press, 2016.
- [55] Zhimin Peng, Yangyang Xu, Ming Yan, and Wotao Yin. Arock: an algorithmic framework for asynchronous parallel coordinate updates. *SIAM Journal on Scientific Computing*, 38(5):A2851–A2879, 2016.
- [56] Robert Hannah and Wotao Yin. More iterations per second, same quality—why asynchronous algorithms may drastically outperform traditional ones. *arXiv preprint arXiv:1708.05136*, 2017.
- [57] Robert Hannah and Wotao Yin. On unbounded delays in asynchronous parallel fixed-point algorithms. *Journal of Scientific Computing*, pages 1–28, 2016.
- [58] Tao Sun, Robert Hannah, and Wotao Yin. Asynchronous coordinate descent under more realistic assumptions. In *Advances in Neural Information Processing Systems*, pages 6183–6191, 2017.
- [59] Remi Leblond, Fabian Pedregosa, and Simon Lacoste-Julien. Asaga: Asynchronous parallel saga. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2017.
- [60] Jianyu Wang and Gauri Joshi. Adaptive communication strategies to achieve the best error-runtime trade-off in local-update SGD. *arXiv preprint arXiv:1810.08313*, 2018.
- [61] Jianyu Wang and Gauri Joshi. Cooperative SGD: A unified framework for the design and analysis of communication-efficient SGD algorithms. *arXiv preprint arXiv:1808.07576*, 2018.
- [62] Jianmin Chen, Rajat Monga, Samy Bengio, and Rafal Jozefowicz. Revisiting distributed synchronous SGD. *CoRR*, abs/1604.00981, 2016.
- [63] Ross Sheldon. *A first course in probability*. Pearson Education India, 2002.
- [64] Robert G. Gallager. *Stochastic Processes: Theory for Applications*. Cambridge University Press, 1st edition, 2013.
- [65] Jian Zhang, Christopher De Sa, Ioannis Mitliagkas, and Christopher Re. Parallel sgd: When does averaging help? *arXiv preprint arXiv:1606.07365*, 2016.
- [66] Dong Yin, Ashwin Pananjady, Max Lam, Dimitris Papailiopoulos, Kannan Ramchandran, and Peter Bartlett. Gradient diversity empowers distributed learning. *arXiv preprint arXiv:1706.05699*, 2017.

- [67] Fan Zhou and Guojing Cong. On the convergence properties of a k -step averaging stochastic gradient descent algorithm for nonconvex optimization. *arXiv preprint arXiv:1708.01012*, 2017.
- [68] Sixin Zhang, Anna E Choromanska, and Yann LeCun. Deep learning with elastic averaging sgd. In *Advances in Neural Information Processing Systems*, pages 685–693, 2015.
- [69] H Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, et al. Communication-efficient learning of deep networks from decentralized data. *arXiv preprint arXiv:1602.05629*, 2016.
- [70] Jakub Konečný, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*, 2016.
- [71] Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. Federated learning: Challenges, methods, and future directions. *arXiv preprint arXiv:1908.07873*, 2019.
- [72] Jianyu Wang, Anit Kumar Sahu, Zhouyi Yang, Gauri Joshi, and Soumya Kar. Matcha: Speeding up decentralized sgd via matching decomposition sampling. *arXiv preprint arXiv:1905.09435*, 2019.
- [73] David M Kreps. *A course in microeconomic theory*, volume 41. JSTOR, 1990.

A Strong Convexity Discussion

Definition 4 (Strong-Convexity). *A function $h(\mathbf{u})$ is defined to be c -strongly convex, if the following holds for all \mathbf{u}_1 and \mathbf{u}_2 in the domain:*

$$h(\mathbf{u}_2) \geq h(\mathbf{u}_1) + [\nabla h(\mathbf{u}_1)]^T (\mathbf{u}_2 - \mathbf{u}_1) + \frac{c}{2} \|\mathbf{u}_2 - \mathbf{u}_1\|_2^2$$

For strongly convex functions, the following result holds for all \mathbf{u} in the domain of $h(\cdot)$.

$$2c(h(\mathbf{u}) - h^*) \leq \|\nabla h(\mathbf{u})\|_2^2 \tag{24}$$

The proof is derived in [8]. For completeness, we give the sketch here.

Proof. Given a particular \mathbf{u} , let us define the quadratic function as follows:

$$q(\mathbf{u}') = h(\mathbf{u}) + \nabla h(\mathbf{u})^T (\mathbf{u}' - \mathbf{u}) + \frac{c}{2} \|\mathbf{u}' - \mathbf{u}\|_2^2$$

Now, $q(\mathbf{u}')$ is minimized at $\mathbf{u}' = \mathbf{u} - \frac{1}{c} \nabla h(\mathbf{u})$ and the value is $h(\mathbf{u}) - \frac{1}{2c} \|\nabla h(\mathbf{u})\|_2^2$. Thus, from the definition of strong convexity we now have,

$$\begin{aligned} h^* &\geq h(\mathbf{u}) + \nabla h(\mathbf{u})^T (\mathbf{u}' - \mathbf{u}) + \frac{c}{2} \|\mathbf{u}' - \mathbf{u}\|_2^2 \\ &\geq h(\mathbf{u}) - \frac{1}{2c} \|\nabla h(\mathbf{u})\|_2^2 \quad [\text{minimum value of } q(\mathbf{u}')]. \end{aligned}$$

□

B Runtime Analysis Proofs

Here we provide all the remaining proofs and supplementary information for the results in Section 3.3.

B.1 Runtime of K-sync SGD

K -th statistic of exponential distributions: Here we give a sketch of why the K -th order statistic of P exponentials scales as $\log(P/P - K)$. A detailed derivation can be obtained in [63]. Consider P i.i.d. exponential distributions with parameter μ . The minimum $X_{1:P}$ of P independent exponential random variables with parameter μ is exponential with parameter $P\mu$. Conditional on $X_{1:P}$, the second smallest value $X_{2:P}$ is distributed like the sum of $X_{1:P}$ and an independent exponential random variable with parameter $(P - 1)\mu$. And so on, until the K -th smallest value $X_{K:P}$ which is distributed like the sum of $X_{(K-1):P}$ and an independent exponential random variable with parameter $(P - K + 1)\mu$. Thus,

$$X_{K:P} = Y_P + Y_{P-1} + \cdots + Y_{P-K+1}$$

where the random variables Y_i s are independent and exponential with parameter $i\mu$. Thus,

$$\mathbb{E}[X_{K:P}] = \sum_{i=P-K+1}^P \frac{1}{i\mu} = \frac{H_P - H_{P-K}}{\mu} \approx \frac{\log \frac{P}{P-K}}{\mu}.$$

Here H_P and H_{P-K} denote the P -th and $(P - K)$ -th harmonic numbers respectively.

For the case where $K = P$, the expectation is given by,

$$\mathbb{E}[X_{P:P}] = \frac{1}{\mu} \sum_{i=1}^P \frac{1}{i} = \frac{1}{\mu} H_P \approx \frac{1}{\mu} \log P.$$

B.2 Runtime of K-batch-sync SGD

Proof of Lemma 2. For new-longer-than-used distributions observe that the following holds:

$$\Pr(X_i > u + t | X_i > t) \leq \Pr(X_i > u). \quad (25)$$

Thus the random variable corresponding to the remaining time, *i.e.*, $X_i - t | X_i > t$ is thus stochastically dominated by X_i .

