
AsGrad: A Sharp Unified Analysis of Asynchronous-SGD Algorithms

Rustem Islamov
University of Basel
Switzerland

Mher Safaryan
ISTA, Austria

Dan Alistarh
ISTA, Austria

Abstract

We analyze asynchronous-type algorithms for distributed SGD in the heterogeneous setting, where each worker has its own computation and communication speeds, as well as data distribution. In these algorithms, workers compute possibly stale and stochastic gradients associated with their local data at some iteration back in history and then return those gradients to the server without synchronizing with other workers. We present a unified convergence theory for non-convex smooth functions in the heterogeneous regime. The proposed analysis provides convergence for pure asynchronous SGD and its various modifications. Moreover, our theory explains what affects the convergence rate and what can be done to improve the performance of asynchronous algorithms. In particular, we introduce a novel asynchronous method based on worker shuffling. As a by-product of our analysis, we also demonstrate convergence guarantees for gradient-type algorithms such as SGD with random reshuffling and shuffle-once, and mini-batch SGD. The derived rates match the best-known results for those algorithms, highlighting the tightness of our approach. Finally, our numerical evaluations support theoretical findings and show the good practical performance of our method.

(SGD) method (Robbins and Monro, 1951) and its various modifications, such as Adam (Kingma and Ba, 2015) and AdaGrad (Duchi et al., 2011) are at the core of machine learning training, due to their easy implementation together with strong practical performance. However, in recent times, both the size of state-of-the-art models and the amount of data required for training have increased significantly. Due to this growth, optimization algorithms have also been adapted to the need for efficient training of large models, and thus distributed and parallel variants of SGD have started playing a crucial role in modern machine learning (Bekkerman et al., 2011). In the distributed regime, training is performed using many computational nodes (e.g., CPUs or GPUs on a cluster) working in parallel and orchestrated by a server. Every worker computes gradients based on available data, and then a server aggregates those gradients to perform one step of an algorithm. Distributed SGD-based algorithms are also adopted to Federated Learning applications (Konečný et al., 2016; Kairouz et al., 2021) where local data is kept private and is not seen by other workers.

Nevertheless, distributed variants of SGD suffer from many practical challenges. For example, approaches such as communication compression (Alistarh et al., 2018; Khirirat et al., 2018; Mishchenko et al., 2019; Gao et al., 2024), performing several local steps before communication (Mishchenko et al., 2022b; Gorbunov et al., 2021; Koloskova et al., 2020), decentralized communication (Koloskova et al., 2019; Kovalev et al., 2021), or their combinations (Condat et al., 2023; Zakerinia et al., 2022) are designed to improve the efficiency of distributed training.

Across all these approaches, workers are synchronized, i.e., the server must wait for the slowest worker before proceeding to the next algorithm iteration. This may drastically slow down the performance of SGD if workers have significantly different computational power. *Asynchronous* communication, which breaks this lock-step behavior, is preferable in practice, since it enables

1 INTRODUCTION

Modern machine learning relies heavily on effective optimization algorithms. The stochastic gradient descent

Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. Copyright 2024 by the author(s).

a more efficient use of resources. In asynchronous-type algorithms, workers do not wait for each other, thus a server immediately updates the current model and assigns a new job to available workers. All of the aforementioned approaches are orthogonal to asynchronicity, and can be combined with it (Nguyen et al., 2022; Zakerinia et al., 2022).

Recently, (Koloskova et al., 2022; Mishchenko et al., 2022a) improved theoretical analysis of pure asynchronous SGD in the *homogeneous data* setting, i.e., when workers have access to the same data, and make a step towards better understanding in a heterogeneous regime. Here, we specifically focus on the more challenging heterogeneous setting, and provide a unified convergence analysis of asynchronous SGD.

We start from the observation that asynchronous SGD can be seen as an instance of SGD with arbitrary data orderings, which include random reshuffling at each step, shuffle once (at the beginning of training) or some incremental order. The only difference between these variants is in the order defined naturally according to the computation speeds of the workers. Data ordering might improve the performance of the algorithm. For example, several works (Mishchenko et al., 2020; Nguyen et al., 2021) show that SGD with random reshuffling and shuffle once is always better than vanilla SGD in the strongly convex case, and can outperform it for both convex and non-convex objectives if the number of epochs is sufficiently large (Mohtashami et al., 2022; Lu et al., 2022). Recently, (Koloskova et al., 2022) analyzed SGD with arbitrary data ordering as an algorithm with linearly correlated noise (Koloskova et al., 2023b). Here, we go further by linking asynchronous SGD and SGD using various data orderings.

1.1 Contributions

In this section, we summarize the key contributions of our work.

- We propose a theoretical framework that allows us to analyze various types of asynchronous SGD in a unified manner. Besides purely-asynchronous SGD, we analyze variants of asynchronous SGD when a server waits for the first b fastest workers or assigns new jobs according to some scheduling procedure. The analysis is performed for a constant stepsize schedule, without bringing any additional hyper-parameters except for stepsize, as all the parameters used in the algorithm are known during training.
- Moreover, this unified framework enables us to design and analyze new asynchronous algorithms. In particular, we propose a new method called *shuffled asynchronous SGD* and show its superiority over com-

petitive methods both theoretically and practically.

