
Sharper Convergence Guarantees for Asynchronous SGD for Distributed and Federated Learning

Anastasia Koloskova
EPFL
anastasia.koloskova@epfl.ch

Sebastian U. Stich
CISPA*
stich@cispa.de

Martin Jaggi
EPFL
martin.jaggi@epfl.ch

Abstract

We study the asynchronous stochastic gradient descent algorithm for distributed training over n workers which have varying computation and communication frequency over time. In this algorithm, workers compute stochastic gradients in parallel at their own pace and return those to the server without any synchronization. Existing convergence rates for this algorithm for non-convex smooth objectives depend on the maximum gradient delay τ_{\max} and show that an ε -stationary point is reached after $\mathcal{O}(\sigma^2\varepsilon^{-2} + \tau_{\max}\varepsilon^{-1})$ iterations, where σ denotes the variance of stochastic gradients.

In this work we obtain (i) a tighter convergence rate of $\mathcal{O}(\sigma^2\varepsilon^{-2} + \sqrt{\tau_{\max}\tau_{\text{avg}}}\varepsilon^{-1})$ without any change in the algorithm, where τ_{avg} is the average delay, which can be significantly smaller than τ_{\max} . We also provide (ii) a simple delay-adaptive learning rate scheme, under which asynchronous SGD achieves a convergence rate of $\mathcal{O}(\sigma^2\varepsilon^{-2} + \tau_{\text{avg}}\varepsilon^{-1})$, and does not require any extra hyperparameter tuning nor extra communications. Our result allows to show *for the first time* that asynchronous SGD is *always faster* than mini-batch SGD. In addition, (iii) we consider the case of heterogeneous functions motivated by federated learning applications and improve the convergence rate by proving a weaker dependence on the maximum delay compared to prior works. In particular, we show that the heterogeneity term in convergence rate is only affected by the average delay within each worker.

1 Introduction

The stochastic gradient descent (SGD) algorithm [43, 13] and its variants (momentum SGD, Adam, etc.) form the foundation of modern machine learning and frequently achieve state of the art results. With recent growth in the size of models and available training data, parallel and distributed versions of SGD are becoming increasingly important [57, 17, 16]. Without those, modern state-of-the-art language models [44], generative models [40, 41], and many others [50] would not be possible. In the distributed setting, also known as data-parallel training, optimization is distributed over many compute devices working in parallel (e.g. cores, or GPUs on a cluster) in order to speed up training. Every worker computes gradients on a subset of the training data, and the resulting gradients are aggregated (averaged) on a server.

The same type of SGD variants also form the core algorithms for federated learning applications [34, 24] where the training process is naturally distributed over many user devices, or clients, that keep their local data private, and only transfer (e.g. encrypted or differentially private) gradients to the server.

A rich literature exists on the convergence theory of above mentioned parallel SGD methods, see e.g. [17, 13] and references therein. Plain parallel SGD still faces many challenges in practice, motivat-

*CISPA Helmholtz Center for Information Security

ing research on various approaches to improve efficiency of distributed learning and mini-batch SGD. This includes for example communication compression techniques [2, 4, 48, 49], decentralized communication [29, 7, 36, 26] or performing several local SGD steps on workers before communicating with the server [30, 32, 34, 47].

These approaches use synchronous communication, where workers in each round are required to wait for the slowest one, before being able to start the next round of computations. In the presence of such straggler nodes or nodes that have different computation speeds, other workers face significant idle times. *Asynchronous* variants of SGD are aimed to solve such inefficiencies and use available workers more effectively. In asynchronous SGD, each worker starts the next computation immediately after finishing computing its own gradient, without waiting for any other workers. This is especially important in the presence of straggler nodes. Asynchronous algorithms were studied both in distributed and federated learning settings [42, 31, 27, 46, 37]. In this paper we focus on such challenging asynchronous variants of SGD and provide an improved theoretical analysis of convergence compared to prior works.

Most existing work has studied the convergence behavior of asynchronous SGD for the setting of homogeneous distributed training data, where worker's objectives are i.i.d.. This assumption however is only realistic e.g. in shared-memory implementations where all processes can access the same data [42]. Under this assumption, it can be proven that asynchronous SGD finds an ε -approximate stationary point (squared gradient norm bounded by ε) in $\mathcal{O}\left(\frac{\sigma^2}{\varepsilon^2} + \frac{\tau_{\max}}{\varepsilon}\right)$ iterations [48], for smooth non-convex objective functions. This complexity bound depends on the maximum delay of the gradients τ_{\max} and the gradient variance $\sigma > 0$. Unfortunately, the maximal delay is a very pessimistic metric, not well reflecting the true behavior in practice. For instance, if a worker struggles just once, the maximum delay is large, while we would still expect reasonable overall convergence.

Two recent works [15, 9] tackle this issue by proposing two new delay-adaptive algorithms that achieve a convergence rate that depends only on the average delay of the applied gradients, with Aviv et al. [9] considering only the convex optimization and Cohen et al. [15] providing a rate of $\mathcal{O}\left(\frac{\sigma^2}{\varepsilon^2} + \frac{\tau_{avg}}{\varepsilon}\right)$ for smooth non-convex functions. The average delay can be much smaller than the maximal delay, and thus these methods are robust to rare stragglers. However, Cohen et al. [15] requires twice more communications at every step, and an extra hyperparameter to tune. Aviv et al. [9] analyze only convex functions and assume a bound on the variance of the delays, which can frequently degrade with the maximum delay τ_{\max} . Moreover, both works require the assumption that gradients are uniformly bounded.

In the realistic case of heterogeneous objective functions, that is in particular relevant in federated learning applications [24], all the existent convergence rates of asynchronous SGD depend on the maximum delay [37].

Contributions.

- For standard asynchronous SGD with constant stepsize, and with non-convex L -smooth homogeneous objective functions, we prove the tighter convergence rate of $\mathcal{O}\left(\frac{\sigma^2}{\varepsilon^2} + \frac{\sqrt{\tau_{avg}\tau_{\max}}}{\varepsilon}\right)$ to ε -small error. Under the additional assumption of bounded gradients, we obtain a convergence rate of $\mathcal{O}\left(\frac{\sigma^2}{\varepsilon^2} + \frac{\tau_{avg}G}{\varepsilon^{3/2}} + \frac{\tau_{avg}}{\varepsilon}\right)$ where G is the bound on the norm of gradients. The previously best known rate was $\mathcal{O}\left(\frac{\sigma^2}{\varepsilon^2} + \frac{\tau_{\max}}{\varepsilon}\right)$.
- With homogeneous objective functions, we provide a delay-adaptive stepsize scheme that does not require tuning of any extra hyperparameters, and converges at the rate of $\mathcal{O}\left(\frac{\sigma^2}{\varepsilon^2} + \frac{\tau_{avg}}{\varepsilon}\right)$ for non-convex L -smooth functions.
- This result allows us to show that asynchronous SGD is always better than mini-batch SGD regardless of the delays pattern (under assumption that the server can perform operations with zero time).
- We also consider distributed optimization with heterogeneous objectives where the delays can depend on the workers and give the convergence rate of $\mathcal{O}\left(\frac{\sigma^2}{\varepsilon^2} + \frac{\zeta^2}{\varepsilon^2} + \frac{\sqrt{\tau_{avg}\frac{1}{n}\sum_{i=1}^n \zeta_i^2 \tau_{avg}^i}}{\varepsilon^{3/2}} + \frac{\sqrt{\tau_{avg}\tau_{\max}}}{\varepsilon}\right)$, where ζ_i 's measure functions heterogeneity and $\bar{\tau}_i$ is the average delay of worker i . This rate improves over the best previously-known results that had worse dependence on the maximum delay τ_{\max} .

2 Related Work

Asynchronous SGD. The research field of asynchronous optimization can be traced back at least to 1989 [11]. Recent works are heavily focused on its SGD variants, such as Hogwild! SGD [39] which deals with coordinate-wise asynchrony. Nguyen et al. [38] provided a tighter convergence analysis by removing the bounded gradient assumption. Our work does not focus on such a coordinate-wise asynchrony as it relies on sparsity assumption that is not realistic in modern machine learning applications. Mania et al. [31] introduces the perturbed iterate framework which enabled theoretical advances with tighter convergence rates [48, 46]. Leblond et al. [27] focus on asynchronous variance-reduced methods.

Many works [1, 14, 20, 6, 45, 28, 48, 18] focused on asynchronous SGD variants where workers communicate with the server without any synchronization, but these communications are considered to be atomic. All of these works provide convergence guarantees that depend on the maximum delay τ_{\max} with [6, 48] providing the first tight convergence rates under assumption that the delays are always constant for quadratic and general (convex, strongly convex and non-convex) functions correspondingly. Stich et al. [46] showed a connection of large batches and delays, although still depending only on the maximum delay. Even et al. [19] consider a continuized view of the time (rather than classical per-iteration time) for asynchronous algorithms on a decentralized network. [3] improve the convergence rates of asynchronous SGD to depend on $\sqrt{\tau_{\max}n}$, however assuming bounded gradients. Under bounded gradients our convergence rates completely remove dependence on the maximum delay τ_{\max} .

Delay-adaptive methods. The works [56, 55, 45, 51, 33, 18] considered delay-adaptive schemes to mitigate adversarial effect of stragglers, however with convergence rates that still depend on the maximum delay τ_{\max} . Only Cohen et al. [15] in the non-convex, and Aviv et al. [9] in the convex case were able to obtain convergence rates depending on the average delay τ_{avg} . Concurrent to our work, Mishchenko et al. [35] provide a delay-adaptive stepsize scheme and derive convergence guarantees similar to ours. Similar to us, they considered asynchronous SGD with constant stepsizes, but under different assumptions: assuming only Lipschitz-continuity of functions instead of Lipschitz-smoothness. They did not discuss the connection of the number of workers to the average delay. For the heterogeneous case they chose a different approach than ours and provide a delay-adaptive learning rate that converges only to an approximate solution, but allows workers to be arbitrarily long delayed (including the case when some of the workers are never responding).

Asynchronous federated learning. In typical federated learning (FL) applications [34], clients or workers frequently have very different computing powers/speed. This makes especially appealing for practitioners to use asynchronous algorithms for FL [47, 37, 54, 8, 53, 23, 10, 22, 52] with many of these works focusing on correcting for unequal participation ratio of different clients [52, 22, 23, 10, 53] by implementing variance reduction techniques on the server. Nguyen et al. [37] introduce the FedBuff algorithm that is very close to the algorithm that we consider in this work and show its practical superiority over classical synchronous FL algorithms.

3 Setup

We consider optimization problems where the components of the objective function (i.e. the data for machine learning problems) is distributed across n workers (or clients),

$$\min_{\mathbf{x} \in \mathbb{R}^d} \left[f(\mathbf{x}) := \frac{1}{n} \sum_{i=1}^n [f_i(\mathbf{x}) := \mathbb{E}_{\xi \sim \mathcal{D}_i} F_i(\mathbf{x}, \xi)] \right]. \quad (1)$$

Here $f_i: \mathbb{R}^d \rightarrow \mathbb{R}$ denotes the local objective function that is accessible to the worker i , $i \in [n] := \{1, \dots, n\}$. Each f_i is a stochastic function $f_i(\mathbf{x}) = \mathbb{E}_{\xi \sim \mathcal{D}_i} F_i(\mathbf{x}, \xi)$ and clients can only access stochastic gradients $\nabla F_i(\mathbf{x}, \xi)$. This setting covers deterministic optimization if $F_i(\mathbf{x}, \xi) = f_i(\mathbf{x})$, $\forall \xi$. It also covers *empirical risk minimization* problems by setting \mathcal{D}_i being a uniform distribution over a local dataset $\{\xi_i^1 \dots \xi_i^{m_i}\}$ of size m_i . In this case the local functions f_i can be written as finite sums: $f_i(\mathbf{x}) = \frac{1}{m_i} \sum_{j=1}^{m_i} F_i(\mathbf{x}, \xi_i^j)$.