Assume that we want to compute the expected computation time of one iteration of K -batch-sync SGD starting at time instant t_0 . The time taken to complete the first mini-batch is the minimum of P iid random variables, denoted as $\mathbb{E}[X_{1:P}]$. Suppose the first mini-batch finishes at time t_1 , and the worker which finished first has index i' . Now, the time taken to complete the next mini-batch is the minimum of the remaining runtimes of all the P workers of which the i' -th worker just started afresh at time t_1 while the $P - 1$ other workers had been computing since time t_0 . Let Y_1, Y_2, \dots, Y_P be the random variables denoting the remaining computation time of the P workers after time t_1 . Thus,

$$Y_i = \begin{cases} X_i - (t_1 - t_0) | X_i > (t_1 - t_0) & \text{for all } i \neq i' \\ X_i & \text{otherwise.} \end{cases} \quad (26)$$

Now each of the Y_i s are independent and are stochastically dominated by X_i s.

$$\Pr(Y_i > u) \leq \Pr(X_i > u) \quad \forall i = 1, 2, \dots, P. \quad (27)$$

The expectation of the minimum of $\{Y_1, Y_2, \dots, Y_P\}$ is the expected runtime of computing the second mini-batch. Let us denote $h_1(x_1, x_2, \dots, x_P)$ as the minimum of P numbers (x_1, x_2, \dots, x_P) . And let us denote $g_{1,s}(x)$ as the minimum of P numbers where $P - 1$ of them are given as $\mathbf{s}_{1 \times (P-1)}$ and x is the P -th number. Thus

$$g_{1,s}(x) = h_1(x, s(1), s(2), \dots, s(P-1)).$$

First observe that $g_{1,\mathbf{s}}(x)$ is an increasing function of x since given all the other $P - 1$ values, the minimum will either stay the same or increase with x . Now we use the property that if Y_i is stochastically dominated by X_i , then for any increasing function $g(\cdot)$, we have

$$\mathbb{E}_{Y_1} [g(Y_1)] \leq \mathbb{E}_{X_1} [g(X_1)].$$

This result is derived in [73].

This implies that for a given \mathbf{s} ,

$$\mathbb{E}_{Y_1} [g_{K,\mathbf{s}}(Y_1)] \leq \mathbb{E}_{X_1} [g_{K,\mathbf{s}}(X_1)].$$

This leads to,

$$\begin{aligned} \mathbb{E}_{Y_1|Y_2=s(1),Y_3=s(2)\dots Y_P=s(P-1)} [h_1(Y_1, Y_2, \dots Y_P)] \\ \leq \mathbb{E}_{X_1|Y_2=s(1),Y_3=s(2)\dots Y_P=s(P-1)} [h_1(X_1, Y_2, \dots Y_P)]. \end{aligned} \quad (28)$$

From this,

$$\begin{aligned} \mathbb{E} [h_1(Y_1, Y_2, \dots Y_P)] &= \mathbb{E}_{Y_2, \dots, Y_P} [\mathbb{E}_{Y_1|Y_2, Y_3, \dots Y_P} [h_1(Y_1, Y_2, \dots Y_P)]] \\ &\leq \mathbb{E}_{Y_2, \dots, Y_P} [\mathbb{E}_{X_1|Y_2, Y_3, \dots Y_P} [h_1(X_1, Y_2, \dots Y_P)]] \\ &= \mathbb{E} [h_1(X_1, Y_2, \dots Y_P)]. \end{aligned} \quad (29)$$

This step proceeds inductively. Thus, similarly

$$\begin{aligned} \mathbb{E} [h_1(X_1, Y_2, \dots Y_P)] &= \mathbb{E}_{X_1, Y_3, \dots, Y_P} [\mathbb{E}_{Y_2|X_1, Y_3, \dots Y_P} [h_1(X_1, Y_2, \dots Y_P)]] \\ &\leq \mathbb{E}_{X_1, Y_3, \dots, Y_P} [\mathbb{E}_{X_2|X_1, Y_3, \dots Y_P} [h_1(X_1, X_2, Y_3, \dots Y_P)]] \\ &= \mathbb{E} [h_1(X_1, X_2, Y_3, \dots Y_P)]. \end{aligned} \quad (30)$$

Thus, finally combining, we have,

$$\begin{aligned} \mathbb{E} [h_1(Y_1, Y_2, \dots Y_P)] &\leq \mathbb{E} [h_1(X_1, Y_2, \dots Y_P)] \\ &\leq \mathbb{E} [h_1(X_1, X_2, Y_3, \dots Y_P)] \\ &\leq \dots \\ &\leq \mathbb{E} [h_1(X_1, X_2, X_3, \dots X_P)]. \end{aligned} \quad (31)$$

Therefore, the expected runtime of the second mini-batch is upper bounded by $\mathbb{E}[X_{1:P}]$. A similar argument can be made for all the K mini-batches, leading to an upper bound of $K\mathbb{E}[X_{1:P}]$ for all the K mini-batches. \square

In general, the exact analytical expression for the expected runtime per iteration of K -batch-sync SGD is not tractable but for the special case of exponentials it follows the distribution $Erlang(K, P\mu)$. This is obtained from the memoryless property of exponentials.

All the workers start their computation together. The expected time taken by the first mini-batch to be completed is the minimum of P i.i.d. exponential random variables $X_1, X_2, \dots, X_P \sim Exp(\mu)$ is another exponential random variable distributed as $Exp(P\mu)$. At the time when the first mini-batch is complete, from the memoryless property of exponentials, it may be viewed as P i.i.d. exponential random variables $X_1, X_2, \dots, X_P \sim Exp(\mu)$ starting afresh again. Thus, the time to complete each mini-batch is distributed as $Exp(P\mu)$, and an iteration being the sum of the time to complete K such mini-batches, has the distribution $Erlang(K, P\mu)$.

B.3 Runtime of K-async SGD

Proof of Lemma 3. For new-longer-than-used distributions observe that the following holds:

$$\Pr(X_i > u + t | X_i > t) \leq \Pr(X_i > u). \quad (32)$$

Thus the random variable $X_i - t | X_i > t$ is thus stochastically dominated by X_i . Now let us assume we want to compute the expected computation time of one iteration of K -async starting at time instant t_0 . Let us also assume that the workers last read their parameter values at time instants t_1, t_2, \dots, t_P respectively where any K of these t_1, t_2, \dots, t_P are equal to t_0 as K out of P workers were updated at time t_0 and the remaining $(P - K)$ of these t_1, t_2, \dots, t_P are $< t_0$. Let Y_1, Y_2, \dots, Y_P be the random variables denoting the computation time of the P workers starting from time t_0 . Thus,

$$Y_i = X_i - (t_0 - t_i) | X_i > (t_0 - t_i) \quad \forall i = 1, 2, \dots, P. \quad (33)$$

Now each of the Y_i s are independent and are stochastically dominated by X_j s.

$$\Pr(Y_i > u) \leq \Pr(X_j > u) \quad \forall i, j = 1, 2, \dots, P. \quad (34)$$

The expectation of the K -th statistic of $\{Y_1, Y_2, \dots, Y_P\}$ is the expected runtime of the iteration. Let us denote $h_K(x_1, x_2, \dots, x_P)$ as the K -th statistic of P numbers (x_1, x_2, \dots, x_P) . And let us denote $g_{K,s}(x)$ as the K -th statistic of P numbers where $P - 1$ of them are given as $\mathbf{s}_{1 \times (P-1)}$ and x is the P -th number. Thus

$$g_{K,s}(x) = h_K(x, s(1), s(2), \dots, s(P-1))$$

First observe that $g_{K,s}(x)$ is an increasing function of x since given the other $P - 1$ values, the K -th order statistic will either stay the same or increase with x . Now we use the property that if Y_i is stochastically dominated by X_i , then for any increasing function $g(\cdot)$, we have

$$\mathbb{E}_{Y_1} [g(Y_1)] \leq \mathbb{E}_{X_1} [g(X_1)].$$

This result is derived in [73].