- Our framework also recovers popular synchronous variants of SGD, such as SGD with random reshuffling, shuffle once, and mini-batch SGD. For these methods, we derive the best-known convergence results, without any changes in the analysis, highlighting the tightness of our framework.
- All of our results have a better or similar dependencies on the maximum delay, compared with existing work. With bounded gradient assumption, we remove entirely dependencies on maximum delay used by prior works.

2 RELATED WORKS

2.1. Asynchronous SGD. The literature on asynchronous-type SGD algorithms is extremely rich, starting from (Baudet, 1978), and (Bertsekas and Tsitsiklis, 1989). The source of asynchrony might be caused by various factors, such as different hardware (Horvath et al., 2022) or failures (Nadiradze et al., 2021).

Most existing works are devoted to the analysis of SGD-based asynchronous algorithms, concentrating on the homogeneous regime (Cohen et al., 2021; Aviv et al., 2021; Stich et al., 2021) which typically holds only for shared memory architectures. In our work, we concentrate on the more challenging heterogeneous setting.

The first convergence guarantees were given for constant delays (Stich and Karimireddy, 2020; Arjevani et al., 2020) which usually is not true in practice. Follow-up studies provide an analysis depending on the maximum delay (Stich et al., 2021; Zakerinia et al., 2022; Nguyen et al., 2022), however, it can be much larger than the average delay. (Alistarh et al., 2018) improved the dependency on the maximum delay under the bounded gradient assumption while (Cohen et al., 2021; Aviv et al., 2021) completely removed that dependency. (Cohen et al., 2021) propose a method which sends delayed models together with computed gradients, and additional parameter tuning while (Aviv et al., 2021) derived rates involving the variance of the delays. Nonetheless, in some cases, this quantity might be proportional to the maximum delay. Recently, (Tyurin and Richtarik, 2023) introduced a new asynchronous algorithm with optimal convergence guarantees; however, under the assumption that the computational speeds of workers are *fixed* throughout the execution. This assumption is too restrictive in many practical scenarios, where computational speeds usually fluctuate during the training. In parallel work, (Even et al., 2024) create a unified framework to analyze asynchronous algorithms for decentralized

networks and derive convergence guarantees in that regime.

Several works consider asynchronous SGD in Federated Learning, where workers frequently have heterogeneous computational power. (Zheng et al., 2017; Zhang et al., 2016) utilize delay-adaptive step-sizes to mitigate the effect of asynchrony while (Yan et al., 2020; Glasgow and Wootters, 2021; Gu et al., 2021) proposed variance-reduction-based mechanisms to handle different worker availability. (Nguyen et al., 2022) proposed a method that incorporates local steps and shows its practical superiority, but the analysis is done under unrealistic assumptions.

Recently, (Koloskova et al., 2022; Mishchenko et al., 2022a) introduced novel analysis based on perturbed iterates framework (Mania et al., 2017) which improves the convergence guarantees for asynchronous SGD for both heterogeneous and homogeneous regimes. They propose similar delay-adaptive stepsize schedules that allow to derive maximum-delay-free rates for asynchronous SGD. Besides, (Koloskova et al., 2022) proposed an algorithm with a special random job assigning procedure to balance the workers in the heterogeneous setting. In the same setting (Mishchenko et al., 2022a) derived convergence guarantees for pure asynchronous SGD. Inspired by (Koloskova et al., 2023a)¹, we propose a new virtual-iterates-based analysis and cover the results of (Koloskova et al., 2022; Mishchenko et al., 2022a) as special cases.

2.1 SGD with Arbitrary Data Ordering

The most practical training SGD schemes utilize a certain order (usually random) of data samples. Schemes like random reshuffling or shuffle once are considered a default choice in many real-world applications. Sampling without replacement allows to leverage of the finite-sum structure of the problem since all data points will contribute equally to the solution. However, the theoretical guarantees for these methods are much less studied than for vanilla SGD (Rakhlin et al., 2012; Nguyen et al., 2021). The main complication comes from the fact that the gradient estimator becomes biased. Biased SGD-type methods are typically considered more challenging to analyze than their unbiased counterparts.

In the last years, particular attention was given to federated and distributed methods with random reshuffling (Malinovsky et al., 2022; Sadiev et al., 2022; Yun

¹We aware of the technical gap in the proofs in (Koloskova et al., 2023a). However, a proper choice of correlation period τ allows to take conditional expectations correctly. We highlight for each special case the choice of τ that leads to correct statements.

et al., 2021; Cho et al., 2023). Data shuffling might provably improve the convergence in this case as well which is important for the training of large-scale models.

Here, (Mishchenko et al., 2020) introduce novel techniques to improve the convergence guarantees in the strongly convex case, whereas (Nguyen et al., 2021; Ahn et al., 2020) demonstrate the same rates showing the superiority of SGD with random reshuffling over vanilla SGD. However, in convex and non-convex regimes this remains an open problem. Existing results demonstrate better performance of SGD with random reshuffling in some special cases (Mishchenko et al., 2020) in convex and non-convex scenarios.

3 SETUP

We consider the classical Empirical Risk Minimization (ERM) problem of the form

$$\min_{x \in \mathbb{R}^d} \left[f(x) := \frac{1}{n} \sum_{i=1}^n f_i(x) \right], \quad (1)$$

which covers many optimization problems in machine learning. Here, $x \in \mathbb{R}^d$ denotes the parameters of a model we aim to train, n is the number of workers participating in the distributed training, and $f_i(x)$ is the loss associated with data \mathcal{D}_i available to worker $i \in [n] := \{1, 2, \dots, n\}$. For example, the local function f_i can be written as an expectation $f_i(x) := \mathbb{E}_{\xi \sim \mathcal{D}_i} [f_i(x, \xi)]$ over dataset \mathcal{D}_i locally stored at worker i . Besides, the formulation (1) also recovers the single-node setting. In this scenario, f represents a loss over a dataset of size n , and each f_i is a loss associated with i -th data point. We denote the minimum of the problem (1) by f^* , and assume that it is finite.