Assumptions. For our convergence analysis we rely on following standard assumptions on the functions f_i and F_i :

Algorithm 1 ASYNCHRONOUS SGD

input Initial value $\mathbf{x}^{(0)} \in \mathbb{R}^d$

- 1: server selects a set of active workers $\mathcal{C}_0 \subseteq [n]$ and sends them $\mathbf{x}^{(0)}$
 - 2: **for** $t = 0, \dots, T - 1$ **do**
 - 3: active workers \mathcal{C}_t are computing stochastic gradients in parallel at the assigned points
 - 4: once a worker j_t finishes compute, it sends $\nabla F(\mathbf{x}^{(t-\tau_t)}, \xi_{t-\tau_t})$ to the server
 - 5: server updates $\mathbf{x}^{(t+1)} = \mathbf{x}^{(t)} - \eta_t \nabla F(\mathbf{x}^{(t-\tau_t)}, \xi_{t-\tau_t})$
 - 6: server selects subset $\mathcal{A}_t \subseteq [n]$ of inactive workers, i.e. $(\mathcal{C}_t \setminus \{j_t\}) \cap \mathcal{A}_t = \emptyset$, and sends them $\mathbf{x}^{(t+1)}$
 - 7: update active worker set $\mathcal{C}_{t+1} = \mathcal{C}_t \setminus \{j_t\} \cup \mathcal{A}_t$
 - 8: **end for**
-

Assumption 1 (Bounded variance). *We assume that there exists a constant $\sigma \geq 0$ such that*

$$\mathbb{E}_{\xi \sim \mathcal{D}_i} \|\nabla F_i(\mathbf{x}, \xi) - \nabla f_i(\mathbf{x})\| \leq \sigma^2, \quad \forall i \in [n], \forall \mathbf{x} \in \mathbb{R}^d. \quad (2)$$

Assumption 2 (Bounded function heterogeneity). *We assume that there exists n constants $\zeta_i \geq 0$, $i \in [n]$ such that*

$$\|\nabla f_i(\mathbf{x}) - \nabla f(\mathbf{x})\|_2^2 \leq \zeta_i^2, \quad \forall \mathbf{x} \in \mathbb{R}^d, \quad \text{and define } \zeta^2 := \frac{1}{n} \sum_{i=1}^n \zeta_i^2. \quad (3)$$

Assumption 3 (L -smoothness). *Each function $f_i: \mathbb{R}^d \rightarrow \mathbb{R}$, $i \in [n]$ is differentiable and there exists a constant $L \geq 0$ such that*

$$\|\nabla f_i(\mathbf{y}) - \nabla f_i(\mathbf{x})\| \leq L \|\mathbf{x} - \mathbf{y}\|. \quad \forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^d. \quad (4)$$

For only *some* of our results we will assume a uniform bound on the gradient norm:

Assumption 4 (Bounded gradient). *Each function $f_i: \mathbb{R}^d \rightarrow \mathbb{R}$, $i \in [n]$ is differentiable and there exists a constant $G \geq 0$ such that*

$$\|\nabla f_i(\mathbf{x})\|_2^2 \leq G^2, \quad \forall \mathbf{x} \in \mathbb{R}^d. \quad (5)$$

4 Homogeneous Distributed Setting

We start with an important special case of problem (1) where the objective functions are identical for all workers, i.e. $f_i(\mathbf{x}) \equiv f_j(\mathbf{x})$ for all $i, j \in [n]$, such as in the case of homogeneously (i.i.d.) distributed training data. Consequently, this implies that Assumption 2 holds with $\zeta_i = 0$, $i \in [n]$. Many classical works have focused on asynchronous algorithms under this homogeneous setting (e.g. [6, 48, 1, 20, 45, 28, 46], see the related work for more references). This setting commonly appears in the datacenter setup for distributed training [16], where all nodes (or GPUs) have access to the full dataset or data distribution. Moreover, this special case allows us to present our main ideas in a simplified way, without complicating the presentation due to data heterogeneity. We will later see that most of the results in this section can also be obtained as a corollary of the more general heterogeneous functions case (Section 5) by setting $\zeta_i = 0$ $i \in [n]$.

4.1 Algorithm

We consider standard asynchronous SGD (also known as delayed SGD, or SGD with stale updates) as presented in Algorithm 1, see e.g. [6, 48, 1, 20, 45, 28, 46]. First, the server initializes training by selecting an initial active worker set \mathcal{C}_0 and assigning $\mathbf{x}^{(0)}$ to these workers. Throughout the algorithm, the active workers compute gradients at their own speed, based on their local data. On line 4, once some worker (which we denote as j_t) finishes computing its gradient, it sends the result to the server. On line 5 the server incorporates the received—possibly delayed—gradient, using a stepsize η_t that can depend on the gradient delay τ_t . The *gradient delay* τ_t is defined as the difference between the iteration at which worker j_t started to compute the gradient and the iteration t at which it got applied. We index the stochastic noise of the gradients ξ_t by iteration t to highlight that previous iterates $\mathbf{x}^{(t')}$ for $t' \leq t$ do not depend on this stochastic noise. However, the client selects the data sample ξ_t at iteration $t - \tau_t$ when the computation starts. After that, on lines 6–7 the server selects the new active workers out of the ones that are currently inactive (including worker j_t) and assigns them the latest iterate $\mathbf{x}^{(t+1)}$.

In contrast to previous works, we explicitly define the set of workers that are busy with computations at every step t as \mathcal{C}_t (the active workers set). Note that this does not pose any restrictions. A main advantage of allowing the sets \mathcal{C}_t to be different at every step t lies in the possibility to also cover mini-batch SGD as a special case, which we discuss in Example 2. Our theoretical results depend on the size of these sets \mathcal{C}_t , a.k.a. the *concurrency*.

Definition 1 (Concurrency). *The concurrency $\tau_C^{(t)}$ at step t is defined as the size of the active worker set \mathcal{C}_t , i.e. $\tau_C^{(t)} = |\mathcal{C}_t|$. We also define the maximum and average concurrency as*

$$\tau_C = \max_t \{\tau_C^{(t)}\}, \quad \bar{\tau}_C = \frac{1}{T+1} \sum_{t=0}^T \tau_C^{(t)}.$$

Note that in many practical scenarios, we have a *constant concurrency* of n over time, meaning that all n workers are active at every step, and thus $\tau_C = \bar{\tau}_C = n$.

We discuss two important practical examples that fit into our Algorithm 1:

Example 2 (Mini-batch SGD). *Mini-batch SGD with batch size n can be seen as a special case of Algorithm 1, as follows: The server (i) in line 1 selects all n workers, $\mathcal{C}_0 = [n]$; (ii) in line 6 does not select new workers while the gradients from the same batch have not been fully applied yet, i.e. $\mathcal{A}_t = \emptyset$ if $t \bmod n \neq 0$; (iii) in line 6 selects $\mathcal{A}_t = [n]$ if $t \bmod n = 0$ to start a new batch.*

Example 3 (Asynchronous SGD with maximum concurrency). *In practical implementations one should always aim to utilize all resources available and thus (i) in line 1 select all available workers $\mathcal{C}_0 = [n]$; (ii) in line 6 select to re-assign the worker that finished its computations $\mathcal{A}_t = \{j_t\}$ so that workers are always kept busy with jobs.*

4.2 Theoretical analysis: Constant stepsizes

We first formally define the average and maximum delays.

Definition 4 (Average and maximum delays). *Let $\{\tau_t\}_{t=0}^{T-1}$ be the delays of the applied gradients in Algorithm 1. We define $\{\tau_i^{C_T}\}_{i \in \mathcal{C}_T \setminus \{j_T\}}$ as the delays of gradients which are in flight at time T , that is they have remained unapplied at the last step. Each $\tau_i^{C_T}$ is equal to the difference between the last iteration T and the iteration at which worker i started to compute its last gradient. We then define the average and the maximum delays as*

$$\tau_{avg} = \frac{1}{T + |\mathcal{C}_T| - 1} \left(\sum_{t=0}^{T-1} \tau_t + \sum_{i \in \mathcal{C}_T \setminus \{j_T\}} \tau_i^{C_T} \right), \quad \tau_{max} = \max \left\{ \max_{t=1, \dots, T-1} \tau_t, \max_{i \in \mathcal{C}_T \setminus \{j_T\}} \tau_i^{C_T} \right\}. \quad (6)$$

We further provide a key observation on the connection between the average delay and the average concurrency. This observation, is one of the essential elements for achieving an improved analysis.

Remark 5 (Key Observation). *In Algorithm 1 the average concurrency $\bar{\tau}_C$ is connected to the average delay τ_{avg} as*

$$\tau_{avg} = \frac{T+1}{T + |\mathcal{C}_T| - 1} \bar{\tau}_C \stackrel{T > |\mathcal{C}_T|}{=} \Theta(\bar{\tau}_C). \quad (7)$$

We explain this observation on a simple example. Assume that the concurrency is constant at every step ($\tau_C = \bar{\tau}_C$), and that all workers except one are responding very rarely. Then on steps 4–5 of Algorithm 1 only this one responding worker would mostly participate. This means that for this one worker the delay τ_t would be frequently equal to zero, and the overall average delay will be small.

Next, we provide our theoretical results. We first focus on the Asynchronous SGD Algorithm 1 under constant stepsizes, i.e. $\eta_t \equiv \eta$. This setting was studied in many works such as [1, 20, 6, 28, 48].

Theorem 6 (Constant stepsizes). *Under Assumptions 1, 3, there exists a constant stepsize $\eta_t \equiv \eta$ such that for Algorithm 1 it holds that $\frac{1}{T+1} \sum_{t=0}^T \|\nabla f(\mathbf{x}^{(t)})\|_2^2 \leq \varepsilon$ after*

$$\mathcal{O} \left(\frac{\sigma^2}{\varepsilon^2} + \frac{\sqrt{\tau_C \tau_{max}}}{\varepsilon} \right) \quad \text{iterations.} \quad (8)$$

If we additionally assume bounded gradient Assumption 4, then $\frac{1}{\sum_{t=0}^T |\mathcal{A}_t|} \sum_{t=0}^T |\mathcal{A}_t| \|\nabla f(\mathbf{x}^{(t)})\|_2^2 \leq \varepsilon$ after

$$\mathcal{O} \left(\frac{\sigma^2}{\varepsilon^2} + \frac{\tau_C G}{\varepsilon^{3/2}} + \frac{\tau_C}{\varepsilon} \right) \quad \text{iterations.} \quad (9)$$

Under constant concurrency, we can directly connect τ_C to the average delay τ_{avg} due to Remark 5. We highlight again that in practice, to get the best utilization of the available resources, practical implementations choose the maximum concurrency possible, which is equal to n .

Corollary 7. *If in Algorithm 1 the concurrency is constant at every step (thus $\tau_C = \bar{\tau}_C$), then under the same conditions as in Theorem 6 the convergence rate of Algorithm 1 is*

$$\mathcal{O}\left(\frac{\sigma^2}{\varepsilon^2} + \frac{\sqrt{\tau_{avg}\tau_{max}}}{\varepsilon}\right) \quad \text{and} \quad \mathcal{O}\left(\frac{\sigma^2}{\varepsilon^2} + \frac{\tau_{avg}G}{\varepsilon^{3/2}} + \frac{\tau_{avg}}{\varepsilon}\right) \quad (10)$$

for the case without and with bounded gradients (Assumption 4) respectively.

The previously best known convergence rate for Asynchronous SGD (Algorithm 1) under constant stepsizes was given in [48] and is equal to $\mathcal{O}\left(\frac{\sigma^2}{\varepsilon^2} + \frac{\tau_{max}}{\varepsilon}\right)$. In our theorem we improved the delay dependence from τ_{max} to $\sqrt{\tau_{avg}\tau_{max}}$ in the last term *without any change in the algorithm*, only by taking into account concurrency that is usually fixed in practical implementations anyways. No other work previously made an assumption on the number of computing workers in their theoretical analysis. $\sqrt{\tau_{avg}\tau_{max}}$ could be much smaller than τ_{max} in the presence of rare straggler devices. With an additional assumption of bounded gradients, the dependence on the maximum delay can be completely removed.

4.3 Theoretical analysis: Delay-adaptive stepsizes

In many cases, the bounded gradient Assumption 4 is unrealistic [38], meaning that the gradient bound G is often large and thus the rate (9) is loose. In this section we show that by weighting the stepsize down for the gradients that have a large delay, once can remove the dependence on the maximum delay τ_{max} without assuming bounded gradients (Assump. 4).

Theorem 8 (Delay-adaptive stepsizes). *There exist a parameter $\eta \leq \frac{1}{4L}$ such that if we set the stepsizes in Algorithm 1 dependent on the delays as*

$$\eta_t = \begin{cases} \eta & \tau_t \leq 2\tau_C, \\ < \min\{\eta, \frac{1}{4L\tau_t}\} & \tau_t > 2\tau_C, \end{cases} \quad (11)$$

then for Algorithm 1, under Assumptions 1, 3 it holds that $\frac{1}{\sum_{t=0}^T \eta_t} \sum_{t=0}^T \eta_t \|\nabla f(\mathbf{x}^{(t)})\|_2^2 \leq \varepsilon$ after

$$\mathcal{O}\left(\frac{\sigma^2}{\varepsilon^2} + \frac{\tau_C}{\varepsilon}\right) \quad \text{iterations.} \quad (12)$$

In our theorem, the stepsize η_t in the case of large delays $\tau_t > 2\tau_C$ can be an arbitrary value between 0 and $\min\{\eta, \frac{1}{4L\tau_t}\}$. Setting the stepsize $\eta_t \equiv 0$ is equivalent to dropping these gradients.

Proof sketch of Theorem 8. We give the intuitive proof sketch for the case when we drop gradients with $\tau_t > 2\tau_C$ and we deal with the general case in the Appendix. We know that $\tau_{avg} \approx \bar{\tau}_C \leq \tau_C$ from Remark 5. It also holds that the number of gradients that have delay larger than the two times the average delay $2\tau_{avg}$ is smaller than half of all the gradients ($\leq \frac{T}{2}$) because delays are bounded below by zero ($\tau_t \geq 0 \forall t$). Thus, dropping the gradients with the delay $\tau_t > 2\tau_C$, or equivalently setting their stepsize $\eta_t \equiv 0$, will degrade the convergence rate at most by half, while the maximum delay among the applied ones now is equal to $2\tau_C$. Thus we can apply the result from [48] with $\tau_{max} = 2\tau_C$. \square

Corollary 9. *If in Algorithm 1 the concurrency is constant at every step (thus $\tau_C = \bar{\tau}_C$), then under the same conditions as in Theorem 8 the convergence rate of Algorithm 1 is*

$$\mathcal{O}\left(\frac{\sigma^2}{\varepsilon^2} + \frac{\tau_{avg}}{\varepsilon}\right). \quad (13)$$

4.4 Discussion

Comparison to synchronous optimization. Mini-batch SGD with batch size n has the same degree of parallelism as Algorithm 1 with constant concurrency n , i.e. it has n workers computing gradients in parallel. Mini-batch SGD needs $\mathcal{O}\left(\frac{\sigma^2}{n\varepsilon^2} + \frac{1}{\varepsilon}\right)$ [21] batches of gradients to reach an ε -stationary point, and thus needs $\mathcal{O}\left(\frac{\sigma^2}{\varepsilon^2} + \frac{n}{\varepsilon}\right)$ gradients, as the batch-size is equal to n . On the

contrary, asynchronous SGD Algorithm 1 with stepsizes chosen as in (11) achieves exactly the same rate (13) since $\tau_{avg} = \tau_C = n$, while its expected per-iteration time is *faster* than that of mini-batch SGD, as no workers have to wait for others. Thus, our result shows that asynchronous SGD is always faster than mini-batch SGD *regardless of the delay pattern*. A small note that in our reasoning we implicitly assumed that the sever can perform its operations in negligible time.