This implies that for a given \mathbf{s} ,

$$\mathbb{E}_{Y_1} [g_{K,s}(Y_1)] \leq \mathbb{E}_{X_1} [g_{K,s}(X_1)].$$

This leads to,

$$\begin{aligned} \mathbb{E}_{Y_1 | Y_2=s(1), Y_3=s(2), \dots, Y_P=s(P-1)} [h_K(Y_1, Y_2, \dots, Y_P)] \\ \leq \mathbb{E}_{X_1 | Y_2=s(1), Y_3=s(2), \dots, Y_P=s(P-1)} [h_K(X_1, Y_2, \dots, Y_P)]. \end{aligned} \quad (35)$$

From this,

$$\begin{aligned} \mathbb{E} [h_K(Y_1, Y_2, \dots, Y_P)] &= \mathbb{E}_{Y_2, \dots, Y_P} [\mathbb{E}_{Y_1 | Y_2, Y_3, \dots, Y_P} [h_K(Y_1, Y_2, \dots, Y_P)]] \\ &\leq \mathbb{E}_{Y_2, \dots, Y_P} [\mathbb{E}_{X_1 | Y_2, Y_3, \dots, Y_P} [h_K(X_1, Y_2, \dots, Y_P)]] \\ &= \mathbb{E} [h_K(X_1, Y_2, \dots, Y_P)]. \end{aligned} \quad (36)$$

This step proceeds inductively. Thus, similarly

$$\begin{aligned} \mathbb{E} [h_K(X_1, Y_2, \dots, Y_P)] &= \mathbb{E}_{X_1, Y_3, \dots, Y_P} [\mathbb{E}_{Y_2 | X_1, Y_3, \dots, Y_P} [h_K(X_1, Y_2, \dots, Y_P)]] \\ &\leq \mathbb{E}_{X_1, Y_3, \dots, Y_P} [\mathbb{E}_{X_2 | X_1, Y_3, \dots, Y_P} [h_K(X_1, X_2, Y_3, \dots, Y_P)]] \\ &= \mathbb{E} [h_K(X_1, X_2, Y_3, \dots, Y_P)]. \end{aligned} \quad (37)$$

Thus, finally combining, we have,

$$\mathbb{E} [h_K(Y_1, Y_2, \dots, Y_P)] \leq \mathbb{E} [h_K(X_1, Y_2, \dots, Y_P)]$$

$$\begin{aligned}
&\leq \mathbb{E} [h_K(X_1, X_2, Y_3 \dots Y_P)] \\
&\leq \dots \\
&\leq \mathbb{E} [h_K(X_1, X_2, X_3 \dots X_P)].
\end{aligned} \tag{38}$$

□

B.3.1 Discussion for special distributions

For exponential distributions, the inequality in Lemma 3 holds with equality. This follows from the memoryless property of exponentials. Let us consider the scenario of the proof of Lemma 3 where we similarly define $Y_i = X_i - (t_0 - t_i) | X_i > (t_0 - t_i)$. From the memoryless property of exponentials [63], if $X_i \sim \text{Exp}(\mu)$, then $Y_i \sim \text{Exp}(\mu)$. Thus, the expectation of the K -th statistic of Y_i s can be easily derived as all the Y_i s are now i.i.d. with distribution $\text{Exp}(\mu)$. Thus, the expected runtime per iteration is given by,

$$\mathbb{E} [T] = \mathbb{E} [Y_{K:P}] = \frac{1}{\mu} \sum_{i=P-K+1}^P \frac{1}{i} \approx \frac{1}{\mu} \log \frac{P}{P-K}.$$

Proof of Lemma 4. Here, we derive an alternate upper bound on the runtime for n consecutive iterations for shifted-exponential distributions where $P = nK$. Suppose, we are looking at the first n consecutive iterations starting from time t_0 . Assuming all the workers launched afresh for the first iteration, the expected time for the first iteration is given by $\mathbb{E} [X_{K:P}] = \Delta + \mathbb{E} [\tilde{X}_{K:P}]$ where $\tilde{X} \sim \text{Exp}(\mu)$.

For the second iteration, the runtime is the K -th order statistic of the remaining runtimes of the P workers, where K of them just launched afresh (hence shifted-exponential) and the remaining $P - K$ of them have been computing since t_0 . The distribution of the remaining runtime for these $P - K$ workers correspond to exponential distributions with no shift ($\tilde{X} \sim \text{Exp}(\mu)$) because they have been running for at least time Δ . The expectation of the K -th order statistic of all the P runtimes can thus be upper-bounded by the K -th order statistic of a subset of the P remaining runtimes, specifically, these $P - K$ workers having the exponential runtime distribution of \tilde{X} .

Similarly, the expected runtime of the third iteration is the K -th order statistic of the remaining runtimes of the P workers of which at least $P - 2K$ of them have been running since t_0 (for at least time Δ) and hence have the distribution of \tilde{X} . Thus, the expected runtime of the third iteration is upper-bounded by the K -th order statistic of these $P - 2K$ exponential random variables. Extending a similar argument for n iterations, the expected runtime for n consecutive iterations can be upper-bounded as

$$\mathbb{E} [T_1 + T_2 + \dots + T_n] \leq \Delta + \sum_{i=0}^{n-1} \mathbb{E} [\tilde{X}_{K:(P-iK)}]. \tag{39}$$

From the new-longer-than-used property, the expected runtime of any n consecutive iterations is upper-bounded by the expected runtime of the first n consecutive iterations when all the P workers are launched afresh. □

B.4 Runtime of K-batch-async SGD

Here we include a discussion on renewal processes for completeness, to provide a background for the proof of Lemma 5, which gives the expected runtime of K -batch-async SGD. The familiar reader can merely skim through this and refer to the proof provided in the main section of the paper in Section 3.3.

Definition 5 (Renewal Process). *A renewal process is an arrival process where the inter-arrival intervals are positive, independent and identically distributed random variables.*

Lemma 8 (Elementary Renewal Theorem). *[64, Chapter 5] Let $\{N(t), t > 0\}$ be a renewal counting process denoting the number of renewals in time t . Let $\mathbb{E} [Z]$ be the mean inter-arrival time. Then,*

$$\lim_{t \rightarrow \infty} \frac{\mathbb{E} [N(t)]}{t} = \frac{1}{\mathbb{E} [Z]}. \tag{40}$$

Observe that for asynchronous SGD or K -batch-async SGD, every gradient push by a worker to the PS can be thought of as an arrival process. The time between two consecutive pushes by a worker follows the distribution of X_i and is independent as computation time has been assumed to be independent across workers and mini-batches. Thus the inter-arrival intervals are positive, independent and identically distributed and hence, the gradient pushes are a renewal process.

C Error Analysis Proofs

In this section, we first discuss and understand the significance of the quantity p_0 that was introduced in the error-analysis in Theorem 3 in Appendix C.1. Next, for ease of understanding, we discuss and analyze the error convergence of asynchronous SGD first in Appendix C.2 because of its simplicity. This will be followed by the proof of the more general result, *i.e.*, the error convergence analysis of K -async SGD (Theorem 3) in Appendix C.3.

C.1 Discussion on p_0

Let us denote the conditional probability of $\tau(l, j) = j$ given all the past delays and parameters as $p_0^{(l, j)}$. Now $p_0 \leq p_0^{(l, j)} \forall l, j$. Clearly the value of $p_0^{(l, j)}$ will differ for different distributions and accordingly the value of p_0 will differ. Here we include a brief discussion on the possible values of p_0 for different distributions. These also hold for K -async and K -batch-async SGD.