3.1 Proposed Algorithmic Framework

To solve the problem (1), we apply general asynchronous SGD (Algorithm 1). The backbone of our algorithmic approach is to analyze updates of the form

$$x_{t+1} = x_t - \gamma g_{i_t}(x_{\pi_t}), \quad (2)$$

where $i_t \in [n]$ is an index of a worker whose potentially stale and stochastic gradient $g_{i_t}(x_{\pi_t})$ computed at an outdated point x_{π_t} is applied at iteration t .

It is crucial to note that the order of applied gradients $g_{i_t}(x_{\pi_t})$ is not controlled by the server. Once an ongoing gradient computation is done by some worker, the gradient is immediately communicated to the server and is used to update the global model without synchronizing with other workers. In fact, the order of

Table 1: Asynchronous algorithms whose convergence analysis is covered by our framework. Here constants ζ^2 , G^2 such that $\|\nabla f_i(x) - \nabla f(x)\|^2 \leq \zeta^2$ and $\|\nabla f_i(x)\|^2 \leq G^2$ for all $i \in [n]$ and $x \in \mathbb{R}^d$. τ_C and τ_{\max} denote the maximum number of active jobs and the maximum delay respectively. F_0 is the initial functional suboptimality, the definitions of other constants are given in Section 4. For shuffled asynchronous SGD $\tau_C = n$. **BG** = requires **Bounded Gradients** assumption.

| Method | Alg # | Citation | BG | Rate ^(a) |
|--|-------|---------------------------|-----|---|
| Pure Asynchronous SGD | Alg 2 | Mishchenko et al. (2022a) | No | $\frac{LF_0\tau_C}{T} + \left(\frac{LF_0\sigma^2}{T}\right)^{1/2} + \zeta^2$ ^(b) |
| | | Ours | No | $\frac{LF_0\sqrt{\tau_{\max}\tau_C}}{T} + \left(\frac{LF_0\sigma^2}{T}\right)^{1/2} + \zeta^2$ |
| | | Ours | Yes | $\frac{LF_0\tau_C}{T} + \left(\frac{LF_0\sigma^2}{T}\right)^{1/2} + \left(\frac{LF_0G\tau_C}{T}\right)^{2/3} + \zeta^2$ |
| Pure Asynchronous SGD with waiting | Alg 3 | Ours | No | $\frac{LF_0\sqrt{\tau_{\max}\tau_C}}{T\sqrt{b}} + \left(\frac{LF_0\sigma^2}{Tb}\right)^{1/2} + \zeta^2$ |
| | | Ours | Yes | $\frac{LF_0\tau_C}{Tb} + \left(\frac{LF_0\sigma^2}{Tb}\right)^{1/2} + \left(\frac{LF_0G\tau_C}{Tb}\right)^{2/3} + \zeta^2$ |
| Random Asynchronous SGD | Alg 4 | Koloskova et al. (2022) | No | $\frac{LF_0\sqrt{\tau_{\max}\tau_C}}{T} + \left(\frac{LF_0\sigma^2}{T}\right)^{1/2} + \left(\frac{LF_0\zeta^2}{T}\right)^{1/2} + \left(\frac{LF_0\tau_C\zeta}{T}\right)^{2/3}$ |
| | | Koloskova et al. (2022) | Yes | $\frac{LF_0\tau_C}{T} + \left(\frac{LF_0\sigma^2}{T}\right)^{1/2} + \left(\frac{LF_0\zeta^2}{T}\right)^{1/2} + \left(\frac{LF_0\tau_C G}{T}\right)^{2/3}$ |
| | | Ours | Yes | $\frac{LF_0\tau_C}{T} + \left(\frac{LF_0\sigma^2}{T}\right)^{1/2} + \left(\frac{LF_0\zeta^2}{T}\right)^{1/2} + \left(\frac{LF_0\tau_C G}{T}\right)^{2/3}$ |
| Random Asynchronous SGD with waiting (FedBuff) | Alg 5 | Nguyen et al. (2022) | Yes | $\frac{LF_0}{T} + \left(\frac{LF_0\sigma^2}{T}\right)^{1/2} + \left(\frac{LF_0\zeta\tau_{\max}}{T}\right)^{2/3} + \left(\frac{LF_0G\tau_{\max}}{T}\right)^{2/3}$ ^(c) |
| | | Ours | Yes | $\frac{LF_0\tau_C}{T} + \left(\frac{LF_0\zeta^2}{Tb}\right)^{1/2} + \left(\frac{LF_0\sigma^2}{Tb}\right)^{1/2} + \left(\frac{LF_0\tau_C G}{Tb}\right)^{2/3}$ |
| Shuffled Asynchronous SGD [NEW] | Alg 6 | Ours | Yes | $\frac{LnF_0}{T} + \left(\frac{LF_0\sigma^2}{T}\right)^{1/2} + \left(\frac{LF_0\sqrt{n}\zeta}{T}\right)^{2/3} + \left(\frac{LF_0Gn}{T}\right)^{2/3}$ |

(a) We present the best-known rates under the same set of assumptions as we use in the analysis.