Tuning the stepsize. It is worth noting that our stepsize rule (11) does not introduce any additional hyperparameters to tune compared to the constant stepsize case or to synchronous SGD. τ_C is usually known and can be easily controlled by the server, especially in the practical constant concurrency case. Thus, to implement such a stepsize rule (11) one needs to tune only stepsize η , and in case of $\tau_t > 2\tau_C$ set stepsize $\eta_t \leq \frac{\eta}{\tau_t}$.

Average v.s. maximum delay. In a homogeneous environment when every worker computes gradients with same speed during the whole training, the average and maximum delays would be almost equal. However, occasional straggler devices will usually be present. In this case the maximum delay is much larger than the average delay.

Consider a simple example with $n = 2$ workers, where the first worker computes gradients very fast, while the second worker returns its gradient only at the end of the training at the last iteration T . In this case the average delay $\tau_{avg} = 2$ is a small constant, while the maximum delay $\tau_{max} = T$. In this case the rate depending only on the maximum delay τ_{max} would guarantee convergence only up to a constant accuracy $\varepsilon = \mathcal{O}(1)$. While both rates with $\sqrt{\tau_{max}\tau_{avg}}$ and with τ_{avg} guarantee convergence up to an arbitrary small accuracy.

Comparison to other methods. Cohen et al. [15] recently proposed the PickySGD algorithm that achieves a convergence rate of $\mathcal{O}\left(\frac{\sigma^2}{\varepsilon^2} + \frac{\tau_{avg}}{\varepsilon}\right)$ (same as (13)). Their algorithm discards gradients based on the distance between the current point and the delayed one $\|\mathbf{x}^{(t)} - \mathbf{x}^{(t-\tau_t)}\|$. The disadvantage of their method is that it requires sending points $\mathbf{x}^{(t-\tau_t)}$ along with the gradients thus incurring twice more communications at every step. Their method also requires tuning an extra hyperparameter. In this work we achieve the same convergence rate with a much simpler method that does not require any additional communications nor additional tuning compared to synchronous SGD.

[9] also recently proposed the delay-adaptive algorithm with convergence rate depending on the average delay τ_{avg} for the convex and strongly convex cases. Although, our convergence rates are for the non-convex case and are not directly comparable to theirs, we highlight some key differences in their analysis. First, their convergence rate depends not only on τ_{avg} but also on the variance σ_τ of the delays, which can degrade with the maximum delay. Second, they require the bounded gradient Assumption 4. In Theorem 6 we show that under Assumption 4 no modifications to the algorithm are needed to completely remove the dependence on the maximum delay τ_{max} (9).

Tightness. As we explained in Example 2, mini-batch SGD is covered by Algorithm 1. We know that mini-batch SGD convergence is lower bounded by $\Theta\left(\frac{\sigma^2}{n\varepsilon^2} + \frac{1}{\varepsilon}\right)$ [5] in terms of batches processed and thus by $\Theta\left(\frac{\sigma^2}{\varepsilon^2} + \frac{n}{\varepsilon}\right)$ in terms of the gradients computed. Our convergence rate given in Theorem 6 coincides with this lower bound as in this case concurrency $\tau_C = n$, $\tau_{avg} = \bar{\tau}_C = \frac{n}{2}$.

5 Heterogeneous Distributed Setting

In this section we consider more general problems of the form (1) where the functions f_i are different on different workers. This setting is motivated for example by federated learning [34, 24], where every worker (client) possesses its own private data, possibly coming from a different data distribution, and thus has its own different local objective function f_i .

The setting here is therefore more general than the one considered in previous Section 4, and we will see that some of the results (with the constant stepsizes) in the homogeneous case follow as a special case of the more general results we present in this section.

5.1 Algorithm

We consider asynchronous SGD as given in Algorithm 2. Close variants of this algorithm were studied in several prior works [37, 47]. In order to simplify the presentation, we consider that concurrency is constant over time (and thus $\tau_C = \bar{\tau}_C$ in Definition 1). In order to allow for client subsampling

often implemented in practical federated learning applications, we allow the concurrency τ_C to be smaller than overall number of workers n . The same concurrency model was recently considered in the practical FedBuff algorithm [37].

Algorithm 2 ASYNCHRONOUS SGD with concurrency τ_C

input Initial value $\mathbf{x}^{(0)} \in \mathbb{R}^d$, n clients, concurrency τ_C

Server:

- 1: server selects *uniformly at random* a set of active clients \mathcal{C}_0 of size τ_C and sends them $\mathbf{x}^{(0)}$
 - 2: **for** $t = 0, \dots, T - 1$ **do**
 - 3: active clients \mathcal{C}_t are computing stochastic gradients in parallel at the assigned points
 - 4: once some client j_t finishes compute, it sends $\nabla F_{j_t}(\mathbf{x}^{(t-\tau_t)}, \xi_{t-\tau_t})$ to the server
 - 5: server updates $\mathbf{x}^{(t+1)} = \mathbf{x}^{(t)} - \eta_t \nabla F_{j_t}(\mathbf{x}^{(t-\tau_t)}, \xi_{t-\tau_t})$
 - 6: server selects a new client $k_t \sim \text{Uniform}[1, n]$ and sends it $\mathbf{x}^{(t+1)}$
 - 7: update the active worker multiset $\mathcal{C}_{t+1} = \mathcal{C}_t \setminus \{j_t\} \cup \{k_t\}$
 - 8: **end for**
-

The algorithm is very similar to the homogeneous Algorithm 1 with two key differences: at line 6, the server selects clients *out of all clients*, and does so *uniformly at random*, regardless of the current active worker set \mathcal{C}_t . This means that the same client can get sampled several times, even if it didn't finish its previous job(s) yet (thus \mathcal{C}_t is a multiset). In this case, the assigned jobs would just pile up on this client.

5.2 Theoretical analysis

We first note that our key observation on the delays (Remark 5) holds for Algorithm 2 as well. Moreover, as we have a constant concurrency τ_C at every step, $\tau_{avg} = \Theta(\tau_C)$.

Definition 10. Denote a (possibly empty) set $\{\tau_k^{C_T, i}\}_k$ to be the set of delays from gradients of the client i that are left unapplied at the last iteration of the Algorithm 2.

We define the average delay of a client i as

$$\tau_{avg}^i = \frac{1}{T_i} \left(\sum_{t: j_t=i} \tau_t + \sum_k \tau_k^{C_T, i} \right)$$

where T_i is the number of times that client i was sampled during lines 1 and 6 of Algorithm 2.

Assumption 5. The average delay τ_{avg}^i is independent from the number of times T_i that client i was sampled.

Theorem 11 (constant stepsizes). *Under Assumptions 1, 2, 3, 5 there exists a constant stepsize $\eta_t \equiv \eta$ such that for Algorithm 2 it holds that $\frac{1}{T+1} \sum_{t=0}^T \|\nabla f(\mathbf{x}^{(t)})\|_2^2 \leq \varepsilon$ after*

$$\mathcal{O} \left(\frac{\sigma^2}{\varepsilon^2} + \frac{\zeta^2}{\varepsilon^2} + \frac{\sqrt{\tau_{avg} \frac{1}{n} \sum_{i=1}^n \zeta_i^2 \tau_{avg}^i}}{\varepsilon^{\frac{3}{2}}} + \frac{\sqrt{\tau_{avg} \tau_{\max}}}{\varepsilon} \right) \quad \text{iterations}, \quad (14)$$

Under Assumptions 1, 2, 3 and additional bounded gradient Assumption 4, it holds that $\frac{1}{T+1} \sum_{t=0}^T \|\nabla f(\mathbf{x}^{(t)})\|_2^2 \leq \varepsilon$ after

$$\mathcal{O} \left(\frac{\sigma^2}{\varepsilon^2} + \frac{\zeta^2}{\varepsilon^2} + \frac{\tau_{avg} G}{\varepsilon^{\frac{3}{2}}} + \frac{\tau_{avg}}{\varepsilon} \right) \quad \text{iterations}. \quad (15)$$

We note that the leading $\frac{1}{\varepsilon^2}$ term is affected by heterogeneity ζ^2 because at every step we apply gradient from only one client. This term is usually present in the federated learning algorithms with client subsampling see e.g. [25].

5.3 Discussion

Comparison to other works. The recent FedBuff algorithm [37] is similar to our Algorithm 2. Their algorithm allows clients to perform several local steps and the server to wait for more than 1

client to finish compute (aka buffering), which we did not include for simplicity as these aspects are orthogonal to the effect of delays.

Disregarding these two orthogonal changes, the FedBuff algorithm is almost equivalent to our Algorithm 2 with a key difference: they assume that the client j_t that finishes computation at every step comes from the uniform distribution over all the clients. This is unrealistic to assume in practice because the server cannot control which clients finish computations at every step. In Algorithm 2 we have the more realistic assumption only on the sampling process of the new clients (on line 6) that *can be controlled* by the server. This reflects practical client sampling in federated learning.

The convergence rate of FedBuff [37] under the bounded gradient assumption is $\mathcal{O}\left(\frac{\sigma^2}{\varepsilon^2} + \frac{\zeta^2}{\varepsilon^2} + \frac{(\zeta^2+1)\tau_{\max}G^2}{\varepsilon}\right)$. In contrast, in Theorem 11 we completely remove the dependence on the maximum delay τ_{\max} under bounded gradients (as in Equation (15)).

Delays. We note that for Theorem 11 we did not impose any assumption on the delays. Thus, our result allows clients and the delays on these clients to be dependent, meaning that some of the clients could be systematically slower than others. Interestingly, the middle heterogeneity term (the term with ζ_i) is not affected by the maximum delay at all, but is affected by the average delay within each individual client. If all the heterogeneity parameters are equal, i.e. $\zeta_i = \zeta_j, \forall i, j$, then the middle term will be affected only by the overall average delay τ_{avg} .

Gradient clipping. Practical implementations of FL algorithms usually apply clipping to the gradients in order to guarantee differential privacy [24]. This automatically bounds the norm of all applied gradients, making the the constant G^2 in Assumption 4 small. Although we do not provide formal convergence guarantees of asynchronous SGD with gradient clipping, we envision that its convergence rate would depend only on the average delay, similar to the bounded gradient case (9), thus making the algorithm robust to stragglers.

Delay-adaptive stepsizes. For homogeneous functions we have shown that delay-adaptive stepsizes result in a convergence rate dependent only on the average delay τ_{avg} without assuming bounded gradients (as in Equation (11)). However in the heterogeneous case this is not so straightforward. Delay-adaptive learning rate schemes will introduce a bias towards the clients that compute quickly, and Algorithm 2 would converge to the wrong objective.

It is interesting to note that current popular schemes implemented in practice for FL over-selects the clients at every iteration [12]. The server waits only for some percentage (e.g. 80%) of sampled clients and discards the rest. Such a scheme also introduces a bias towards fast workers. A delay-adaptive learning rate scheme is expected to introduce less bias as the gradients are still applied but with the smaller weight. We leave this question for future practical investigations, as it is not the focus of our current work.

Independent delays. If the delays and the clients are independent (e.g. coming from the same distribution for all of the clients), then the convergence rate of Algorithm 2 will simplify to $\mathcal{O}\left(\frac{\sigma^2}{\varepsilon^2} + \frac{\zeta\tau_{avg}}{\varepsilon^{\frac{3}{2}}} + \frac{\sqrt{\tau_{avg}\tau_{\max}}}{\varepsilon}\right)$ (without needing bounded gradient assumption). In this case it is also possible to use delay-adaptive stepsizes (similar to Theorem 8) to completely remove the dependence on the maximum delays τ_{\max} without assuming bounded gradients.

Extensions. We can extend the Algorithm 2 and our theoretical analysis to allow clients to perform several local steps, before sending back the change in \mathbf{x} . We can also extend Algorithm 2 to allow the server to wait for the first K clients to finish computations rather than just one, similar to [37]. These extensions are straightforward and we excluded them here for simplicity of presentation.

Finally, we can also extend Algorithm 2 to sample new clients as soon as some previous client finished compute, without waiting for the server update on the line 5.

6 Experiments

In this section we aim to empirically evaluate the effectiveness of our proposed delay-adaptive stepsizes (11) in the homogeneous case. We use w1a dataset from LIBSVM data library, and evaluate performance on logistic regression function defined as $f(\mathbf{x}) = \frac{1}{m} \sum_{j=1}^m \log(1 + \exp(-b_j \mathbf{a}_j^\top \mathbf{x})) +$

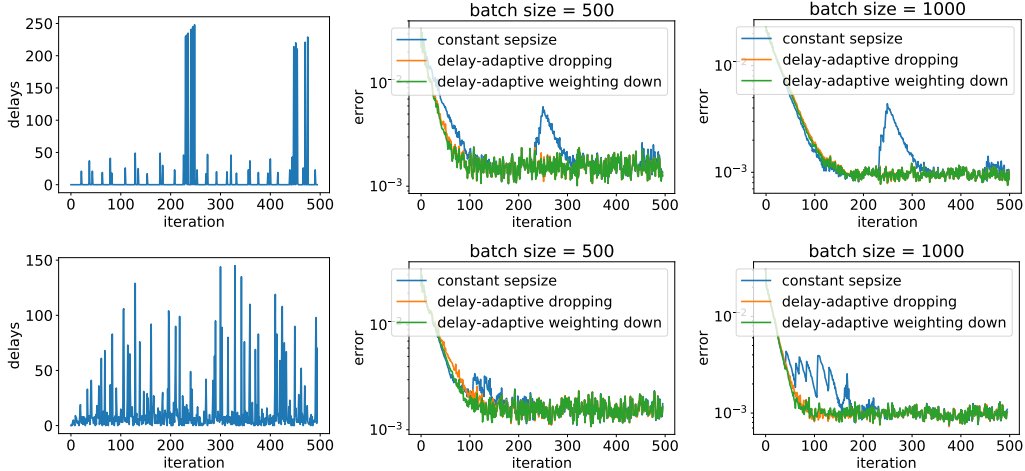


Figure 1: Comparison of three different stepsize strategies in asynchronous SGD. (left) delay pattern, (middle) convergence curve with batch size 500, (right) convergence curves with batch size 1000. Both delay-adaptive stepsize schemes are more robust to occasional large delays.