Proof of Lemma 6. Let t_0 be the time when the j -th iteration occurs, and suppose that worker i' pushed its gradient in the j -th iteration. Now similar to the proof of Lemma 3, let us also assume that the workers last read their parameter values at time instants t_1, t_2, \dots, t_P respectively where $t_i = t_0$ and the remaining $(P - 1)$ of these t_i s are $< t_0$. Let Y_1, Y_2, \dots, Y_P be the random variables denoting the computation time of the P workers starting from time t_0 . Thus, $Y_i = X_i - (t_0 - t_i) | X_i > (t_0 - t_i)$. For exponentials, from the memoryless property, all these Y_i s become i.i.d. and thus from symmetry the probability of i' finishing before all the others is equal, *i.e.* $\frac{1}{P}$. Thus, $p_0^{(j)} = p_0 = \frac{1}{P}$. For new-longer-than-used distributions, as we have discussed before all the Y_i s with $i \neq i'$ will be stochastically dominated by $Y_{i'} = X_{i'}$. Thus, probability of i s with $i \neq i'$ finishing first is higher than i' . Thus, $p_0^{(j)} \leq \frac{1}{P}$ and so is p_0 . Similarly, for new-shorter-than-used distributions, $Y_{i'}$ is stochastically dominated by all the Y_i s and thus probability of i' finishing first is more. So, $p_0^{(j)} \geq \frac{1}{P}$ and so is p_0 . \square

We use Lemma 9 below to prove Theorem 3.

Lemma 9. *Suppose that $p_0^{(l, j)}$ is the conditional probability that $\tau(l, j) = j$ given all the past delays and all the previous \mathbf{w} , and $p_0 \leq p_0^{(l, j)}$ for all j . Then,*

$$\mathbb{E} [\|\nabla F(\mathbf{w}_{\tau(l, j)})\|_2^2] \geq p_0 \mathbb{E} [\|\nabla F(\mathbf{w}_j)\|_2^2]. \quad (41)$$

Proof. By the law of total expectation,

$$\begin{aligned} & \mathbb{E} [\|\nabla F(\mathbf{w}_{\tau(l, j)})\|_2^2] \\ &= p_0^{(l, j)} \mathbb{E} [\|\nabla F(\mathbf{w}_{\tau(l, j)})\|_2^2 | \tau(j) = j] + (1 - p_0^{(l, j)}) \mathbb{E} [\|\nabla F(\mathbf{w}_{\tau(l, j)})\|_2^2 | \tau(j) \neq j] \\ &\geq p_0 \mathbb{E} [\|\nabla F(\mathbf{w}_j)\|_2^2]. \end{aligned}$$

\square

C.2 Async-SGD with fixed learning rate

First we prove a simplified version of Theorem 3 for the case $K = 1$. While this is actually a corollary of the more general Theorem 3, we prove this first for ease of understanding and simplicity. The proof of the more general Theorem 3 is then provided in Appendix C.3.

The corollary is as follows:

Corollary 2. *Suppose that the objective function $F(\mathbf{w})$ is strongly convex with parameter c and the learning rate $\eta \leq \frac{1}{2L(\frac{MG}{m}+1)}$. Also assume that $\mathbb{E} [\|\nabla F(\mathbf{w}_j) - \nabla F(\mathbf{w}_{\tau(j)})\|_2^2] \leq \gamma \mathbb{E} [\|\nabla F(\mathbf{w}_j)\|_2^2]$ for some constant $\gamma \leq 1$. Then, the error after J iterations of Async SGD is given by,*

$$\mathbb{E} [F(\mathbf{w}_J)] - F^* \leq \frac{\eta L \sigma^2}{2c\gamma' m} + (1 - \eta c \gamma')^J \left(\mathbb{E} [F(\mathbf{w}_0)] - F^* - \frac{\eta L \sigma^2}{2c\gamma' m} \right),$$

where $\gamma' = 1 - \gamma + \frac{p_0}{2}$ and p_0 is a non-negative lower bound on the conditional probability that $\tau(j) = j$ given all the past delays and parameters.

To prove the result, we will use the following lemma.

Lemma 10. *Let us denote $\mathbf{v}_j = g(\mathbf{w}_{\tau(j)}, \xi_j)$, and assume that $\mathbb{E}_{\xi_j|\mathbf{w}} [g(\mathbf{w}, \xi_j)] = \nabla F(\mathbf{w})$. Then,*

$$\begin{aligned} \mathbb{E} [\|\nabla F(\mathbf{w}_j) - \mathbf{v}_j\|_2^2] &\leq \mathbb{E} [\|\mathbf{v}_j\|_2^2] - \mathbb{E} [\|\nabla F(\mathbf{w}_{\tau(j)})\|_2^2] \\ &\quad + \mathbb{E} [\|\nabla F(\mathbf{w}_j) - \nabla F(\mathbf{w}_{\tau(j)})\|_2^2]. \end{aligned}$$

Proof of Lemma 10. Observe that,

$$\begin{aligned} \mathbb{E} [\|\nabla F(\mathbf{w}_j) - \mathbf{v}_j\|_2^2] &= \mathbb{E} [\|\nabla F(\mathbf{w}_j) - \nabla F(\mathbf{w}_{\tau(j)}) + \nabla F(\mathbf{w}_{\tau(j)}) - \mathbf{v}_j\|_2^2] \\ &= \mathbb{E} [\|\nabla F(\mathbf{w}_j) - \nabla F(\mathbf{w}_{\tau(j)})\|_2^2] + \mathbb{E} [\|\mathbf{v}_j - \nabla F(\mathbf{w}_{\tau(j)})\|_2^2]. \end{aligned} \quad (42)$$

The last line holds since the cross term is 0 as derived below.

$$\begin{aligned} &\mathbb{E} [(\nabla F(\mathbf{w}_j) - \nabla F(\mathbf{w}_{\tau(j)}))^T (\mathbf{v}_j - \nabla F(\mathbf{w}_{\tau(j)}))] \\ &= \mathbb{E}_{\mathbf{w}_{\tau(j)}, \mathbf{w}_j} [(\nabla F(\mathbf{w}_j) - \nabla F(\mathbf{w}_{\tau(j)}))^T \mathbb{E}_{\xi_j|\mathbf{w}_{\tau(j)}, \mathbf{w}_j} [(\mathbf{v}_j - \nabla F(\mathbf{w}_{\tau(j)}))]] \\ &= \mathbb{E}_{\mathbf{w}_{\tau(j)}, \mathbf{w}_j} [(\nabla F(\mathbf{w}_j) - \nabla F(\mathbf{w}_{\tau(j)}))^T (\mathbb{E}_{\xi_j|\mathbf{w}_{\tau(j)}} [\mathbf{v}_j] - \nabla F(\mathbf{w}_{\tau(j)}))] \\ &= 0. \end{aligned}$$

Here again the last line follows from Assumption 2 in Section 2 which states that

$$\mathbb{E}_{\xi_j|\mathbf{w}_{\tau(j)}} [\mathbf{v}_j] = \nabla F(\mathbf{w}_{\tau(j)}).$$

Returning to (42), observe that the second term can be further decomposed as,

$$\begin{aligned} \mathbb{E} [\|\mathbf{v}_j - \nabla F(\mathbf{w}_{\tau(j)})\|_2^2] &= \mathbb{E}_{\mathbf{w}_{\tau(j)}} \left[\mathbb{E}_{\xi_j|\mathbf{w}_{\tau(j)}} [\|\mathbf{v}_j - \nabla F(\mathbf{w}_{\tau(j)})\|_2^2] \right] \\ &= \mathbb{E}_{\mathbf{w}_{\tau(j)}} \left[\mathbb{E}_{\xi_j|\mathbf{w}_{\tau(j)}} [\|\mathbf{v}_j\|_2^2] \right] - 2 \mathbb{E}_{\mathbf{w}_{\tau(j)}} \left[\mathbb{E}_{\xi_j|\mathbf{w}_{\tau(j)}} [\mathbf{v}_j^T \nabla F(\mathbf{w}_{\tau(j)})] \right] \\ &\quad + \mathbb{E}_{\mathbf{w}_{\tau(j)}} \left[\mathbb{E}_{\xi_j|\mathbf{w}_{\tau(j)}} [\|\nabla F(\mathbf{w}_{\tau(j)})\|_2^2] \right] \\ &= \mathbb{E} [\|\mathbf{v}_j\|_2^2] - 2 \mathbb{E} [\|\nabla F(\mathbf{w}_{\tau(j)})\|_2^2] + \mathbb{E} [\|\nabla F(\mathbf{w}_{\tau(j)})\|_2^2] \\ &= \mathbb{E} [\|\mathbf{v}_j\|_2^2] - \mathbb{E} [\|\nabla F(\mathbf{w}_{\tau(j)})\|_2^2]. \end{aligned}$$

□

We will also be proving a K -worker version of this lemma Appendix C.3 to prove Theorem 3. Now we proceed to provide the proof of Corollary 2.