(b) Mishchenko et al. (2022a) uses delay adaptive stepsizes to get rid of the dependency on τ_{\max} .

(c) If we set $\eta_l = \frac{\gamma}{b}$, $\eta_g = b$, $Q = 1$ in Theorem 1 Nguyen et al. (2022). The analysis is done under the unrealistic assumption that $\{i_t\}_{t=0}^{T-1}$ are distributed uniformly at random.

received gradients depends on the speeds of the workers which can change during the training. We denote by \mathcal{R}_t the set of all received jobs² before iteration t . Hence, before training $\mathcal{R}_0 = \emptyset$, and once a new job (i_t, π_t) is completed (i.e., the gradient $g_{i_t}(x_{\pi_t})$ is computed) we update the set $\mathcal{R}_{t+1} = \mathcal{R}_t \cup \{(i_t, \pi_t)\}$.

After receiving a completed job (i_t, π_t) and updating the model parameters x_{t+1} , the server proceeds to assign a new job. This is the stage when the server can influence the training by controlling the assigned jobs. Specifically, the server is allowed to assign any worker $k_{t+1} \in [n]$ to compute a stochastic gradient at any point $x_{\alpha_{t+1}}$ in the history of model parameters. Similar to the set of received jobs, denote by \mathcal{A}_t the set of all assigned jobs before iteration t . Thus, before the training no job is assigned, i.e., $\mathcal{A}_0 = \emptyset$. Then, the training starts by the initial job assignments $\mathcal{A}_1 = \{(i, 0) : \text{for some workers } i \in [n]\}$. Afterward, once a new job (k_{t+1}, α_{t+1}) is assigned by the server, we update the set $\mathcal{A}_{t+2} = \mathcal{A}_{t+1} \cup \{(k_{t+1}, \alpha_{t+1})\}$.

By definition, $\mathcal{R}_t \subseteq \mathcal{A}_t$ since only assigned jobs might be finished. Besides, the set $\mathcal{A}_{t+1} \setminus \mathcal{R}_t$ represents the

²By job we define a pair (i, j) such that worker i is assigned to compute $\nabla f_i(x_j)$ for a model from iteration j .

jobs that are ‘‘in flight’’ at iteration t . In particular, $\mathcal{A}_{T+1} \setminus \mathcal{R}_T$ consists of all jobs that are not finished within the optimization process.

The power of our algorithmic framework is its versatility to recover various variants of both asynchronous and synchronous SGD algorithms in one method. The strength of our approach is that the theoretical analysis for all of them is covered by unified theory, as well as the rates obtained by our theory match the best-known results for those cases or improve them. We present convergence guarantees for several synchronous and asynchronous methods covered by our analysis in Table 1 and Section 4 in more details.

3.2 Special Cases

Below we list some interesting special cases covered by our framework. Detailed description is provided in Sections C.3 and D.3 while convergence rates are presented in Table 1.

Pure Asynchronous SGD. In the beginning of the training the server assigns jobs to all workers at x_0 . Then, a new job (for a freshly updated model) is assigned back to the same worker which completed the

Algorithm 1 General Asynchronous SGD

- 1: **Input:** $x_0 \in \mathbb{R}^d$, stepsize $\gamma > 0$, set of assigned jobs $\mathcal{A}_0 = \emptyset$, set of received jobs $\mathcal{R}_0 = \emptyset$
 - 2: **Initialization:** for all jobs $(i, 0) \in \mathcal{A}_1$, the server assigns worker i to compute a stochastic gradient $g_i(x_0)$
 - 3: **for** $t = 0, 1, 2, \dots, T - 1$ **do**
 - 4: once worker i_t finishes a job $(i_t, \pi_t) \in \mathcal{A}_{t+1}$, it sends $g_{i_t}(x_{\pi_t})$ to the server
 - 5: server updates the current model $x_{t+1} = x_t - \gamma g_{i_t}(x_{\pi_t})$ and the set $\mathcal{R}_{t+1} = \mathcal{R}_t \cup \{(i_t, \pi_t)\}$
 - 6: server assigns worker k_{t+1} to compute a gradient $g_{k_{t+1}}(x_{\alpha_{t+1}})$
 - 7: server updates the set $\mathcal{A}_{t+2} = \mathcal{A}_{t+1} \cup \{(k_{t+1}, \alpha_{t+1})\}$
 - 8: **end for**
-

previous job, i.e., $(k_{t+1}, \alpha_{t+1}) \equiv (i_t, t + 1)$. We derive improved square root dependency on τ_{\max} in contrast to previous works, and remove this dependency completely with bounded gradients assumption. Our rate matches the best-known result (Koloskova et al., 2022) in homogeneous regime (i.e., $\zeta^2 = 0$).

Pure Asynchronous SGD with waiting. In contrast to the previous case, the server waits for $b \geq 1$ workers to finish their jobs, and then it assigns to them (i.e., $k_{t+1} = i_t$ as before) new jobs for the same updated model with $\alpha_{t+1} = \lfloor t+1/b \rfloor b$. In comparison with previous method, waiting for b leads to faster convergence while keeping the same dependency τ_{\max} .