$\frac{1}{2m} \|\mathbf{x}\|^2$, where $\mathbf{a}_j \in \mathbb{R}^d$ and $b_j \in \{-1, 1\}$ are data samples. w1a dataset has $m = 2477$ examples and dimension $d = 300$. We use fixed concurrency with $n = 10$ workers.

We compare the following three stepsize strategies:

- constant stepsizes $\eta_t \equiv \eta$
- delay-adaptive stepsizes with dropping large-delayed gradients

$$\eta_t = \begin{cases} \eta & \tau_t \leq 2n, \\ 0 & \tau_t > 2n, \end{cases}$$

- delay-adaptive stepsizes with weighting down large-delayed gradients

$$\eta_t = \begin{cases} \eta & \tau_t \leq 2n, \\ \eta \frac{\tau_C}{\tau_t} & \tau_t > 2n, \end{cases}$$

Note that the two last stepsize schemes satisfy the conditions in (11). For each of the strategies we tune only one parameter η separately from the logarithmic grid and ensure that the found stepsize is not on the edge of this grid.

Figure 1 presents the results for two different delay patterns. We can see that both of the delay-adaptive stepsizes have comparable performance and are more robust to the large delays than the constant stepsize.

7 Conclusion

In this paper we study the asynchronous SGD algorithm both in homogeneous and heterogeneous settings. By leveraging the notion of concurrency—the number of workers that compute gradients in parallel—we show a much faster convergence rate for asynchronous SGD, improving the dependence on the maximum delay τ_{\max} over prior works, for both homogeneous and heterogeneous objectives. Our proof technique also allows to design a simple delay-adaptive stepsize rule (11) that attains a convergence rate depending only on the average delay τ_{avg} that neither requires any additional tuning, nor additional communication. Our techniques allows us to demonstrate that *asynchronous SGD is faster than mini-batch SGD for any delay pattern*.

Acknowledgments

AK was supported by a Google PhD Fellowship. We thank Brendan McMahan, Thijs Vogels, Hadrien Hendrikx, Aditya Vardhan Varre and Maria-Luiza Vladarean for useful discussions and their feedback on the manuscript.

References

- [1] Alekh Agarwal and John C Duchi. Distributed delayed stochastic optimization. In *Advances in Neural Information Processing Systems 24*, pages 873–881. Curran Associates, Inc., 2011.
- [2] Dan Alistarh, Demjan Grubic, Jerry Li, Ryota Tomioka, and Milan Vojnovic. QSGD: Communication-efficient SGD via gradient quantization and encoding. In *Advances in Neural Information Processing Systems 30*, pages 1709–1720. Curran Associates, Inc., 2017.
- [3] Dan Alistarh, Christopher De Sa, and Nikola Konstantinov. The convergence of stochastic gradient descent in asynchronous shared memory. In *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing*, PODC '18, page 169–178, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450357951. doi: 10.1145/3212734.3212763.
- [4] Dan Alistarh, Torsten Hoefler, Mikael Johansson, Nikola Konstantinov, Sarit Khirirat, and Cedric Renggli. The convergence of sparsified gradient methods. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *NeurIPS - Advances in Neural Information Processing Systems 31*, pages 5977–5987. Curran Associates, Inc., 2018.
- [5] Yossi Arjevani, Yair Carmon, John C. Duchi, Dylan J. Foster, Nathan Srebro, and Blake E. Woodworth. Lower bounds for non-convex stochastic optimization. *ArXiv*, abs/1912.02365, 2019.
- [6] Yossi Arjevani, Ohad Shamir, and Nathan Srebro. A tight convergence analysis for stochastic gradient descent with delayed updates. In Aryeh Kontorovich and Gergely Neu, editors, *Proceedings of the 31st International Conference on Algorithmic Learning Theory*, volume 117 of *Proceedings of Machine Learning Research*, pages 111–132. PMLR, 08 Feb–11 Feb 2020.
- [7] Mahmoud Assran, Nicolas Loizou, Nicolas Ballas, and Michael Rabbat. Stochastic gradient push for distributed deep learning. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*. PMLR, 2019.
- [8] Dmitrii Avdiukhin and Shiva Kasiviswanathan. Federated learning under arbitrary communication patterns. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 425–435. PMLR, 18–24 Jul 2021.
- [9] Rotem Zamir Aviv, Ido Hakimi, Assaf Schuster, and Kfir Yehuda Levy. Learning under delayed feedback: Implicitly adapting to gradient delays. In *Proceedings of the 38th International Conference on Machine Learning*. PMLR, 2021.
- [10] Arda Aytekin, Hamid Reza Feyzmahdavian, and Mikael Johansson. Analysis and implementation of an asynchronous optimization algorithm for the parameter server, 2016.
- [11] D.P. Bertsekas and J.N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Prentice-Hall, 1989.
- [12] Keith Bonawitz, Hubert Eichner, Wolfgang Grieskamp, Dzmitry Huba, Alex Ingerman, Vladimir Ivanov, Chloé Kiddon, Jakub Konečný, Stefano Mazzocchi, Brendan McMahan, Timon Van Overveldt, David Petrou, Daniel Ramage, and Jason Roselander. Towards federated learning at scale: System design. In A. Talwalkar, V. Smith, and M. Zaharia, editors, *Proceedings of Machine Learning and Systems*, volume 1, pages 374–388, 2019.
- [13] L. Bottou, F. Curtis, and J. Nocedal. Optimization methods for large-scale machine learning. *SIAM Review*, 60(2):223–311, 2018.
- [14] Sorathan Chaturapruek, John C Duchi, and Christopher Ré. Asynchronous stochastic convex optimization: the noise is in the noise and SGD don’t care. In *Advances in Neural Information Processing Systems 28*, pages 1531–1539. Curran Associates, Inc., 2015.
- [15] Alon Cohen, Amit Daniely, Yoel Drori, Tomer Koren, and Mariano Schain. Asynchronous stochastic optimization robust to arbitrary delays. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 9024–9035. Curran Associates, Inc., 2021.
- [16] Jeffrey Dean, Greg Corrado, Rajat Monga, Kai Chen, Matthieu Devin, Mark Mao, Marc’ aurelio Ranzato, Andrew Senior, Paul Tucker, Ke Yang, Quoc V. Le, and Andrew Y. Ng. Large scale distributed deep networks. In *NIPS - Advances in Neural Information Processing Systems*, pages 1223–1231, 2012.

- [17] Ofer Dekel, Ran Gilad-Bachrach, Ohad Shamir, and Lin Xiao. Optimal distributed online prediction using mini-batches. *Journal of Machine Learning Research (JMLR)*, 13(1):165–202, 2012.
- [18] Sanghamitra Dutta, Gauri Joshi, Soumyadip Ghosh, Parijat Dube, and Priya Nagpurkar. Slow and stale gradients can win the race: Error-runtime trade-offs in distributed sgd. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, pages 803–812. PMLR, 2018.
- [19] Mathieu Even, Hadrien Hendrikx, and Laurent Massoulié. Decentralized optimization with heterogeneous delays: a continuous-time approach, 2021.
- [20] H. R. Feyzmahdavian, A. Aytekin, and M. Johansson. An asynchronous mini-batch algorithm for regularized stochastic optimization. *IEEE Transactions on Automatic Control*, 61(12):3740–3754, Dec 2016. ISSN 0018-9286.
- [21] Saeed Ghadimi and Guanghui Lan. Stochastic first- and zeroth-order methods for nonconvex stochastic programming. *SIAM J. Optim.*, 23:2341–2368, 2013.
- [22] Margalit Glasgow and Mary Wootters. Asynchronous distributed optimization with stochastic delays, 2020.
- [23] Xinran Gu, Kaixuan Huang, Jingzhao Zhang, and Longbo Huang. Fast federated learning in the presence of arbitrary device unavailability. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 12052–12064. Curran Associates, Inc., 2021.
- [24] Peter Kairouz, H. Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Keith Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, Rafael G. L. D’Oliveira, Hubert Eichner, Salim El Rouayheb, David Evans, Josh Gardner, Zachary Garrett, Adrià Gascón, Badih Ghazi, Phillip B. Gibbons, Marco Gruteser, Zaid Harchaoui, Chaoyang He, Lie He, Zhouyuan Huo, Ben Hutchinson, Justin Hsu, Martin Jaggi, Tara Javidi, Gauri Joshi, Mikhail Khodak, Jakub Konečný, Aleksandra Korolova, Farinaz Koushanfar, Sanmi Koyejo, Tancrède Lepoint, Yang Liu, Prateek Mittal, Mehryar Mohri, Richard Nock, Ayfer Özgür, Rasmus Pagh, Mariana Raykova, Hang Qi, Daniel Ramage, Ramesh Raskar, Dawn Song, Weikang Song, Sebastian U. Stich, Ziteng Sun, Ananda Theertha Suresh, Florian Tramèr, Praneeth Vepakomma, Jianyu Wang, Li Xiong, Zheng Xu, Qiang Yang, Felix X. Yu, Han Yu, and Sen Zhao. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2):1–210, 2021.
- [25] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank J. Reddi, Sebastian U. Stich, and Ananda Theertha Suresh. SCAFFOLD: stochastic controlled averaging for on-device federated learning. *CoRR*, abs/1910.06378, 2019.
- [26] Anastasia Koloskova, Nicolas Loizou, Sadra Boreiri, Martin Jaggi, and Sebastian U. Stich. A unified theory of decentralized sgd with changing topology and local updates. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*. PMLR, 2020.
- [27] Remi Leblond, Fabian Pedregosa, and Simon Lacoste-Julien. Improved asynchronous parallel optimization analysis for stochastic incremental methods. *Journal of Machine Learning Research*, 19(81):1–68, 2018.
- [28] 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.
- [29] Xiangru Lian, Ce Zhang, Huan Zhang, Cho-Jui Hsieh, Wei Zhang, and Ji Liu. Can decentralized algorithms outperform centralized algorithms? a case study for decentralized parallel stochastic gradient descent. In *Advances in Neural Information Processing Systems 30*, pages 5330–5340. Curran Associates, Inc., 2017.
- [30] Olvi L. Mangasarian and Mikhail V. Solodov. Backpropagation convergence via deterministic nonmonotone perturbed minimization. In J. Cowan, G. Tesauro, and J. Alspector, editors, *Advances in Neural Information Processing Systems*, volume 6. Morgan-Kaufmann, 1994.
- [31] 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. doi: 10.1137/16M1057000.
- [32] Ryan McDonald, Keith Hall, and Gideon Mann. Distributed training strategies for the structured perceptron. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 456–464. Association for Computational Linguistics, 2010.

- [33] Brendan McMahan and Matthew Streeter. Delay-tolerant algorithms for asynchronous distributed online learning. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014.
- [34] H. Brendan McMahan, Eider Moore, Daniel Ramage, and Blaise Agüera y Arcas. Federated learning of deep networks using model averaging. *arXiv preprint arXiv:1602.05629*, 2016.
- [35] Konstantin Mishchenko, Francis Bach, Mathieu Even, and Blake Woodworth. Asynchronous SGD beats minibatch SGD under arbitrary delays. In *Advances in Neural Information Processing Systems*. arXiv preprint arXiv:2206.07638, 2022.
- [36] Angelia Nedić. Distributed gradient methods for convex machine learning problems in networks: Distributed optimization. *IEEE Signal Processing Magazine*, 37(3):92–101, 2020.
- [37] John Nguyen, Kshitiz Malik, Hongyua Zhan, Ashka Yousefpour, Mike Rabbat, Mani Malek, and Dzmitry Huba. Federated learning with buffered asynchronous aggregation. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*. PMLR, 2022.
- [38] Lam Nguyen, Phuong Ha Nguyen, Marten van Dijk, Peter Richtarik, Katya Scheinberg, and Martin Takac. SGD and Hogwild! Convergence without the bounded gradients assumption. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 3750–3758. PMLR, 10–15 Jul 2018.
- [39] Feng Niu, Benjamin Recht, Christopher Re, and Stephen J. Wright. HOGWILD!: A lock-free approach to parallelizing stochastic gradient descent. In *Proceedings of the 24th International Conference on Neural Information Processing Systems*, pages 693–701. Curran Associates Inc., 2011.
- [40] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021.
- [41] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- [42] Benjamin Recht, Christopher Re, Stephen Wright, and Feng Niu. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 693–701. Curran Associates, Inc., 2011.
- [43] Herbert Robbins and Sutton Monro. A Stochastic Approximation Method. *The Annals of Mathematical Statistics*, 22(3):400–407, September 1951.
- [44] Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. Megatron-lm: Training multi-billion parameter language models using model parallelism. *arXiv preprint arXiv:1909.08053*, 2019.
- [45] Suvrit Sra, Adams Wei Yu, Mu Li, and Alex Smola. Adadelayer: Delay adaptive distributed stochastic optimization. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pages 957–965. PMLR, 2016.
- [46] Sebastian Stich, Amirkeivan Mohtashami, and Martin Jaggi. Critical parameters for scalable distributed learning with large batches and asynchronous updates. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, pages 4042–4050. PMLR, 2021.
- [47] Sebastian U. Stich. Local SGD converges fast and communicates little. *ICLR - International Conference on Learning Representations*, art. arXiv:1805.09767, 2019.
- [48] Sebastian U. Stich and Sai Praneeth Karimireddy. The error-feedback framework: SGD with delayed gradients. *Journal of Machine Learning Research*, 21(237):1–36, 2020.
- [49] Thijs Vogels, Sai Praneeth Karimireddy, and Martin Jaggi. PowerSGD: Practical low-rank gradient compression for distributed optimization. In *Advances in Neural Information Processing Systems 32 (NeurIPS)*, pages 1626–1636. Curran Associates, Inc., 2019.
- [50] Meng Wang, Weijie Fu, Xiangnan He, Shijie Hao, and Xindong Wu. A survey on large-scale machine learning. *IEEE Transactions on Knowledge and Data Engineering*, 2020.
- [51] Xuyang Wu, Sindri Magnusson, Hamid Reza Feyzmahdavian, and Mikael Johansson. Delay-adaptive step-sizes for asynchronous learning, 2022.