Proof of Corollary 2.

$$\begin{aligned}
F(\mathbf{w}_{j+1}) &\leq F(\mathbf{w}_j) + (\mathbf{w}_{j+1} - \mathbf{w}_j)^T \nabla F(\mathbf{w}_j) + \frac{L}{2} \|\mathbf{w}_{j+1} - \mathbf{w}_j\|_2^2 \\
&= F(\mathbf{w}_j) + (-\eta \mathbf{v}_j)^T \nabla F(\mathbf{w}_j) + \frac{L\eta^2}{2} \|\mathbf{v}_j\|_2^2 \\
&= F(\mathbf{w}_j) - \frac{\eta}{2} \|\nabla F(\mathbf{w}_j)\|_2^2 - \frac{\eta}{2} \|\mathbf{v}_j\|_2^2 + \frac{\eta}{2} \|\nabla F(\mathbf{w}_j) - \mathbf{v}_j\|_2^2 + \frac{L\eta^2}{2} \|\mathbf{v}_j\|_2^2.
\end{aligned} \tag{43}$$

Here the last line follows from $2\mathbf{a}^T \mathbf{b} = \|\mathbf{a}\|_2^2 + \|\mathbf{b}\|_2^2 - \|\mathbf{a} - \mathbf{b}\|_2^2$. Taking expectation,

$$\begin{aligned}
\mathbb{E}[F(\mathbf{w}_{j+1})] &\leq \mathbb{E}[F(\mathbf{w}_j)] - \frac{\eta}{2} \mathbb{E}[\|\nabla F(\mathbf{w}_j)\|_2^2] - \frac{\eta}{2} \mathbb{E}[\|\mathbf{v}_j\|_2^2] + \frac{\eta}{2} \mathbb{E}[\|\nabla F(\mathbf{w}_j) - \mathbf{v}_j\|_2^2] + \frac{L\eta^2}{2} \mathbb{E}[\|\mathbf{v}_j\|_2^2] \\
&\stackrel{(a)}{\leq} \mathbb{E}[F(\mathbf{w}_j)] - \frac{\eta}{2} \mathbb{E}[\|\nabla F(\mathbf{w}_j)\|_2^2] - \frac{\eta}{2} \mathbb{E}[\|\mathbf{v}_j\|_2^2] + \frac{\eta}{2} \mathbb{E}[\|\mathbf{v}_j\|_2^2] - \frac{\eta}{2} \mathbb{E}[\|\nabla F(\mathbf{w}_{\tau(j)})\|_2^2] \\
&\quad + \frac{\eta}{2} \mathbb{E}[\|\nabla F(\mathbf{w}_j) - \nabla F(\mathbf{w}_{\tau(j)})\|_2^2] + \frac{L\eta^2}{2} \mathbb{E}[\|\mathbf{v}_j\|_2^2].
\end{aligned} \tag{44}$$

Here, (a) follows from Lemma 10 that we just derived. Now, again bounding from (44), we have

$$\begin{aligned}
\mathbb{E}[F(\mathbf{w}_{j+1})] &\stackrel{(b)}{\leq} \mathbb{E}[F(\mathbf{w}_j)] - \frac{\eta}{2} \mathbb{E}[\|\nabla F(\mathbf{w}_j)\|_2^2] - \frac{\eta}{2} \mathbb{E}[\|\nabla F(\mathbf{w}_{\tau(j)})\|_2^2] + \frac{\eta}{2} \gamma \mathbb{E}[\|\nabla F(\mathbf{w}_j)\|_2^2] \\
&\quad + \frac{L\eta^2}{2} \mathbb{E}[\|\mathbf{v}_j\|_2^2] \\
&\stackrel{(c)}{\leq} \mathbb{E}[F(\mathbf{w}_j)] - \frac{\eta}{2} (1 - \gamma) \mathbb{E}[\|\nabla F(\mathbf{w}_j)\|_2^2] + \frac{L\eta^2 \sigma^2}{2m} \\
&\quad - \frac{\eta}{2} \left(1 - L\eta \frac{M_G}{m} + 1\right) \mathbb{E}[\|\nabla F(\mathbf{w}_{\tau(j)})\|_2^2] \\
&\stackrel{(d)}{\leq} \mathbb{E}[F(\mathbf{w}_j)] - \frac{\eta}{2} (1 - \gamma) \mathbb{E}[\|\nabla F(\mathbf{w}_j)\|_2^2] + \frac{L\eta^2 \sigma^2}{2m} - \frac{\eta}{4} \mathbb{E}[\|\nabla F(\mathbf{w}_{\tau(j)})\|_2^2] \\
&\stackrel{(e)}{\leq} \mathbb{E}[F(\mathbf{w}_j)] - \frac{\eta}{2} (1 - \gamma) \mathbb{E}[\|\nabla F(\mathbf{w}_j)\|_2^2] + \frac{L\eta^2 \sigma^2}{2m} - \frac{\eta}{4} p_0 \mathbb{E}[\|\nabla F(\mathbf{w}_j)\|_2^2].
\end{aligned} \tag{45}$$

Here (b) follows from the statement of the theorem that

$$\mathbb{E}[\|\nabla F(\mathbf{w}_j) - \nabla F(\mathbf{w}_{\tau(j)})\|_2^2] \leq \gamma \mathbb{E}[\|\nabla F(\mathbf{w}_j)\|_2^2]$$

for some constant $\gamma \leq 1$. The next step (c) follows from Assumption 4 in Section 2 which lead to

$$\mathbb{E}[\|\mathbf{v}_j\|_2^2] \leq \frac{\sigma^2}{m} + \left(\frac{M_G}{m} + 1\right) \mathbb{E}[\|\nabla F(\mathbf{w}_{\tau(j)})\|_2^2].$$

Step (d) follows from choosing $\eta < \frac{1}{2L(\frac{M_G}{m} + 1)}$ and finally (e) follows from Lemma 9.

Now one might recall that the function $F(w)$ was defined to be strongly convex with parameter c .

Using the standard result of strong-convexity (16) in (45), we obtain the following result:

$$\mathbb{E}[F(\mathbf{w}_{j+1})] - F^* \leq \frac{\eta^2 L \sigma^2}{2m} + (1 - \eta c (1 - \gamma + \frac{p_0}{2})) (\mathbb{E}[F(\mathbf{w}_j)] - F^*).$$

Let us denote $\gamma' = (1 - \gamma + \frac{p_0}{2})$. Then, using the above recursion, we thus have,

$$\mathbb{E}[F(\mathbf{w}_J)] - F^* \leq \frac{\eta L \sigma^2}{2c\gamma' m} + (1 - \eta\gamma' c)^J (\mathbb{E}[F(\mathbf{w}_0)] - F^* - \frac{\eta L \sigma^2}{2c\gamma' m}).$$

□

C.3 K-async SGD under fixed learning rate

In this subsection, we provide a proof of Theorem 3.