Random Asynchronous SGD Koloskova et al. (2022). In this version, a new job is assigned to the worker $k_{t+1} \sim \text{Uni}[1, 2, \dots, n]$ chosen independently and uniformly at random among all workers for the latest model, i.e., $\alpha_{t+1} = t + 1$. Thus, some workers might receive new jobs without completing the current one. We obtain the same rates as in (Koloskova et al., 2022) using more general theory which indicates the sharpness of our approach.

Random Asynchronous SGD with waiting Nguyen et al. (2022). This method is a special case of the FedBuff algorithm Nguyen et al. (2022) with $Q = 1$ local steps. Besides, it can be seen as a combination of the previous two special cases, namely asynchronous SGD with waiting and random assignments. The server waits for the first b fastest workers and then assigns new jobs with the same model $\alpha_{t+1} = \lfloor t+1/b \rfloor b$ to b randomly chosen workers $k_{t+1} \sim \text{Uni}[1, \dots, n]$. Unlike (Nguyen et al., 2022), we derive convergence under realistic assumptions. Moreover, we derive τ_{\max} -free rate and show the benefit from waiting for few workers as the rate improves with b .

Shuffled Asynchronous SGD [NEW]. In this case, we assume that all participating workers $[n]$ are active in the training. Similar to random asynchronous SGD described above, new jobs are always for the latest model, i.e., $\alpha_{t+1} = t + 1$. However, new jobs are not assigned to workers independently but rather based on a random permutation of workers that can be re-

sampled after each cycle or sampled once and reused throughout the training. More specifically, if χ is a random permutation of indices $[n]$, then $k_{t+1} = \chi(j)$, where $j - 1 = t \pmod n$ is the remainder of t when divided by n . In Section D.3.3 we demonstrate that new method outperforms its random counterpart in highly heterogeneous regime $\zeta \geq \sqrt{ne}^{1/2}$ which typically holds in Federated Learning.

Mini-batch SGD. This is the standard variety of SGD method, and a popular method for a single-node setting. We show that the update rule of mini-batch SGD can be modified to suit update rule (2) and derive standard convergence rate. In particular, if we treat each data point as a separate client then mini-batch SGD can be viewed as random asynchronous SGD with waiting where initial number of jobs assigned by the server is b .

SGD with Random Reshuffling Nedic and Bertsekas (2001). SGD with random reshuffling is one the most used and sometimes a default algorithm in practice to train neural networks. At the beginning of each epoch, the dataset is randomly shuffled, and gradients are computed following that random order. Similar to the analogy described for mini-batch SGD, we can view SGD with random reshuffling as a special case of shuffled asynchronous SGD. Our rate matches the best known guarantees showing the tightness of our approach.

4 CONVERGENCE THEORY

4.1 Theoretical assumptions

Below we list the assumptions we use in the theoretical analysis. All of them are standard in the distributed non-convex optimization literature.

Assumption 1. *Local functions f_i are differentiable and L -smooth for some positive constant L , namely,*

$$\|\nabla f_i(x) - \nabla f_i(y)\| \leq L\|x - y\|, \quad \forall x, y \in \mathbb{R}^d. \quad (3)$$

For some of our results, we also need a bounded variance assumption on stochastic gradients.

Assumption 2. Stochastic gradients $\nabla f_i(x, \xi)$ are unbiased estimators of $\nabla f_i(x)$, i.e.,

$$\mathbb{E}_{\xi \sim \mathcal{D}_i} [\nabla f_i(x, \xi)] = \nabla f_i(x), \quad \forall x \in \mathbb{R}^d, \quad (4)$$

and have bounded variance $\sigma^2 \geq 0$, namely,

$$\mathbb{E}_{\xi \sim \mathcal{D}_i} [\|\nabla f_i(x, \xi) - \nabla f_i(x)\|^2] \leq \sigma^2, \quad \forall x \in \mathbb{R}^d. \quad (5)$$

This is a typical assumption in the literature, and it holds, for example, when we have access to the gradients with Gaussian noise. We denote a realization of $\nabla f_i(x, \xi)$ by $g_i(x)$ for shortness.

Next, we also assume that the bounded function heterogeneity assumption holds since in general case it is not possible to derive any convergence guarantees for asynchronous algorithms.

Assumption 3. Local gradients $\nabla f_i(x)$ satisfy bounded heterogeneity condition for some $\zeta^2 \geq 0$, i.e.,

$$\|\nabla f_i(x) - \nabla f(x)\|^2 \leq \zeta^2, \quad \forall x \in \mathbb{R}^d. \quad (6)$$

Several results require the Lipschitzness of local loss functions.

Assumption 4. Local functions $f_i(x)$ are G -Lipschitz, i.e. for some positive constant G they satisfy

$$|f_i(x) - f_i(y)| \leq G\|x - y\| \quad \forall x, y \in \mathbb{R}^d. \quad (7)$$

Note that, in the case of differentiable f_i , this assumption implies that local gradients are bounded, i.e., for all $x \in \mathbb{R}^d$ $\|\nabla f_i(x)\| \leq G$ [Bubeck \(2015\)](#). In contrast to [Mishchenko et al. \(2022a\)](#), we do not assume the boundedness of stochastic gradients. Practical implementations frequently resort to using clipping in the presence of Byzantine workers or stragglers. The clipping automatically bounds the norms of applied gradient, forcing the constant G^2 to be small.