- [52] Yikai Yan, Chaoyue Niu, Yucheng Ding, Zhenzhe Zheng, Fan Wu, Guihai Chen, Shaojie Tang, and Zhihua Wu. Distributed non-convex optimization with sublinear speedup under intermittent client availability, 2020.
- [53] Haibo Yang, Xin Zhang, Prashant Khanduri, and Jia Liu. Anarchic federated learning, 2021.
- [54] Hossein Zakerinia, Shayan Talaei, Giorgi Nadiradze, and Dan Alistarh. QuAFL: Federated averaging can be both asynchronous and communication-efficient. *arXiv preprint arXiv:2206.10032*, 2022.
- [55] Wei Zhang, Suyog Gupta, Xiangru Lian, and Ji Liu. Staleness-aware async-sgd for distributed deep learning. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-16)*, 2016.
- [56] Shuxin Zheng, Qi Meng, Taifeng Wang, Wei Chen, Nenghai Yu, Zhi-Ming Ma, and Tie-Yan Liu. Asynchronous stochastic gradient descent with delay compensation. In *Proceedings of the 34th International Conference on Machine Learning*. PMLR, 2017.
- [57] Martin Zinkevich, Markus Weimer, Lihong Li, and Alex J Smola. Parallelized stochastic gradient descent. In *Advances in Neural Information Processing Systems*, pages 2595–2603, 2010.

A Proofs

In this section we provide the proofs of all the theoretical results stated in the main paper.

A.1 Proof of Remark 5

First, we prove our key observation given in Remark 5.

Remark 5 (Key Observation). *In Algorithm 1 the average concurrency $\bar{\tau}_C$ is connected to the average delay τ_{avg} as*

$$\tau_{avg} = \frac{T+1}{T+|C_T|-1} \bar{\tau}_C \stackrel{T > |C_T|}{=} \Theta(\bar{\tau}_C). \quad (7)$$

Proof. Define $\{\tau_i^{C_t}\}_{i \in C_t}$ as the set of delays of the gradients that are left in the active worker set before iteration t is performed, i.e. each $\tau_i^{C_t}$ is equal to the difference between the current iteration t and the iteration at which worker i started to compute its current gradient for $t > 0$, and $\tau_i^{C_0} = 1$ for all $i \in C_0$, that is the initial set of active workers. For simplicity we denote

$$\tau_{sum}^{active,t} := \sum_{i \in C_t} \tau_i^{C_t}.$$

We also define $\tau_{sum}^{applied,t}$ as the sum of all delays of gradients applied before iteration t is performed, i.e.

$$\tau_{sum}^{applied,t} := \sum_{j=0}^{t-1} \tau_j.$$

At the zero-th iteration we have that

$$\tau_{sum}^{applied,0} = 0, \quad \tau_{sum}^{active,0} = \tau_C^{(0)}, \quad (16)$$

as no gradients were applied yet.

We claim that

$$\tau_{sum}^{applied,t+1} + \tau_{sum}^{active,t+1} = \tau_{sum}^{applied,t} + \tau_{sum}^{active,t} + \tau_C^{(t+1)}. \quad (17)$$

Indeed, one of the gradients from C_t got applied and its delay moved from $\tau_{sum}^{active,t}$ to $\tau_{sum}^{applied,t+1}$. The newly selected active workers in line 6 of Algorithm 1 have delay zero, as they just started their computations in this step. And all of the current active workers in C_{t+1} (of size $|C_{t+1}| = \tau_C^{(t+1)}$) got an increase by 1 due to increase of the iteration count from t to $t+1$.

Using the initial conditions (16) and (17) we can conclude that

$$\tau_{sum}^{applied,T} + \tau_{sum}^{active,T} = \sum_{t=0}^{T-1} \tau_C^{(t)} = (T+1)\bar{\tau}_C.$$

Note that the left hand side is exactly equal to $(T+|C_T|-1)\tau_{avg}$ from our Definition 4. Thus,

$$\tau_{avg} = \frac{T+1}{T+|C_T|-1} \bar{\tau}_C = \mathcal{O}(\bar{\tau}_C),$$

where the last equality holds if $T > |C_T|$. □

A.2 Useful inequalities

Lemma 12. *For an arbitrary set of n vectors $\{\mathbf{a}_i\}_{i=1}^n$, $\mathbf{a}_i \in \mathbb{R}^d$*

$$\left\| \sum_{i=1}^n \mathbf{a}_i \right\|^2 \leq n \sum_{i=1}^n \|\mathbf{a}_i\|^2. \quad (18)$$

A.3 Proof of Theorems 6, (8) and 8

We first recall both of the theorems

Theorem 6 (Constant stepsizes). *Under Assumptions 1, 3, there exists a constant stepsize $\eta_t \equiv \eta$ such that for Algorithm 1 it holds that $\frac{1}{T+1} \sum_{t=0}^T \|\nabla f(\mathbf{x}^{(t)})\|_2^2 \leq \varepsilon$ after*

$$\mathcal{O}\left(\frac{\sigma^2}{\varepsilon^2} + \frac{\sqrt{\tau_C \tau_{\max}}}{\varepsilon}\right) \quad \text{iterations.} \quad (8)$$

If we additionally assume bounded gradient Assumption 4, then $\frac{1}{\sum_{t=0}^T |\mathcal{A}_t|} \sum_{t=0}^T |\mathcal{A}_t| \|\nabla f(\mathbf{x}^{(t)})\|_2^2 \leq \varepsilon$ after

$$\mathcal{O}\left(\frac{\sigma^2}{\varepsilon^2} + \frac{\tau_C G}{\varepsilon^{3/2}} + \frac{\tau_C}{\varepsilon}\right) \quad \text{iterations.} \quad (9)$$

Theorem 8 (Delay-adaptive stepsizes). *There exist a parameter $\eta \leq \frac{1}{4L}$ such that if we set the stepsizes in Algorithm 1 dependent on the delays as*

$$\eta_t = \begin{cases} \eta & \tau_t \leq 2\tau_C, \\ < \min\{\eta, \frac{1}{4L\tau_t}\} & \tau_t > 2\tau_C, \end{cases} \quad (11)$$

then for Algorithm 1, under Assumptions 1, 3 it holds that $\frac{1}{\sum_{t=0}^T \eta_t} \sum_{t=0}^T \eta_t \|\nabla f(\mathbf{x}^{(t)})\|_2^2 \leq \varepsilon$ after

$$\mathcal{O}\left(\frac{\sigma^2}{\varepsilon^2} + \frac{\tau_C}{\varepsilon}\right) \quad \text{iterations.} \quad (12)$$

We first give a common lemma that will be used in the proofs for both of the theorems.

Lemma 13 (Descent Lemma). *Under Assumptions 1 and 3, if in Algorithm 1 the stepsize $\eta_t < \frac{1}{2L}$ then it holds that*

$$\mathbb{E}_{t+1} f(\mathbf{x}^{(t+1)}) \leq f(\mathbf{x}^{(t)}) - \frac{\eta_t}{2} \|\nabla f(\mathbf{x}^{(t)})\|_2^2 - \frac{\eta_t}{4} \|\nabla f(\mathbf{x}^{(t-\tau_t)})\|_2^2 + L\eta_t^2 \sigma^2 + \frac{\eta_t L^2}{2} \|\mathbf{x}^{(t)} - \mathbf{x}^{(t-\tau_t)}\|_2^2.$$

Proof. Because the function f is L -smooth, we have

$$\begin{aligned} \mathbb{E}_{t+1} f(\mathbf{x}^{(t+1)}) &= \mathbb{E}_{t+1} f\left(\mathbf{x}^{(t)} - \eta_t \nabla F(\mathbf{x}^{(t-\tau_t)}, \xi_{t-\tau_t})\right) \\ &\leq f(\mathbf{x}^{(t)}) - \underbrace{\eta_t \mathbb{E}_{t+1} \langle \nabla f(\mathbf{x}^{(t)}), \nabla F(\mathbf{x}^{(t-\tau_t)}, \xi_{t-\tau_t}) \rangle}_{=: T_1} + \underbrace{\mathbb{E}_{t+1} \frac{L}{2} \eta_t^2 \|\nabla F(\mathbf{x}^{(t-\tau_t)}, \xi_{t-\tau_t})\|_2^2}_{=: T_2} \end{aligned}$$

We first estimate the second term as

$$T_1 = -\eta_t \langle \nabla f(\mathbf{x}^{(t)}), \nabla f(\mathbf{x}^{(t-\tau_t)}) \rangle = -\frac{\eta_t}{2} \|\nabla f(\mathbf{x}^{(t)})\|_2^2 - \frac{\eta_t}{2} \|\nabla f(\mathbf{x}^{(t-\tau_t)})\|_2^2 + \frac{\eta_t}{2} \|\nabla f(\mathbf{x}^{(t)}) - \nabla f(\mathbf{x}^{(t-\tau_t)})\|_2^2$$

For the last term, we add and subtract $\nabla f(\mathbf{x}^{(t-\tau_t)})$, and use that $\mathbb{E}_{t+1} \nabla F(\mathbf{x}^{(t-\tau_t)}, \xi_{t-\tau_t}) - \nabla f(\mathbf{x}^{(t-\tau_t)}) = 0$

$$\begin{aligned} T_2 &= \mathbb{E}_{t+1} \left\| \nabla F(\mathbf{x}^{(t-\tau_t)}, \xi_{t-\tau_t}) - \nabla f(\mathbf{x}^{(t-\tau_t)}) \right\|_2^2 + \|\nabla f(\mathbf{x}^{(t-\tau_t)})\|_2^2 \\ &\stackrel{(2)}{\leq} \sigma^2 + \|\nabla f(\mathbf{x}^{(t-\tau_t)})\|_2^2. \end{aligned}$$

Combining this together and using L -smoothness to estimate $\|\nabla f(\mathbf{x}^{(t)}) - \nabla f(\mathbf{x}^{(t-\tau_t)})\|_2^2$,

$$\begin{aligned} \mathbb{E}_{t+1} f(\mathbf{x}^{(t+1)}) &\leq f(\mathbf{x}^{(t)}) - \eta_t \|\nabla f(\mathbf{x}^{(t)})\|_2^2 - \frac{\eta_t}{2} (1 - L\eta_t) \|\nabla f(\mathbf{x}^{(t-\tau_t)})\|_2^2 + \frac{\eta_t L^2}{2} \|\mathbf{x}^{(t)} - \mathbf{x}^{(t-\tau_t)}\|_2^2 \\ &\quad + L\eta_t^2 \sigma^2. \end{aligned}$$

Applying $\eta < \frac{1}{2L}$ we get statement of the lemma. \square

A.3.1 Proof of Theorem 6, convergence rate (8)

Lemma 14 (Estimation of the residual). *Under Assumptions 1 and 3, the iterates of Algorithm 1 with the constant stepsize $\eta_t \equiv \eta$ with $\eta \leq \frac{1}{2L\sqrt{\tau_{\max}\tau_C}}$ satisfy*

$$\frac{1}{T+1} \sum_{t=0}^T \mathbb{E} \left\| \mathbf{x}^{(t)} - \mathbf{x}^{(t-\tau_t)} \right\|_2^2 \leq \frac{1}{4L^2(T+1)} \sum_{t=0}^T \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t-\tau_t)}) \right\|^2 + \frac{\sigma^2 \eta}{2L}.$$

Proof. We start with unrolling the difference and use that $\mathbb{E} \nabla F(\mathbf{x}^{(j-\tau_j)}, \xi^{(j-\tau_j)}) = \nabla f(\mathbf{x}^{(j-\tau_j)})$.

$$\begin{aligned} \mathbb{E} \left\| \mathbf{x}^{(t)} - \mathbf{x}^{(t-\tau_t)} \right\|_2^2 &= \mathbb{E} \left\| \sum_{j=t-\tau_t}^{t-1} \eta \nabla F(\mathbf{x}^{(j-\tau_j)}, \xi_{j-\tau_j}) \right\|_2^2 \stackrel{(5)}{\leq} \mathbb{E} \left\| \sum_{j=t-\tau_t}^{t-1} \eta \nabla f(\mathbf{x}^{(j-\tau_j)}) \right\|_2^2 + \tau_t \eta^2 \sigma^2 \\ &\stackrel{(18)}{\leq} \tau_t \mathbb{E} \sum_{j=t-\tau_t}^{t-1} \eta^2 \left\| \nabla f(\mathbf{x}^{(j-\tau_j)}) \right\|^2 + \tau_t \eta^2 \sigma^2. \end{aligned}$$

Using that $\eta \leq \frac{1}{2L\sqrt{\tau_{\max}\tau_C}}$,

$$\mathbb{E} \left\| \mathbf{x}^{(t)} - \mathbf{x}^{(t-\tau_t)} \right\|_2^2 \leq \frac{1}{4L^2\tau_C} \sum_{j=t-\tau_t}^{t-1} \mathbb{E} \left\| \nabla f(\mathbf{x}^{(j-\tau_j)}) \right\|^2 + \tau_t \eta^2 \sigma^2.$$

Summing over T ,

$$\begin{aligned} \sum_{t=0}^T \mathbb{E} \left\| \mathbf{x}^{(t)} - \mathbf{x}^{(t-\tau_t)} \right\|_2^2 &\leq \frac{1}{4L^2\tau_C} \sum_{t=0}^T \sum_{j=t-\tau_t}^{t-1} \mathbb{E} \left\| \nabla f(\mathbf{x}^{(j-\tau_j)}) \right\|^2 + \sum_{t=0}^T \tau_t \eta^2 \sigma^2 \\ &\leq \frac{1}{4L^2\tau_C} \sum_{t=0}^T \sum_{j=t-\tau_t}^{t-1} \mathbb{E} \left\| \nabla f(\mathbf{x}^{(j-\tau_j)}) \right\|^2 + (T+1) \tau_{avg} \eta^2 \sigma^2. \end{aligned}$$

We now observe that the number of times each of the gradients $\left\| \nabla f(\mathbf{x}^{(j-\tau_j)}) \right\|^2$ appears in the right hand side is bounded by $\tau_C^{(j)} - 1$ because this many gradients started to be computed before the iteration j and will get applied at some iteration $t > j$. Thus,

$$\sum_{t=0}^T \mathbb{E} \left\| \mathbf{x}^{(t)} - \mathbf{x}^{(t-\tau_t)} \right\|_2^2 \leq \frac{1}{4L^2} \sum_{t=0}^T \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t-\tau_t)}) \right\|^2 + (T+1) \frac{\sigma^2 \eta}{L},$$

where for the last σ term we estimated $\eta \leq \frac{1}{2L\sqrt{\tau_{\max}\tau_C}}$ and used that both $\tau_{avg} \leq 2\tau_C$ and $\tau_{avg} \leq \tau_{\max}$. Dividing the inequality by $T+1$ we get the statement of the lemma. \square

Next, we give the proof of the first part of Theorem 6.