Before we proceed to the proof of this theorem, we first extend our Assumption 4 from the variance of a single stochastic gradient to sum of stochastic gradients in the following Lemma.

Lemma 11. *If the variance of the stochastic updates is bounded as*

$$\mathbb{E}_{\xi_j | \mathbf{w}_{\tau(l,j)}} [\|g(\mathbf{w}_{\tau(l,j)}, \xi_{l,j}) - \nabla F(\mathbf{w}_{\tau(l,j)})\|_2^2] \leq \frac{\sigma^2}{m} + \frac{M_G}{m} \|\nabla F(\mathbf{w}_{\tau(l,j)})\|_2^2 \quad \forall \tau(l,j) \leq j,$$

then for K-async, the variance of the sum of stochastic updates given all the parameter values $\mathbf{w}_{\tau(l,j)}$ is also bounded as follows:

$$\begin{aligned} & \mathbb{E}_{\xi_{1,j}, \dots, \xi_{K,j} | \mathbf{w}_{\tau(1,j)} \dots \mathbf{w}_{\tau(K,j)}} \left[\left\| \sum_{l=1}^K g(\mathbf{w}_{l,j}, \xi_{l,j}) \right\|_2^2 \right] \\ & \leq \frac{K\sigma^2}{m} + \left(\frac{M_G}{m} + K \right) \left\| \sum_{l=1}^K \nabla F(\mathbf{w}_{\tau(l,j)}) \right\|_2^2. \end{aligned} \quad (46)$$

Proof. First let us consider the expectation of any cross term such that $l \neq l'$. For the ease of writing, let $\Omega = \{\mathbf{w}_{\tau(1,j)} \dots \mathbf{w}_{\tau(K,j)}\}$.

Now observe the conditional expectation of the cross term as follows:

$$\begin{aligned} & \mathbb{E}_{\xi_{1,j}, \dots, \xi_{K,j} | \Omega} [(g(\mathbf{w}_{l,j}, \xi_{l,j}) - \nabla F(\mathbf{w}_{\tau(l,j)}))^T (g(\mathbf{w}_{l',j}, \xi_{l',j}) - \nabla F(\mathbf{w}_{\tau(l',j)}))] \\ & = \mathbb{E}_{\xi_{l,j}, \xi_{l',j} | \Omega} [(g(\mathbf{w}_{l,j}, \xi_{l,j}) - \nabla F(\mathbf{w}_{\tau(l,j)}))^T (g(\mathbf{w}_{l',j}, \xi_{l',j}) - \nabla F(\mathbf{w}_{\tau(l',j)}))] \\ & = \mathbb{E}_{\xi_{l',j} | \Omega} [\mathbb{E}_{\xi_{l,j} | \xi_{l',j}, \Omega} [(g(\mathbf{w}_{l,j}, \xi_{l,j}) - \nabla F(\mathbf{w}_{\tau(l,j)}))^T (g(\mathbf{w}_{l',j}, \xi_{l',j}) - \nabla F(\mathbf{w}_{\tau(l',j)}))] \\ & = \mathbb{E}_{\xi_{l',j} | \Omega} [0^T (g(\mathbf{w}_{l',j}, \xi_{l',j}) - \nabla F(\mathbf{w}_{\tau(l',j)}))] = 0. \end{aligned} \quad (47)$$

Thus the cross terms are all 0. So the expression simplifies as,

$$\begin{aligned} & \mathbb{E}_{\xi_{1,j}, \dots, \xi_{K,j} | \Omega} \left[\left\| \sum_{l=1}^K g(\mathbf{w}_{l,j}, \xi_{l,j}) - F(\mathbf{w}_{\tau(l,j)}) \right\|_2^2 \right] \\ & \stackrel{(a)}{=} \sum_{l=1}^K \mathbb{E}_{\xi_{1,j}, \dots, \xi_{K,j} | \Omega} [\|g(\mathbf{w}_{l,j}, \xi_{l,j}) - F(\mathbf{w}_{\tau(l,j)})\|_2^2] \\ & \leq \sum_{l=1}^K \frac{\sigma^2}{m} + \frac{M_G}{m} \|\nabla F(\mathbf{w}_{\tau(l,j)})\|_2^2. \end{aligned} \quad (48)$$

Thus,

$$\begin{aligned} & \mathbb{E}_{\xi_{1,j}, \dots, \xi_{K,j} | \Omega} \left[\left\| \sum_{l=1}^K g(\mathbf{w}_{l,j}, \xi_{l,j}) \right\|_2^2 \right] \\ & = \mathbb{E}_{\xi_{1,j}, \dots, \xi_{K,j} | \Omega} \left[\left\| \sum_{l=1}^K g(\mathbf{w}_{l,j}, \xi_{l,j}) - F(\mathbf{w}_{\tau(l,j)}) \right\|_2^2 \right] + \mathbb{E}_{\xi_{1,j}, \dots, \xi_{K,j} | \Omega} \left[\left\| \sum_{l=1}^K F(\mathbf{w}_{\tau(l,j)}) \right\|_2^2 \right] \\ & \leq \frac{K\sigma^2}{m} + \sum_{l=1}^K \frac{M_G}{m} \|F(\mathbf{w}_{\tau(l,j)})\|_2^2 + \left\| \sum_{l=1}^K F(\mathbf{w}_{\tau(l,j)}) \right\|_2^2 \\ & \leq \frac{K\sigma^2}{m} + \sum_{l=1}^K \frac{M_G}{m} \|F(\mathbf{w}_{\tau(l,j)})\|_2^2 + \sum_{l=1}^K K \|F(\mathbf{w}_{\tau(l,j)})\|_2^2. \end{aligned} \quad (49)$$

□

Now we return to the proof of the theorem.

Proof of Theorem 3. Let $\mathbf{v}_j = \frac{1}{K} \sum_{l=1}^K g(\mathbf{w}_{l,j}, \xi_{l,j})$. Following steps similar to the Async-SGD proof, from Lipschitz continuity we have the following.

$$\begin{aligned}
F(\mathbf{w}_{j+1}) &\leq F(\mathbf{w}_j) + (\mathbf{w}_{j+1} - \mathbf{w}_j)^T \nabla F(\mathbf{w}_j) + \frac{L}{2} \|\mathbf{w}_{j+1} - \mathbf{w}_j\|_2^2 \\
&= F(\mathbf{w}_j) - \frac{\eta}{K} \sum_{l=1}^K g(\mathbf{w}_{l,j}, \xi_{l,j})^T \nabla F(\mathbf{w}_j) + \frac{L}{2} \|\eta \mathbf{v}_j\|_2^2 \\
&\stackrel{(a)}{=} F(\mathbf{w}_j) - \frac{\eta}{2K} \sum_{l=1}^K \|\nabla F(\mathbf{w}_j)\|_2^2 - \frac{\eta}{2K} \sum_{l=1}^K \|g(\mathbf{w}_{l,j}, \xi_{l,j})\|_2^2 \\
&\quad + \frac{\eta}{2K} \sum_{l=1}^K \|g(\mathbf{w}_{l,j}, \xi_{l,j})\|_2^2 - \frac{\eta}{2K} \sum_{l=1}^K \|\nabla F(\mathbf{w}_j)\|_2^2 + \frac{L\eta^2}{2} \|\mathbf{v}_j\|_2^2 \\
&= F(\mathbf{w}_j) - \frac{\eta}{2} \|\nabla F(\mathbf{w}_j)\|_2^2 - \frac{\eta}{2K} \sum_{l=1}^K \|g(\mathbf{w}_{l,j}, \xi_{l,j})\|_2^2 \\
&\quad + \frac{\eta}{2K} \sum_{l=1}^K \|g(\mathbf{w}_{l,j}, \xi_{l,j}) - \nabla F(\mathbf{w}_j)\|_2^2 + \frac{L\eta^2}{2} \|\mathbf{v}_j\|_2^2. \tag{50}
\end{aligned}$$