4.2 Notation

Generally, we do not make any assumptions on the delays — gradients might be received in any random or deterministic order. We assume that the server can receive and assign jobs with delays, namely, $\pi_t := t - \tau_t$ and $\alpha_t := t - \tilde{\tau}_t$, where $\tau_t, \tilde{\tau}_t \geq 0$ are corresponding delays.³ The order might be natural as in the case of pure asynchronous SGD or be pre-set as for mini-batch SGD; see Sections [C.3](#) and [D.3](#) for more examples of how $\{i_t\}_{t=0}^{T-1}$ might look like. Besides, we introduce the notion of maximum and average delays similar to [Koloskova et al. \(2022\)](#).

³If $\tau_t \equiv \tilde{\tau}_t \equiv 0$, then there is no delay.

Definition 1. Let $\{\tau_t\}_{t=0}^{T-1}$ be the delays of all applied gradients. The average and maximum delays are defined as follows

$$\begin{aligned} \tau_{\text{avg}} &:= \frac{1}{|\mathcal{A}_{T+1}|} \left(\sum_{t=0}^{T-1} \tau_t + \sum_{(i,j) \in \mathcal{A}_{T+1} \setminus \mathcal{R}_T} T - j \right), \\ \tau_{\text{max}} &:= \max \left\{ \max_{0 \leq t < T} \tau_t, \max_{(i,j) \in \mathcal{A}_{T+1} \setminus \mathcal{R}_T} T - j \right\}. \end{aligned} \quad (8)$$

Quantities $\tilde{\tau}_{\text{avg}}$ and $\tilde{\tau}_{\text{max}}$ are defined analogously with respect to the delays $\{\tilde{\tau}_t\}_{t=0}^{T-1}$ related to assigning process. Moreover, in the analysis, we use the maximum number of active jobs or concurrency τ_C . This quantity indicates the maximum number of jobs already assigned, but not yet completed (i.e., active jobs) during the optimization process.

Definition 2. The maximum number of active jobs or concurrency is defined as

$$\tau_C := \max_{0 \leq t \leq T} |\mathcal{A}_{t+1} \setminus \mathcal{R}_t|. \quad (9)$$

To utilize available resources in a more efficient way, in practice, all workers are always busy, i.e., $\tau_C = n$. Nevertheless, it might happen that a fraction of all workers can be unavailable from time to time.

The received $\{i_t\}$ and assigned $\{k_t\}$ orders of functions define the convergence properties of Algorithm 1. Hence, we are interested in the correlation between functions within a certain correlation interval τ . The final rate depends on how much the functions within the correlation interval differ from the averaged gradient. To mathematically describe the aforesaid, we define the sequence correlation [Koloskova et al. \(2023a\)](#) below.

Definition 3. For any given correlation period $\tau \geq 1$, we successively split the set of received gradient indices $\{i_t\}_{t=0}^{T-1}$ into $\lceil \frac{T}{\tau} \rceil$ chunks of size τ . Then, the sequence correlation of received jobs within k -th period is defined as

$$\sigma_{k,\tau}^2 := \max_{0 \leq j < \tau} \mathbb{E} \left[\left\| \sum_{t=k\tau}^{\min\{k\tau+j, T-1\}} \Delta_{i_t}(x_{k\tau}) \right\|^2 \right], \quad (10)$$

where $\Delta_i(x) := \nabla f_i(x) - \nabla f(x)$.

Next, the magnitude of delays affects the resulting convergence rate. We measure the effect of the delays by the quantity defined below. Note that it does not involve any correlation period since it is designed to track how π_t impacts the rate.

Definition 4. For the sequence of received gradient indices $\{i_t\}_{t=0}^{T-1}$ the delay variance is defined as

$$\nu^2 := \sum_{t=0}^{T-1} \mathbb{E} \left[\left\| \sum_{j=\pi_t}^{t-1} \Delta_{i_j}(x_{\pi_j}) \right\|^2 \right].$$

4.3 Analysis of Gradient Receiving Process

In the first theorem, we analyze the gradient receiving process, i.e., how the order of received gradients $\{i_t\}_{t=0}^{T-1}$ influences the convergence. We would like to highlight that this ordering can not be controlled by the server; it is built naturally depending on the speeds of the workers.

For the analysis, we additionally construct a sequence of virtual iterates Mania et al. (2017) with restarts following the approach created in Koloskova et al. (2023a). The sequence is defined as $\tilde{x}_0 = x_0$ and

$$\tilde{x}_{t+1} = \begin{cases} \tilde{x}_t - \gamma \nabla f(x_t) & \text{if } t+1 \neq 0 \pmod{\tau}, \\ x_{t+1} & \text{if } t+1 = 0 \pmod{\tau}. \end{cases} \quad (11)$$

In contrast to the sequence $\{x_t\}_{t=0}^T$, virtual iterates $\{\tilde{x}_t\}_{t=0}^{T-1}$ are updated using full gradient always evaluated for the last model x_t . Moreover, we restart the virtual iterates once in τ iterations, so we can track the progress of real iterates within one correlation period. More particularly, we use $\tau = \Theta(\frac{1}{L\gamma})$, where L and γ are the smoothness constant and the stepsize respectively.