Proof of Theorem 6, convergence rate (8). We start by averaging with T and dividing by η the descent Lemma 13.

$$\begin{aligned} \frac{1}{T+1} \sum_{t=0}^T \left(\frac{1}{2} \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 + \frac{1}{4} \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t-\tau_t)}) \right\|_2^2 \right) &\leq \frac{1}{\eta(T+1)} \left(f(\mathbf{x}^{(0)}) - f^* \right) + L\eta\sigma^2 \\ &\quad + \frac{1}{T+1} \frac{L^2}{2} \sum_{t=0}^T \mathbb{E} \left\| \mathbf{x}^{(t)} - \mathbf{x}^{(t-\tau_t)} \right\|_2^2. \end{aligned}$$

We next apply Lemma 14 to the last term and get

$$\begin{aligned} \frac{1}{T+1} \sum_{t=0}^T \left(\frac{1}{2} \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 + \frac{1}{4} \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t-\tau_t)}) \right\|_2^2 \right) &\leq \frac{1}{\eta(T+1)} \left(f(\mathbf{x}^{(0)}) - f^* \right) + L\eta\sigma^2 \\ &\quad + \frac{1}{8(T+1)} \sum_{t=0}^T \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t-\tau_t)}) \right\|_2^2 + \frac{L\eta\sigma^2}{2}. \end{aligned}$$

And thus,

$$\frac{1}{T+1} \sum_{t=0}^T \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 \leq \frac{2}{\eta(T+1)} \left(f(\mathbf{x}^{(0)}) - f^* \right) + 4L\eta\sigma^2.$$

It is only left to choose a stepsize η . Similar to previous works [48], we chose it as

$$\eta = \min \left\{ \frac{1}{2L\sqrt{\tau_{\max}\tau_C}}; \left(\frac{r_0}{2L\sigma^2(T+1)} \right)^{\frac{1}{2}} \right\} \leq \frac{1}{2L\sqrt{\tau_{\max}\tau_C}},$$

where we defined $r_0 = f(\mathbf{x}^{(0)}) - f^*$. With this choice of stepsize we indeed have that

- If $\frac{1}{2L\sqrt{\tau_{\max}\tau_C}} \leq \left(\frac{r_0}{2L\sigma^2(T+1)} \right)^{\frac{1}{2}}$ then $\eta = \frac{1}{2L\sqrt{\tau_{\max}\tau_C}}$, and

$$\frac{1}{T+1} \sum_{t=0}^T \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 \leq \frac{4Lr_0\sqrt{\tau_{\max}\tau_C}}{T+1} + \left(\frac{r_0}{2L\sigma^2(T+1)} \right)^{\frac{1}{2}} 4L\sigma^2 = \mathcal{O} \left(\frac{\sigma}{\sqrt{T}} + \frac{\sqrt{\tau_{\max}\tau_C}}{T} \right)$$
- Otherwise if $\frac{1}{2L\sqrt{\tau_{\max}\tau_C}} > \left(\frac{r_0}{2L\sigma^2(T+1)} \right)^{\frac{1}{2}}$

$$\frac{1}{T+1} \sum_{t=0}^T \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 \leq 2 \left(\frac{8L\sigma^2 r_0}{(T+1)} \right)^{\frac{1}{2}} = \mathcal{O} \left(\frac{\sigma}{\sqrt{T}} \right)$$

□

A.3.2 Proof of Theorem 8

Lemma 15 (Estimation of the residual). *Under Assumptions 1 and 3, the iterates of Algorithm 1 with the stepsizes η_t chosen as in (11), which we repeat here for readability*

$$\eta_t = \begin{cases} \eta & \tau_t \leq 2\tau_C, \\ < \min\{\eta, \frac{1}{4L\tau_t}\} & \tau_t > 2\tau_C, \end{cases}$$

with $\eta \leq \frac{1}{4L\tau_C}$ satisfy

$$\sum_{t=0}^T \eta_t \left\| \mathbf{x}^{(t)} - \mathbf{x}^{(t-\tau_t)} \right\|_2^2 \leq \frac{1}{16L^2} \sum_{t=0}^T \eta_t \left\| \nabla f(\mathbf{x}^{(t-\tau_t)}) \right\|_2^2 + \frac{\sigma^2}{4L} \sum_{t=0}^T \eta_t^2.$$

Proof.

$$\begin{aligned} \eta_t \left\| \mathbf{x}^{(t)} - \mathbf{x}^{(t-\tau_t)} \right\|_2^2 &= \eta_t \left\| \sum_{j=t-\tau_t}^{t-1} \eta_j \nabla F(\mathbf{x}^{(j-\tau_j)}, \xi_{j-\tau_j}) \right\|_2^2 \stackrel{(2)}{\leq} \eta_t \left\| \sum_{j=t-\tau_t}^{t-1} \eta_j \nabla f(\mathbf{x}^{(j-\tau_j)}) \right\|_2^2 + \eta_t \sum_{j=t-\tau_t}^{t-1} \eta_j^2 \sigma^2 \\ &\stackrel{(18)}{\leq} \eta_t \tau_t \sum_{j=t-\tau_t}^{t-1} \eta_j^2 \left\| \nabla f(\mathbf{x}^{(j-\tau_j)}) \right\|_2^2 + \eta_t \sum_{j=t-\tau_t}^{t-1} \eta_j^2 \sigma^2. \end{aligned}$$

We use that each of the stepsizes $\eta_t \leq \frac{1}{4L \max\{\tau_t, \tau_C\}}$. Thus,

$$\eta_t \left\| \mathbf{x}^{(t)} - \mathbf{x}^{(t-\tau_t)} \right\|_2^2 \leq \frac{1}{4L} \sum_{j=t-\tau_t}^{t-1} \eta_j^2 \left\| \nabla f(\mathbf{x}^{(j-\tau_j)}) \right\|_2^2 + \frac{1}{4L\tau_C} \sum_{j=t-\tau_t}^{t-1} \eta_j^2 \sigma^2.$$

Summing over T , and using that each of the gradients $\left\| \nabla f(\mathbf{x}^{(j-\tau_j)}) \right\|_2^2$ would appear at most $\tau_C^{(j)} - 1$ times (see the discussion in the proof of Lemma 14)

$$\sum_{t=0}^T \eta_t \left\| \mathbf{x}^{(t)} - \mathbf{x}^{(t-\tau_t)} \right\|_2^2 \leq \frac{1}{4L} \sum_{t=0}^T \tau_C \eta_t^2 \left\| \nabla f(\mathbf{x}^{(t-\tau_t)}) \right\|_2^2 + \frac{\sigma^2}{4L} \sum_{t=0}^T \eta_t^2.$$

Using again that $\eta_t \leq \frac{1}{4L \max\{\tau_t, \tau_C\}}$ we get the statement of the lemma. □

Proof of Theorem 8. We start by summing the descent Lemma 13 over the iterations $t = 0, \dots, T$.

$$\sum_{t=0}^T \eta_t \left(\frac{1}{2} \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 + \frac{1}{4} \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t-\tau_t)}) \right\|_2^2 \right) \leq \left(f(\mathbf{x}^{(0)}) - f^* \right) + L\sigma^2 \sum_{t=0}^T \eta_t^2 + \frac{L^2}{2} \sum_{t=0}^T \eta_t \left\| \mathbf{x}^{(t)} - \mathbf{x}^{(t-\tau_t)} \right\|_2^2.$$

Next, we substitute Lemma 15 into the last term,

$$\begin{aligned} \sum_{t=0}^T \eta_t \left(\frac{1}{2} \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 + \frac{1}{4} \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t-\tau_t)}) \right\|_2^2 \right) &\leq \left(f(\mathbf{x}^{(0)}) - f^* \right) + L\sigma^2 \sum_{t=0}^T \eta_t^2 \\ &\quad + \frac{1}{32} \sum_{t=0}^T \eta_t \left\| \nabla f(\mathbf{x}^{(t-\tau_t)}) \right\|_2^2 + \frac{\sigma^2 L}{8} \sum_{t=0}^T \eta_t^2. \end{aligned}$$

Rearranging we thus get

$$\sum_{t=0}^T \eta_t \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 \leq 2 \left(f(\mathbf{x}^{(0)}) - f^* \right) + 4L\sigma^2 \sum_{t=0}^T \eta_t^2.$$

We note that due to our choice of stepsizes (11), $\eta_t \leq \eta$, it also holds that $\sum_{t=0}^T \eta_t \geq \sum_{t:\tau_t \leq 2\tau_C} \eta \geq \frac{T+1}{2} \eta$ since there are at least half of the iterations with the delay smaller than two times the average.

Using this, we estimate

$$\frac{1}{\sum_{t=0}^T \eta_t} \sum_{t=0}^T \eta_t \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 \leq \frac{4}{(T+1)\eta} \left(f(\mathbf{x}^{(0)}) - f^* \right) + 8L\sigma^2 \eta^2.$$

It remains to tune the stepsize η , i.e. to pick is such as to minimize the right hand side of this expression. See Lemma 17 in [26]. \square

A.4 Proof of Theorem 6, convergence rate (9)

To prove the last claim of Theorem 6 we take another approach and follow the perturbed iterate analysis [31].

We introduce a virtual sequence $\tilde{\mathbf{x}}^t$ defined as

$$\tilde{\mathbf{x}}^{(0)} = \mathbf{x}^{(0)}, \quad \tilde{\mathbf{x}}^{(t+1)} = \tilde{\mathbf{x}}^{(t)} - \eta \sum_{i \in \mathcal{A}_t} \nabla F(\mathbf{x}^{(t)}, \xi_t),$$

where we define $\mathcal{A}_0 := \mathcal{C}_0$, and $\hat{\tau}_t^i$ is the delay with which the corresponding gradient will be computed. That is, if we denote $j = t + \hat{\tau}_t^i$, then it will hold that $j - \tau_j = t$. This defines a virtual sequence and we do not have access to it during the execution of Algorithm 1.