Here (a) follows from $2\mathbf{a}^T \mathbf{b} = \|\mathbf{a}\|_2^2 + \|\mathbf{b}\|_2^2 - \|\mathbf{a} - \mathbf{b}\|_2^2$. Taking expectation,

$$\begin{aligned}
\mathbb{E}[F(\mathbf{w}_{j+1})] &\leq \mathbb{E}[F(\mathbf{w}_j)] - \frac{\eta}{2} \mathbb{E}[\|\nabla F(\mathbf{w}_j)\|_2^2] - \frac{\eta}{2K} \sum_{l=1}^K \mathbb{E}[\|g(\mathbf{w}_{l,j}, \xi_{l,j})\|_2^2] \\
&\quad + \frac{\eta}{2K} \sum_{l=1}^K \mathbb{E}[\|\nabla F(\mathbf{w}_j) - g(\mathbf{w}_{l,j}, \xi_{l,j})\|_2^2] + \frac{L\eta^2}{2} \mathbb{E}[\|\mathbf{v}_j\|_2^2] \\
&\stackrel{(a)}{\leq} \mathbb{E}[F(\mathbf{w}_j)] - \frac{\eta}{2} \mathbb{E}[\|\nabla F(\mathbf{w}_j)\|_2^2] - \frac{\eta}{2K} \sum_{l=1}^K \mathbb{E}[\|g(\mathbf{w}_{l,j}, \xi_{l,j})\|_2^2] \\
&\quad + \frac{\eta}{2K} \sum_{l=1}^K \mathbb{E}[\|g(\mathbf{w}_{l,j}, \xi_{l,j})\|_2^2] - \frac{\eta}{2K} \sum_{l=1}^K \mathbb{E}[\|\nabla F(\mathbf{w}_{\tau(l,j)})\|_2^2] \\
&\quad + \frac{\eta}{2K} \sum_{l=1}^K \mathbb{E}[\|\nabla F(\mathbf{w}_j) - \nabla F(\mathbf{w}_{\tau(l,j)})\|_2^2] + \frac{L\eta^2}{2} \mathbb{E}[\|\mathbf{v}_j\|_2^2] \tag{51} \\
&\stackrel{(b)}{\leq} \mathbb{E}[F(\mathbf{w}_j)] - \frac{\eta}{2} \mathbb{E}[\|\nabla F(\mathbf{w}_j)\|_2^2] - \frac{\eta}{2K} \sum_{l=1}^K \mathbb{E}[\|\nabla F(\mathbf{w}_{\tau(l,j)})\|_2^2] + \frac{\eta}{2} \gamma \mathbb{E}[\|\nabla F(\mathbf{w}_j)\|_2^2] \\
&\quad + \frac{L\eta^2}{2} \mathbb{E}[\|\mathbf{v}_j\|_2^2] \\
&\stackrel{(c)}{\leq} \mathbb{E}[F(\mathbf{w}_j)] - \frac{\eta}{2} (1 - \gamma) \mathbb{E}[\|\nabla F(\mathbf{w}_j)\|_2^2] + \frac{L\eta^2 \sigma^2}{2Km} \\
&\quad - \frac{\eta}{2K} \sum_{l=1}^K \left(1 - L\eta \left(\frac{M_G}{Km} + \frac{1}{K}\right)\right) \mathbb{E}[\|\nabla F(\mathbf{w}_{\tau(l,j)})\|_2^2] \\
&\stackrel{(d)}{\leq} \mathbb{E}[F(\mathbf{w}_j)] - \frac{\eta}{2} (1 - \gamma) \mathbb{E}[\|\nabla F(\mathbf{w}_j)\|_2^2] + \frac{L\eta^2 \sigma^2}{2Km} - \frac{\eta}{4K} \sum_{l=1}^K \mathbb{E}[\|\nabla F(\mathbf{w}_{\tau(l,j)})\|_2^2]
\end{aligned}$$

$$\stackrel{(e)}{\leq} \mathbb{E}[F(\mathbf{w}_j)] - \frac{\eta}{2}(1 - \gamma)\mathbb{E}[\|\nabla F(\mathbf{w}_j)\|_2^2] + \frac{L\eta^2\sigma^2}{2Km} - \frac{\eta}{4}p_0\mathbb{E}[\|\nabla F(\mathbf{w}_j)\|_2^2]. \quad (52)$$

Here step (a) follows from Lemma 10 and step (b) follows from the assumption that

$$\mathbb{E}[\|\nabla F(\mathbf{w}_j) - \nabla F(\mathbf{w}_{\tau(l,j)})\|_2^2] \leq \gamma\mathbb{E}[\|\nabla F(\mathbf{w}_j)\|_2^2]$$

for some constant $\gamma \leq 1$. The next step (c) follows from the Lemma 11 that bounds the variance of the sum of stochastic gradients. Step (d) follows from choosing $\eta < \frac{1}{2L(\frac{M}{Km} + \frac{1}{K})}$ and finally (e) follows from Lemma 9 in Section 4 that says $\mathbb{E}[\|\nabla F(\mathbf{w}_{\tau(l,j)})\|_2^2] \geq p_0\mathbb{E}[\|\nabla F(\mathbf{w}_j)\|_2^2]$ for some non-negative constant p_0 which is a lower bound on the conditional probability that $\tau(l, j) = j$ given all past delays and parameter values.

Finally, since $F(\mathbf{w})$ is strongly convex, using the inequality $2c(F(\mathbf{w}) - F^*) \leq \|\nabla F(\mathbf{w})\|_2^2$ in (52), we finally obtain the desired result. \square

C.4 Extension to Non-Convex case

Proof. Recall the recursion derived in the last proof in (52). After re-arrangement, we obtain the following:

$$\mathbb{E}[\|\nabla F(\mathbf{w}_j)\|_2^2] \leq \frac{2(\mathbb{E}[F(\mathbf{w}_j)] - \mathbb{E}[F(\mathbf{w}_{j+1})])}{\eta\gamma'} + \frac{L\eta\sigma^2}{Km\gamma'}. \quad (53)$$

Taking summation from $j = 0$ to $j = J - 1$, we get,

$$\begin{aligned} \frac{1}{J} \sum_{j=0}^{J-1} \mathbb{E}[\|\nabla F(\mathbf{w}_j)\|_2^2] &\leq \frac{2(\mathbb{E}[F(\mathbf{w}_0)] - \mathbb{E}[F(\mathbf{w}_J)])}{J\eta\gamma'} + \frac{L\eta\sigma^2}{Km\gamma'} \\ &\stackrel{(a)}{\leq} \frac{2(F(\mathbf{w}_0) - F^*)}{J\eta\gamma'} + \frac{L\eta\sigma^2}{Km\gamma'}. \end{aligned} \quad (54)$$

Here (a) follows since we assume \mathbf{w}_0 to be known and also from $\mathbb{E}[F(\mathbf{w}_J)] \geq F^*$. \square

D AdaSync: Derivation of optimal K for achieving the best error-runtime trade-off

D.1 K-sync SGD

For general distributions, we are required to solve for

$$\frac{du(K)}{dK} = \frac{2(F(\mathbf{w}_{start}))}{t\eta} \frac{d\mathbb{E}[X_{K:P}]}{dK} - \frac{L\eta\sigma^2}{K^2m} = 0.$$

For the case of exponential or shifted-exponential distributions, this reduces to a quadratic equation in K as follows:

$$\left(\frac{2(F(\mathbf{w}_{start}))}{t\eta}\right) K^2 = \frac{L\eta\sigma^2\mu}{m}(P - K).$$

To actually solve for this equation, we would need the values of the Lipschitz constant, variance of the gradient etc. which are not always available. So, we propose a heuristic here.