Theorem 1. *Let Assumptions 1 and 2 hold. Let the stepsize γ satisfy inequalities $6L\gamma \leq 1$ and $20L\gamma\sqrt{\tau_{\max}\tau_C} \leq 1$, the correlation period $\tau = \lfloor \frac{1}{20L\gamma} \rfloor$, and quantities $\{\sigma_{k,\tau}^2\}_{k=0}^{\lfloor \frac{T}{\tau} \rfloor}$ and ν^2 are finite. Then*

$$\mathbb{E} [\|\nabla f(\hat{x}_T)\|^2] \leq \mathcal{O} \left(\frac{F_0}{\gamma T} + L\gamma\sigma^2 + L^2\gamma^2\Phi \right), \quad (12)$$

where $F_0 := f(x_0) - f^*$, $\Phi := \frac{1}{\lfloor \frac{T}{\tau} \rfloor} \sum_{k=0}^{\lfloor \frac{T}{\tau} \rfloor} \sigma_{k,\tau}^2 + \frac{1}{\tau}\nu^2$, and \hat{x}_T is chosen uniformly at random from $\{x_0, \dots, x_{T-1}\}$.

We give a detailed proof of the theorem in Appendix, Section C. We observe that the rate consists of three terms. The first one, $\frac{F_0}{\gamma T}$, is a standard optimization term that always appears for vanilla GD, and is shown to be optimal. The second term, $L\gamma\sigma^2$, appears because of the stochastic nature of the gradients. The third term, $L^2\gamma^2\Phi$, the most intriguing one, consists of two parts — the first one represents the effect of function ordering while the second part shows the impact of delays on the convergence.

Since the first and second terms are standard for gradient-based algorithms, one asynchronous algorithm differs from another one with a bound on $L^2\gamma^2\Phi$. Intuitively, our goal is to create an algorithm for which this term is as small as possible in order to guarantee better convergence. This can be achieved by properly balancing the workers' contributions. Moreover, the

stepsize γ is decreased by $\tau_{\max}^{1/2}$ to mitigate the effect of the delays.

Note that all quantities $\{\sigma_{k,\tau}^2\}_{k=0}^{\lfloor \frac{T}{\tau} \rfloor}$ and ν^2 depend on τ , and consequently, on the stepsize γ as well. Hence, for the general case, the inequality (12) is implicit. However, for some special cases, we are able to compute all quantities and derive convergence guarantees.

4.4 Analysis of Gradient Assigning Process

Now we switch to the analysis of the order $\{k_t\}_{t=1}^T$ which the server uses to decide the order of assigning new jobs. Recall that the server is able to control this order. That is why we can use various randomization procedures to balance the workers.

The analysis in this case is based on the virtual iterates $\{y_t\}_{t=0}^T$ that follow the assigning process. Formally, we define $y_0 = x_0$ and

$$y_{t+1} \stackrel{t \geq 0}{=} y_t - \gamma \sum_{(i,j) \in \mathcal{A}_{t+1} \setminus \mathcal{A}_t} g_i(x_j) - \gamma g_{k_t}(x_{\alpha_t}). \quad (13)$$

Here we highlight that the server may decide to send a job at outdated point x_{α_t} with bounded by $\tilde{\tau}_{\max}$ delay.⁴ This enables us to investigate methods where several workers compute gradients at the same point, e.g., mini-batch SGD. Hence, it brings even more flexibility to our framework. The real and virtual iterates can not be arbitrarily far away from each other. The next Lemma reveals the connection between them.

Lemma 1. Let real $\{x_t\}_{t=0}^T$ and virtual $\{y_t\}_{t=0}^T$ iterates be defined in (2) and (13) respectively. Then

$$x_t - y_t = \gamma \sum_{(i,j) \in \mathcal{A}_t \setminus \mathcal{R}_t} g_i(x_j).$$

Based on the sequence $\{y_t\}_{t=0}^T$ we construct the restarting sequence $\{\tilde{y}_t\}_{t=0}^T$ similarly to (11)

$$\begin{aligned} \tilde{y}_1 &= y_1, \\ \tilde{y}_{t+1} &= \begin{cases} \tilde{y}_t - \gamma \nabla f(x_t) & \text{if } t \neq 0 \pmod{\tau}, \\ y_{t+1} & \text{if } t = 0 \pmod{\tau}. \end{cases} \end{aligned} \quad (14)$$

Analogously, we can define the sequence correlation and the delay variance for the sequence of assigned gradient indices $\{k_t\}_{t=1}^T$. We denote them by $\{\tilde{\sigma}_{k,\tau}^2\}_{t=0}^{\lfloor \frac{T}{\tau} \rfloor}$ and $\tilde{\nu}^2$ respectively. Based on this, we present our second theorem.

Theorem 2. *Let Assumptions 1, 2, and 4 hold. Let the stepsize γ satisfies inequalities $6L\gamma \leq 1$ and $30L\gamma \max\{\tilde{\tau}_{\max}, \tau_C\} \leq 1$, the correlation period $\tau =$*

⁴In other words, $\alpha_t = t - \tilde{\tau}_t$ where $\tilde{\tau}_t \geq 0$.




Figure 1: Comparison of pure, random, and shuffled asynchronous SGD with tuned stepsizes and full gradient computation on w7a dataset with various delay patterns. Here $n = 10$, $\lambda = 0.1$, $d = 300$, $m = 2505$.