Lemma 16 (Descent lemma). *Under Assumptions 1 and 3, if in Algorithm 1 the stepsize $\eta_t < \frac{1}{2L\tau_C}$ then it holds that*

$$\mathbb{E}_{t+1} f(\tilde{\mathbf{x}}^{(t+1)}) \leq f(\tilde{\mathbf{x}}^{(t)}) - \frac{\eta}{4} |\mathcal{A}_t| \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 + \frac{\eta}{2} |\mathcal{A}_t| L^2 \left\| \mathbf{x}^{(t)} - \tilde{\mathbf{x}}^{(t)} \right\|_2^2 + \frac{L\eta^2 \sigma^2 |\mathcal{A}_t|}{2}.$$

Proof. Because function f is L -smooth, we have

$$\begin{aligned} \mathbb{E}_{t+1} f(\tilde{\mathbf{x}}^{(t+1)}) &= \mathbb{E}_{t+1} f \left(\tilde{\mathbf{x}}^{(t)} - \eta \sum_{i \in \mathcal{A}_t} \nabla F(\mathbf{x}^{(t)}, \xi^{(t)}) \right) \\ &\leq f(\tilde{\mathbf{x}}^{(t)}) - \eta |\mathcal{A}_t| \underbrace{\langle \nabla f(\tilde{\mathbf{x}}^{(t)}), \nabla f(\mathbf{x}^{(t)}) \rangle}_{=: T_1} + \mathbb{E}_{t+1} \frac{L}{2} \eta^2 \left\| \sum_{i \in \mathcal{A}_t} \nabla F(\mathbf{x}^{(t)}, \xi_t) \right\|_2^2. \end{aligned}$$

$\underbrace{\hspace{10em}}_{=: T_2}$

We estimate the second term as

$$\begin{aligned} T_2 &= - \langle \nabla f(\mathbf{x}^{(t)}), \nabla f(\tilde{\mathbf{x}}^{(t)}) \rangle = - \frac{1}{2} \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 - \frac{1}{2} \left\| \nabla f(\tilde{\mathbf{x}}^{(t)}) \right\|_2^2 + \frac{1}{2} \left\| \nabla f(\mathbf{x}^{(t)}) - \nabla f(\tilde{\mathbf{x}}^{(t)}) \right\|_2^2 \\ &\leq - \frac{1}{2} \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 + \frac{1}{2} \left\| \nabla f(\mathbf{x}^{(t)}) - \nabla f(\tilde{\mathbf{x}}^{(t)}) \right\|_2^2. \end{aligned}$$

For the last term, using the notation $\pm a = a - a = 0 \quad \forall a$,

$$\begin{aligned} T_2 &= \mathbb{E}_{t+1} \left\| \sum_{i \in \mathcal{A}_t} \nabla F(\mathbf{x}^{(t)}, \xi_i) \pm |\mathcal{A}_t| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 \\ &\stackrel{(2)}{\leq} |\mathcal{A}_t| \sigma^2 + |\mathcal{A}_t|^2 \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2. \end{aligned}$$

Combining this together, using L -smoothness to estimate $\left\| \nabla f(\mathbf{x}^{(t)}) - \nabla f(\tilde{\mathbf{x}}^{(t)}) \right\|_2^2$ we get

$$\mathbb{E}_{t+1} f(\tilde{\mathbf{x}}^{(t+1)}) \leq f(\tilde{\mathbf{x}}^{(t)}) - \left(\frac{\eta}{2} |\mathcal{A}_t| - \frac{\eta^2 L |\mathcal{A}_t|^2}{2} \right) \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 + \frac{\eta}{2} |\mathcal{A}_t| L^2 \left\| \mathbf{x}^{(t)} - \tilde{\mathbf{x}}^{(t)} \right\|_2^2 + \frac{L \eta^2 \sigma^2 |\mathcal{A}_t|}{2}.$$

Using that $\eta \leq \frac{1}{2L\tau_C} \leq \frac{1}{2L|\mathcal{A}_t|}$ we get statement of the Lemma. \square

Lemma 17 (Estimation of the residual). *Under Assumptions 1, 3, iterated of Algorithm 1 with the constant stepsize $\eta_t \equiv \eta$ with $\eta \leq \frac{1}{2L\tau_C}$ satisfy*

$$\mathbb{E} \left\| \mathbf{x}^{(t)} - \tilde{\mathbf{x}}^{(t)} \right\|_2^2 \leq \tau_C^2 \eta^2 G^2 + \eta^2 \tau_C \sigma^2.$$

Proof.

$$\begin{aligned} \mathbb{E} \left\| \mathbf{x}^{(t)} - \tilde{\mathbf{x}}^{(t)} \right\|_2^2 &= \mathbb{E} \left\| \sum_{j \in \mathcal{C}_t} \eta \nabla F(\mathbf{x}^{(j)}, \xi_j) \right\|_2^2 \stackrel{(2)}{\leq} \mathbb{E} \left\| \sum_{j \in \mathcal{C}_t} \eta \nabla f(\mathbf{x}^{(j)}) \right\|_2^2 + \eta^2 \tau_C^{(t)} \sigma^2 \\ &\stackrel{(18)}{\leq} \tau_C^{(t)} \sum_{j \in \mathcal{C}_t} \eta^2 \mathbb{E} \left\| \nabla f(\mathbf{x}^{(j)}) \right\|_2^2 + \eta^2 \tau_C^{(t)} \sigma^2 \\ &\stackrel{(5)}{\leq} (\tau_C^{(t)})^2 \eta^2 G^2 + \eta^2 \tau_C^{(t)} \sigma^2. \end{aligned} \quad \square$$

We are now ready to prove the second claim of Theorem 6.

Proof of Theorem 6, convergence rate (9). We start by summing over $t = 0, \dots, T$ the descent Lemma 16. We also divide it by η ,

$$\sum_{t=0}^T \frac{1}{4} |\mathcal{A}_t| \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 \leq \frac{1}{\eta} \left(f(\mathbf{x}^{(0)}) - f^* \right) + \frac{L\eta\sigma^2}{2} \sum_{t=0}^T |\mathcal{A}_t| + \frac{L^2}{2} \sum_{t=0}^T |\mathcal{A}_t| \left\| \mathbf{x}^{(t)} - \tilde{\mathbf{x}}^{(t)} \right\|_2^2.$$

We further use Lemma 17 for the last term

$$\sum_{t=0}^T \frac{1}{4} |\mathcal{A}_t| \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 \leq \frac{1}{\eta} \left(f(\mathbf{x}^{(0)}) - f^* \right) + \frac{L\eta\sigma^2}{2} \sum_{t=0}^T |\mathcal{A}_t| + \frac{L^2}{2} (\tau_C^2 \eta^2 G^2 + \eta^2 \tau_C \sigma^2) \sum_{t=0}^T |\mathcal{A}_t|.$$

We further use that $\eta \leq \frac{1}{2L\tau_C}$ for the last σ term and divide the full inequality by $\frac{1}{4} \mathcal{W}_T$, where we defined $\mathcal{W}_T = \sum_{t=0}^T |\mathcal{A}_t|$

$$\frac{1}{\mathcal{W}_T} \sum_{t=0}^T |\mathcal{A}_t| \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 \leq \frac{4}{\eta \mathcal{W}_T} \left(f(\mathbf{x}^{(0)}) - f^* \right) + 4L\eta\sigma^2 + 2L^2 \tau_C^2 \eta^2 G^2.$$

Note that because at every step t only one of the gradients is getting applied, $T \leq \sum_{t=0}^T |\mathcal{A}_t| \leq T + \tau_C \leq 2T$ for $T \geq \tau_C$.

It is left to tune the stepsize using Lemma 17 in [26] to get the final convergence rate. \square

A.5 Proof of the Theorem 11

We first re-state the theorem

Theorem 11 (constant stepsizes). *Under Assumptions 1, 2, 3, 5 there exists a constant stepsize $\eta_t \equiv \eta$ such that for Algorithm 2 it holds that $\frac{1}{T+1} \sum_{t=0}^T \|\nabla f(\mathbf{x}^{(t)})\|_2^2 \leq \varepsilon$ after*

$$\mathcal{O}\left(\frac{\sigma^2}{\varepsilon^2} + \frac{\zeta^2}{\varepsilon^2} + \frac{\sqrt{\tau_{avg} \frac{1}{n} \sum_{i=1}^n \zeta_i^2 \tau_{avg}^i}}{\varepsilon^{\frac{3}{2}}} + \frac{\sqrt{\tau_{avg} \tau_{\max}}}{\varepsilon}\right) \quad \text{iterations,} \quad (14)$$

Under Assumptions 1, 2, 3 and additional bounded gradient Assumption 4, it holds that $\frac{1}{T+1} \sum_{t=0}^T \|\nabla f(\mathbf{x}^{(t)})\|_2^2 \leq \varepsilon$ after

$$\mathcal{O}\left(\frac{\sigma^2}{\varepsilon^2} + \frac{\zeta^2}{\varepsilon^2} + \frac{\tau_{avg} G}{\varepsilon^{\frac{3}{2}}} + \frac{\tau_{avg}}{\varepsilon}\right) \quad \text{iterations.} \quad (15)$$

We utilize again the perturbed iterate technique [31]. We introduce a virtual sequence $\tilde{\mathbf{x}}^{(t)}$ as

$$\tilde{\mathbf{x}}^{(0)} = \mathbf{x}^{(0)} \quad \tilde{\mathbf{x}}^{(t+1)} = \tilde{\mathbf{x}}^{(t)} - \eta \nabla F_{k_t}(\mathbf{x}^{(t)}, \xi_t),$$

where we define $\hat{\tau}_t$ as the delay with which the corresponding gradient will be computed. If we denote $j = t + \hat{\tau}_t$, then it holds that $j - \tau_j = t$.

Lemma 18 (Descent Lemma). *Under Assumptions 1, 2, 3, for Algorithm 2 with the stepsize $\eta_t \leq \frac{1}{4L}$ it holds that*

$$\mathbb{E}_{t+1} f(\tilde{\mathbf{x}}^{(t+1)}) \leq f(\tilde{\mathbf{x}}^{(t)}) - \frac{\eta}{4} \|\nabla f(\mathbf{x}^{(t)})\|_2^2 + \frac{L\eta^2\sigma^2}{2} + L\eta^2\zeta^2 + \frac{\eta L^2}{2} \|\mathbf{x}^{(t)} - \tilde{\mathbf{x}}^{(t)}\|_2^2. \quad (19)$$

Proof. Because the function f is L -smooth, we have

$$\begin{aligned} \mathbb{E}_{t+1} f(\tilde{\mathbf{x}}^{(t+1)}) &= \mathbb{E}_{t+1} f\left(\tilde{\mathbf{x}}^{(t)} - \eta \nabla F_{k_t}(\mathbf{x}^{(t)}, \xi_t)\right) \\ &\leq f(\tilde{\mathbf{x}}^{(t)}) - \underbrace{\eta \langle \nabla f(\tilde{\mathbf{x}}^{(t)}), \nabla f(\mathbf{x}^{(t)}) \rangle}_{=: T_1} + \mathbb{E}_{t+1} \underbrace{\frac{L}{2} \eta^2 \|\nabla F_{k_t}(\mathbf{x}^{(t)}, \xi_t)\|_2^2}_{=: T_2}, \end{aligned}$$

where expectation is taken over both the stochastic noise ξ and sampled index j_t . We estimate terms T_1 and T_2 separately

$$\begin{aligned} T_1 &= -\frac{\eta}{2} \|\nabla f(\mathbf{x}^{(t)})\|_2^2 - \frac{\eta}{2} \|\nabla f(\tilde{\mathbf{x}}^{(t)})\|_2^2 + \frac{\eta}{2} \|\nabla f(\mathbf{x}^{(t)}) - \nabla f(\tilde{\mathbf{x}}^{(t)})\|_2^2 \\ &\leq -\frac{\eta}{2} \|\nabla f(\mathbf{x}^{(t)})\|_2^2 + \frac{\eta}{2} \|\nabla f(\mathbf{x}^{(t)}) - \nabla f(\tilde{\mathbf{x}}^{(t)})\|_2^2. \end{aligned}$$

For the last term, using the notation $\pm a = a - a = 0 \quad \forall a$,

$$\begin{aligned} T_2 &= \mathbb{E}_{t+1} \left\| \nabla F_{k_t}(\mathbf{x}^{(t)}, \xi_t) \pm \nabla f_{j_t}(\mathbf{x}^{(t)}) \pm \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 \\ &\stackrel{(2)}{\leq} \sigma^2 + 2\mathbb{E}_{k_t} \left\| \nabla f_{k_t}(\mathbf{x}^{(t)}) - \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 + 2 \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 \\ &\stackrel{(3)}{\leq} \sigma^2 + 2\zeta^2 + 2 \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2. \end{aligned}$$

Combining this together and using L -smoothness to estimate $\|\nabla f(\mathbf{x}^{(t)}) - \nabla f(\tilde{\mathbf{x}}^{(t)})\|_2^2$ we get

$$\mathbb{E}_{t+1} f(\tilde{\mathbf{x}}^{(t+1)}) \leq f(\tilde{\mathbf{x}}^{(t)}) - \left(\frac{\eta}{2} - L\eta^2\right) \|\nabla f(\mathbf{x}^{(t)})\|_2^2 + \frac{\eta}{2} L^2 \|\mathbf{x}^{(t)} - \tilde{\mathbf{x}}^{(t)}\|_2^2 + \frac{L\eta^2\sigma^2}{2} + L\eta^2\zeta^2.$$

Applying $\eta \leq \frac{1}{4L}$ we get statement of the lemma. \square

A.5.1 Proof of Theorem 11, convergence rate (14)

Lemma 19 (Estimation of the distance $\|\mathbf{x}^{(t)} - \tilde{\mathbf{x}}^{(t)}\|_2^2$). *Under Assumptions 1, 2, 3, for Algorithm 2 with the stepsize $\eta_t \leq \frac{1}{4L\sqrt{\tau_C\tau_{\max}}}$ it holds that*

$$\frac{1}{T+1} \sum_{t=0}^T \mathbb{E} \left\| \mathbf{x}^{(t)} - \tilde{\mathbf{x}}^{(t)} \right\|_2^2 \leq \frac{\eta\sigma^2}{4L} + \frac{2\eta^2\tau_C}{T+1} \frac{1}{n} \sum_{j=1}^n \zeta_j^2 \bar{\tau}_j + \frac{1}{8L^2(T+1)} \sum_{t=0}^T \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2.$$

Proof.

$$\begin{aligned} \mathbb{E} \left\| \mathbf{x}^{(t)} - \tilde{\mathbf{x}}^{(t)} \right\|_2^2 &= \mathbb{E} \eta^2 \left\| \sum_{i \in \mathcal{C}_t} \nabla F_{j_i}(\mathbf{x}^{(i)}, \xi_i) \right\|_2^2 \stackrel{(2)}{\leq} \eta^2 \tau_C \sigma^2 + \eta^2 \mathbb{E} \left\| \sum_{i \in \mathcal{C}_t} \nabla f_{j_i}(\mathbf{x}^{(i)}) \right\|_2^2 \\ &\stackrel{(18)}{\leq} \eta^2 \tau_C \sigma^2 + 2\eta^2 \mathbb{E} \left\| \sum_{i \in \mathcal{C}_t} \nabla f_{j_i}(\mathbf{x}^{(i)}) - \nabla f(\mathbf{x}^{(i)}) \right\|_2^2 + 2\eta^2 \mathbb{E} \left\| \sum_{i \in \mathcal{C}_t} \nabla f(\mathbf{x}^{(i)}) \right\|_2^2 \\ &\stackrel{(18)}{\leq} \eta^2 \tau_C \sigma^2 + 2\eta^2 \tau_C^{(t)} \mathbb{E} \sum_{i \in \mathcal{C}_t} \zeta_{j_i}^2 + 2\eta^2 \tau_C \mathbb{E} \sum_{i \in \mathcal{C}_t} \left\| \nabla f(\mathbf{x}^{(i)}) \right\|_2^2. \end{aligned}$$