Observe that, the larger is $F(\mathbf{w}_{start})$, the smaller is the value of K required to minimize $u(K)$. We assume that $F(\mathbf{w}_{start})$ is maximum at the beginning of training, *i.e.*, when $\mathbf{w}_{start} = \mathbf{w}_0$. Hence we start with the smallest initial K , e.g., $K_0 = 1$. Subsequently, after every time interval t , we choose K as follows:

$$\frac{K^2}{P - K} = \frac{K_0^2}{P - K_0} \frac{F(\mathbf{w}_0)}{F(\mathbf{w}_{start})},$$

where we assume that K_0 satisfies $\left(\frac{2(F(\mathbf{w}_0))}{t\eta}\right) K_0^2 = \frac{L\eta\sigma^2\mu}{m}(P - K_0)$. Thus, to summarize, our method for gradually increasing K is as follows:

- Start with an initial K_0 (typically 1).
- After time t , update K by solving the quadratic equation: $K^2 = \frac{K_0^2 F(\mathbf{w}_0)}{(P - K_0) F(\mathbf{w}_{start})} (P - K)$.

We also verify that the second derivative is positive, *i.e.*,

$$\frac{d^2 u(K)}{dK^2} = \frac{2(F(\mathbf{w}_{start}))}{t\eta(P - K)^2} + \frac{2L\eta\sigma^2}{K^3 m} > 0.$$

D.2 K-batch-sync SGD

For general distributions, if we assume $\mathbb{E}[T] \approx K\mathbb{E}[X_{1:P}]$ (see Lemma 2) we are required to solve for

$$\frac{du(K)}{dK} = \frac{2(F(\mathbf{w}_{start}))}{t\eta} \frac{dK\mathbb{E}[X_{1:P}]}{dK} - \frac{L\eta\sigma^2}{K^2 m} = 0.$$

This leads to

$$K^2 = \frac{L\eta\sigma^2 t\eta}{2m(F(\mathbf{w}_{start}))\mathbb{E}[X_{1:P}]}.$$

The interesting observation is that even if we do not know the constants or $\mathbb{E}[X_{1:P}]$, we can still use the trick that we used in the previous case. We can start with the smallest initial K , e.g., $K_0 = 1$. Then, after each time interval t , we update K by solving for

$$K^2 = K_0^2 \frac{F(\mathbf{w}_0)}{F(\mathbf{w}_{start})}.$$

Thus, our method of gradually varying synchronicity is as follows:

- Start with an initial K_0 (typically 1).
- After time t , update K as follows: $K = K_0 \sqrt{\frac{F(\mathbf{w}_0)}{F(\mathbf{w}_{start})}}$.

We also verify that the second derivative is positive, *i.e.*,

$$\frac{d^2 u(K)}{dK^2} = \frac{2L\eta\sigma^2}{K^3 m} > 0.$$

D.3 K-async SGD

For general distributions, the runtime of K -async is upper bounded by that of K -sync (see Lemma 3). For exponential distributions, the two become equal and an algorithm similar to K -sync works. Here, we examine the interesting case of shifted-exponential distribution. We approximate $\mathbb{E}[T] \approx \frac{K\Delta}{P} + \frac{K \log P}{P\mu}$ (see Lemma 4). This leads to

$$K^2 = \frac{L\eta\sigma^2 t\eta P\mu}{2m(F(\mathbf{w}_{start}))(\Delta\mu + \log P)}.$$

Thus, we could start with a small K_0 and after each time interval t , we can update K by solving for

$$K^2 = K_0^2 \frac{F(\mathbf{w}_0)}{F(\mathbf{w}_{start})},$$

in a manner similar to K -batch-sync. We can also verify that the second derivative is positive, *i.e.*,

$$\frac{d^2 u(K)}{dK^2} = \frac{2L\eta\sigma^2}{K^3 m} > 0.$$

D.4 K-batch-async SGD

For any general distribution, the runtime is given by $\mathbb{E}[T] = \frac{K\mathbb{E}[X]}{P}$. This leads to

$$K^2 = \frac{L\eta\sigma^2 t\eta P}{2m(F(\mathbf{w}_{start}))(\mathbb{E}[X])}.$$

Again, similar to the previous case, we could start with a small K_0 and after each time interval t , we can update K by solving for

$$K^2 = K_0^2 \frac{F(\mathbf{w}_0)}{F(\mathbf{w}_{start})}.$$

We can also verify that the second derivative is positive, *i.e.*,

$$\frac{d^2 u(K)}{dK^2} = \frac{2L\eta\sigma^2}{K^3 m} > 0.$$

E Additional Experimental Results

Here, we include some additional experimental results to complement the results in the main paper.

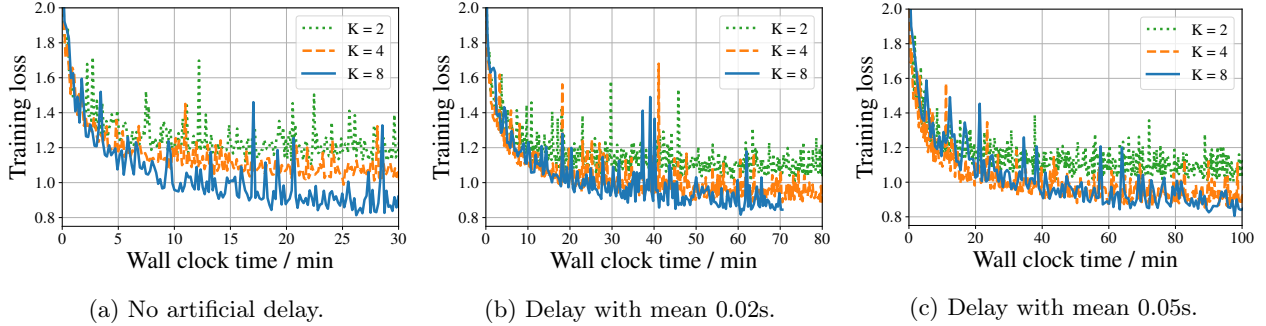


Figure 12: Training loss of **K-sync SGD** on CIFAR-10 with 8 worker nodes.

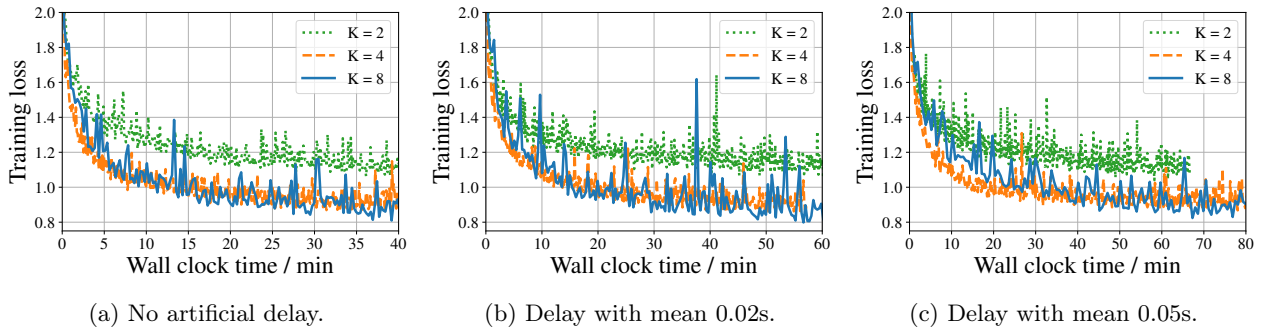


Figure 13: Training loss of **K-async SGD** on CIFAR-10 with 8 worker nodes.