Averaging over T , we get

$$\frac{1}{T+1} \sum_{t=0}^T \mathbb{E} \left\| \mathbf{x}^{(t)} - \tilde{\mathbf{x}}^{(t)} \right\|_2^2 \leq \eta^2 \tau_C \sigma^2 + 2\eta^2 \tau_C \frac{1}{T+1} \sum_{t=0}^T \mathbb{E} \sum_{i \in \mathcal{C}_t} \zeta_{j_i}^2 + 2\eta^2 \tau_C \frac{1}{T+1} \sum_{t=0}^T \mathbb{E} \sum_{i \in \mathcal{C}_t} \left\| \nabla f(\mathbf{x}^{(i)}) \right\|_2^2.$$

We note that in the second term each of ζ_j appears exactly τ_j^{sum} times, where τ_j^{sum} is the sum of all the delays that happened on the node j . In the last term, we estimate the number of appearance of each of $\left\| \nabla f(\mathbf{x}^{(i)}) \right\|_2^2$ by τ_{\max} , thus

$$\frac{1}{T+1} \sum_{t=0}^T \mathbb{E} \left\| \mathbf{x}^{(t)} - \tilde{\mathbf{x}}^{(t)} \right\|_2^2 \leq \eta^2 \tau_C \sigma^2 + 2\eta^2 \tau_C \mathbb{E} \frac{1}{T+1} \sum_{j=1}^n \zeta_j^2 \tau_j^{sum} + 2\eta^2 \tau_C \tau_{\max} \frac{1}{T+1} \sum_{t=0}^T \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2,$$

we further use that number of times T_j that every node j got sampled are equal in expectation because of uniform sampling in line 6 of Algorithm 2. Thus,

$$\frac{1}{T+1} \sum_{t=0}^T \mathbb{E} \left\| \mathbf{x}^{(t)} - \tilde{\mathbf{x}}^{(t)} \right\|_2^2 \leq \eta^2 \tau_C \sigma^2 + 2\eta^2 \tau_C \frac{1}{T+1} \frac{1}{n} \sum_{j=1}^n \zeta_j^2 \bar{\tau}_j + 2\eta^2 \tau_C \tau_{\max} \frac{1}{T+1} \sum_{t=0}^T \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2.$$

Using that $\eta \leq \frac{1}{4L\sqrt{\tau_C\tau_{\max}}}$ we get the statement of the lemma. \square

Proof of Theorem 11, (14). First, averaging the descent Lemma 16,

$$\frac{1}{T+1} \sum_{t=0}^T \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 \leq \frac{4}{\eta(T+1)} (f(\mathbf{x}^0) - f(\mathbf{x}^T)) + 2L\eta\sigma^2 + 4L\eta\zeta^2 + \frac{2L^2}{T+1} \sum_{t=0}^T \mathbb{E} \left\| \mathbf{x}^{(t)} - \tilde{\mathbf{x}}^{(t)} \right\|_2^2.$$

Now plugging in the result of Lemma 19, we get

$$\begin{aligned} \frac{1}{T+1} \sum_{t=0}^T \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 &\leq \frac{4}{\eta(T+1)} (f(\mathbf{x}^0) - f(\mathbf{x}^T)) + 2L\eta\sigma^2 + 4L\eta\zeta^2 + \frac{L\eta\sigma^2}{2} \\ &\quad + \frac{4L^2\eta^2\tau_C}{T+1} \frac{1}{n} \sum_{j=1}^n \zeta_j^2 \bar{\tau}_j + \frac{1}{4(T+1)} \sum_{t=0}^T \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2. \end{aligned}$$

Rearranging terms we thus get

$$\frac{1}{2(T+1)} \sum_{t=0}^T \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 \leq \frac{4}{\eta(T+1)} (f(\mathbf{x}^0) - f(\mathbf{x}^T)) + 3L\eta\sigma^2 + 4L\eta\zeta^2 + \frac{4L^2\eta^2\tau_C}{T+1} \frac{1}{n} \sum_{j=1}^n \zeta_j^2 \bar{\tau}_j$$

It is only left to tune the stepsize using Lemma 17 in [26]. \square

A.5.2 Proof of Theorem 11, convergence rate (15).

Lemma 20 (Estimation of the distance $\|\mathbf{x}^{(t)} - \tilde{\mathbf{x}}^{(t)}\|_2^2$). *Under Assumptions 1, 2, 3, 4 for Algorithm 2 with the stepsize $\eta_t \equiv \eta \leq \frac{1}{4L\tau_C}$ it holds that*

$$\frac{1}{T+1} \sum_{t=0}^T \mathbb{E} \left\| \mathbf{x}^{(t)} - \tilde{\mathbf{x}}^{(t)} \right\|_2^2 \leq \frac{\eta\sigma^2}{4L} + \eta^2\tau_C^2 G^2.$$

Proof. We start our proof similar way as before

$$\begin{aligned} \mathbb{E} \left\| \mathbf{x}^{(t)} - \tilde{\mathbf{x}}^{(t)} \right\|_2^2 &= \mathbb{E} \eta^2 \left\| \sum_{i \in \mathcal{C}_t} \nabla F_{j_i}(\mathbf{x}^{(i)}, \xi_i) \right\|_2^2 \stackrel{(2)}{\leq} \eta^2 \tau_C \sigma^2 + \eta^2 \mathbb{E} \left\| \sum_{i \in \mathcal{C}_t} \nabla f_{j_i}(\mathbf{x}^{(i)}) \right\|_2^2 \\ &\stackrel{(18)}{\leq} \eta^2 \tau_C \sigma^2 + \eta^2 \tau_C \sum_{i \in \mathcal{C}_t} \mathbb{E} \left\| \nabla f_{j_i}(\mathbf{x}^{(i)}) \right\|_2^2 \\ &\stackrel{(5)}{\leq} \eta^2 \tau_C \sigma^2 + \eta^2 \tau_C^2 G^2 \\ &\leq \frac{\eta\sigma^2}{4L} + \eta^2 \tau_C^2 G^2 \end{aligned}$$

where on the last line we used that stepsize $\eta \leq \frac{1}{4L\tau_C}$. \square

Proof of the Theorem 11, (15). We start by averaging the descent Lemma 16,

$$\frac{1}{T+1} \sum_{t=0}^T \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 \leq \frac{4}{\eta(T+1)} (f(\mathbf{x}^0) - f(\mathbf{x}^T)) + 2L\eta\sigma^2 + 4L\eta\zeta^2 + \frac{2L^2}{T+1} \sum_{t=0}^T \mathbb{E} \left\| \mathbf{x}^{(t)} - \tilde{\mathbf{x}}^{(t)} \right\|_2^2.$$

We now plug in the results of Lemma 20 and get

$$\frac{1}{T+1} \sum_{t=0}^T \mathbb{E} \left\| \nabla f(\mathbf{x}^{(t)}) \right\|_2^2 \leq \frac{4}{\eta(T+1)} (f(\mathbf{x}^0) - f(\mathbf{x}^T)) + 3L\eta\sigma^2 + 4L\eta\zeta^2 + 2L^2\eta^2\tau_C^2 G^2.$$

It is only left to tune the stepsize using Lemma 17 in [26]. \square

B Estimating Speedup over Synchronous SGD

Assume that we have n workers with identical functions $f_i \equiv f_j \equiv f \forall i, j$ (homogeneous case). Assume that the objective function has a sum-structure $f(\mathbf{x}) = \frac{1}{m} \sum_{j=1}^m F(\mathbf{x}, \xi_j)$. This setting is common in machine learning where each $F(\mathbf{x}, \xi_j)$ represent a loss function of a model \mathbf{x} on a datapoint $\xi_j \in \mathcal{D}$, $|\mathcal{D}| = m$. Assume that to compute stochastic gradient $\nabla F(\mathbf{x}, \xi_j)$ each worker needs a constant time Δ_j . W.l.o.g. we assume that Δ_i are ordered as $\Delta_1 \leq \Delta_2 \leq \dots \leq \Delta_m$.

Lemma 21. *In expectation, the asynchronous Algorithm 2 needs*

$$\bar{\Delta} = \frac{1}{m} \sum_{i=1}^m \Delta_i$$

time to compute n gradients, while mini-batch SGD with batch size n needs

$$\tilde{\Delta} = \sum_{i=1}^m \alpha_i \Delta_i$$

time to compute a batch of n gradients, where $\alpha_i = \frac{i^n - (i-1)^n}{m^n}$. It is also always holds that $\bar{\Delta} \leq \tilde{\Delta}$.

With this lemma we can precisely estimate how much faster the asynchronous algorithm is compared to the classic synchronous mini-batch one. Note that α_i are increasing with i with a rate of $\mathcal{O}(i^n)$, thus in mini-batch SGD, the large delays get a much higher weight than the small delays, especially when the batch size n is large.

For example, consider 1000 clients, 900 of which compute their update every 10s, while 100 of them computes their update every 60s. Then the expected time for n gradients of the asynchronous algorithm will be 15s, while synchronous mini-batch SGD (with $n = 10$) will take a significantly longer time of 42.5s for the same number of gradients.

B.1 Proof of Lemma 21

In this section we prove Lemma 21.

Proof. We start by proving the first claim.

Time of Asynchronous Algorithm 2. Assume the concurrency is $n = 1$. Then, Algorithm 2 is synchronous, and as we sample every client with equal probability (line 6 of Algorithm 2), the expected time to compute one gradient is equal to $\frac{1}{m} \sum_{i=1}^m \Delta_i$.

To calculate the estimated time with concurrency $n > 1$ we can view the Algorithm 2 as having n independent copies of the previous process run in parallel. Thus, in the same time $\frac{1}{m} \sum_{i=1}^m \Delta_i$ in expectation Algorithm 2 will compute n gradients.

Time of Mini-batch SGD. The expected time of mini-batch SGD of size n is equal to $\mathbb{E} \max\{\Delta_{i_1}, \dots, \Delta_{i_n}\}$ with each $i_j \sim \text{Uniform}[1, m]$. Denote a random variable $X = \max\{\Delta_{i_1}, \dots, \Delta_{i_n}\}$ that takes values within $\Delta_1, \dots, \Delta_m$. Since i_j are independent from each other,

$$\Pr [X \leq \Delta_k] = \prod_{j=1}^n \Pr [\Delta_{i_j} \leq \Delta_k] = \prod_{j=1}^n \Pr [X \leq \Delta_k] = \prod_{j=1}^n \Pr [i_j \leq k] = \left(\frac{k}{m}\right)^n.$$

Thus,

$$\Pr [X = \Delta_k] = \frac{k^n - (k-1)^n}{m^n}.$$

And therefore,

$$\mathbb{E} X = \sum_{k=1}^m \Pr [X = \Delta_k] \Delta_k. \quad \square$$

C Experiments

In this set of experiments we aim to illustrate the dependence on the maximum delay τ_{\max} in Theorem 6, as depicted in Equation (8). For this, we set the stochastic noise σ to zero. In this case Theorem 6, Equation (8) predicts that to reach an ε accuracy, Algorithm 1 needs $T = \mathcal{O}\left(\frac{\sqrt{\tau_{\max} \tau_C}}{\varepsilon}\right)$ iterations. In our experiments we fix $\tau_C = 2$, $\varepsilon = 10^{-14}$. Since $\tau_C = 2$, we have two workers. We vary the relative speed of the second worker, and thus affecting the maximum delay: if the second worker is x times slower than the first worker, then the maximum delay $\tau_{\max} = x$. We measure the time T to reach the accuracy ε . Since all the other parameters are constant, it holds that $T = C_1 \sqrt{\tau_{\max}}$.

We perform experiments on two different functions:

- (i) quadratic function $f(\mathbf{x}) = \frac{1}{2} \|A\mathbf{x} - \mathbf{b}\|_2^2$, $\mathbf{x}, \mathbf{b} \in \mathbb{R}^{10}$, $b_i \sim \mathcal{N}(0, 1)$, $i \in [1, 10]$, $A \in \mathbb{R}^{10 \times 10}$ is a random matrix with $\lambda_{\max}(A) = 2$, $\lambda_{\min}(A) = 1$ and the rest of eigenvalues are equally spaced in between.
- (ii) logistic regression function $f(\mathbf{x}) = \frac{1}{m} \sum_{j=1}^m \log(1 + \exp(-b_j \mathbf{a}_j^\top \mathbf{x}))$, where each b_j is sampled uniformly at random from the set $\{-1, 1\}$, and $\mathbf{a}_j \sim \mathcal{N}(0, 1)^{20}$, $\mathbf{x} \in \mathbb{R}^{20}$, $m = 100$.

We estimate the error as the average over the last 30 iterations $\hat{\varepsilon}_T = \frac{1}{30} \sum_{i=0}^{29} \|\nabla f(\mathbf{x}_{T-i})\|_2$. We tune the stepsize η for every experiment *separately* over the logarithmic grid between 10^{-5} and 10^2 ensuring that the optimal stepsize value is not on the edge of the grid.

Figure 2 shows the resulting dependence of T on τ_{\max} for the quadratic function (i), and Figure 3 for the logistic regression function (ii). In both cases we see that T has linear dependence on $\sqrt{\tau_{\max}}$ confirming our theory.




Figure 2: Verification of $\sqrt{\tau_{\max}}$ dependence on random quadratic function (i). We see that T has linear dependence on $\sqrt{\tau_{\max}}$




Figure 3: Verification of $\sqrt{\tau_{\max}}$ dependence on random logistic regression function (ii). We see that T has linear dependence on $\sqrt{\tau_{\max}}$