$\left\lfloor \frac{1}{30L\gamma} \right\rfloor$, quantities $\{\tilde{\sigma}_{k,\tau}^2\}_{k=0}^{\lfloor T/\tau \rfloor}$ and $\tilde{\nu}^2$ are finite. Then

$$\mathbb{E} [\|\nabla f(\hat{x}_T)\|^2] \leq \mathcal{O}\left(\frac{F_1}{\gamma T} + L\gamma\sigma^2 + L^2\gamma^2\tilde{\Phi} + L^2\gamma^2(\tau_C - 1)^2G^2\right), \quad (15)$$

where $F_1 := f(y_1) - f^*$, $\tilde{\Phi} := \frac{1}{\lfloor T/\tau \rfloor} \sum_{k=0}^{\lfloor T/\tau \rfloor} \tilde{\sigma}_{k,\tau}^2 + \frac{1}{T}\tilde{\nu}^2$, and \hat{x}_T is chosen uniformly at random from $\{x_1, \dots, x_T\}$.

The proof of Theorem 2 can be found in Section D. It gives similar convergence behaviour as Theorem 1, but for the assigning sequence $\{k_t\}_{t=1}^T$. The difference is in the last $L^2\gamma^2(\tau_C - 1)^2G^2$ term which is not present in (12), and it appears because of the bound between real and virtual iterates.

We emphasize that if $\tau_C = 1$, i.e., there is only one active job at each iteration, then the fourth term is zero. For example, this holds for single-node algorithms such as SGD with random reshuffling or shuffle once. Therefore, we recover the same rates for those methods. Besides, we highlight the fact that the rate (15) admits the absence of the dependency on τ_{\max} in particular special cases; the delays affect the rate through averaged delays only.

Moreover, the term $\tilde{\Phi}$ depends on sequences $\{k_t\}$ and $\{\alpha_t\}$ which are fully under the control of the server. Therefore, the server may change its job assignment strategy if it observes that the current one leads to poor performance. In Section D.3 we demonstrate precisely which assignment strategies can be used.

5 EXPERIMENTS

We conduct experiments on logistic regression with nonconvex regularization, namely, we consider the problem (1) with local functions

$$f_i(x) = \frac{1}{m} \sum_{j=1}^m \log\left(1 + e^{-b_{ij}a_{ij}^\top x}\right) + \lambda \sum_{j=1}^d \frac{x_j^2}{1+x_j^2}.$$

We test the performance of FedBuff (Nguyen et al., 2022), pure, random, and shuffled versions of asynchronous SGD methods. We use the w7a dataset from

the LibSVM library Chang and Lin (2011), and additionally construct synthetic datasets following Li et al. (2018); Safaryan et al. (2022) with control of statistical heterogeneity through parameters α and β .⁵ The detailed dataset generation and parameters setup are given in Appendix A.⁶

We assume that all workers start computing gradients in the beginning, i.e. $\tau_C = n$. We fine-tune the stepsize and illustrate the performance with the fine-tuned stepsize only choosing the best one that achieves smaller gradient norm and fluctuations.

We conducted experiments in four different settings. In particular, each worker has a positive parameter s_i . The way these parameters are used is as follows

- **Fixed:** in this case, $r \equiv s_i$. Such fixed timing implies that the delay pattern is fixed as well.
- **Poisson:** for each worker we sample $r \sim \text{Po}(s_i)$.
- **Normal:** similarly, we sample $s \sim \mathcal{N}(s_i, s_i)$, and then set $r = |s| + 1$.
- **Uniform:** in this case, we sample $r \sim \text{Uni}(0, s_i)$.

The number r indicates how many seconds worker i spends to compute one gradient. Such timing patterns simulate possible workers' behavior in practice.

First, we present the practical performance of the three aforesaid methods with full gradient computation in Figure 1. We observe that pure asynchronous SGD gets stuck when the gradient norm becomes about $10^{-1} - 10^{-2}$. The same result holds for FedBuff. The rest two methods achieve better stationary points with smaller gradient norms. We also observe that FedBuff does not benefit from performing several local steps. Moreover, shuffled asynchronous SGD starts converging faster than its random counterpart around $10^{-1.5}$. Further, shuffled asynchronous SGD achieves a stationary point with roughly 10 times smaller gradient norm.

⁵The corresponding dataset is denoted as $\text{Syn}(\alpha, \beta)$.

⁶Our code is available at <https://github.com/Rustem-Islamov/AsGrad-public/>.




Figure 2: Comparison of pure, random, and shuffled asynchronous SGD with tuned stepsizes and stochastic gradient computation on synthetic datasets with different levels of heterogeneity and poisson delay pattern. Here $n = 10$, $\lambda = 0.1$, $d = 300$, $m = 200$, and batch size $m/10$.

Next, we conduct the same set of experiments with stochastic gradients on synthetic datasets changing parameters α and β ; see results in Figure 2.

We observe similar performance in stochastic setting. Random and shuffled asynchronous SGD are superior to pure asynchronous SGD. Our proposed method always finds the stationary point with the smallest error, and in several cases, the convergence is faster than that of random asynchronous SGD.

6 CONCLUSION

Our framework provides a deeper understanding of asynchronous SGD-type algorithms. In contrast to previous works, we derive a theoretical analysis for few method simultaneously with improved dependency on τ_{\max} . In addition, our theory provides the intuition of what affects the convergence of asynchronous-type algorithms, and how we can improve the convergence by balancing the workers’ jobs. The analysis highlights two terms appearing in the rate due to delays and data ordering caused by asynchronicity. We show that different prior techniques and our new proposed approach, in fact, make those two terms in the rate smaller, and as a consequence, lead to better practical performance. Moreover, we do not impose any assumption on the delay pattern.

Finally, experiments support our theoretical findings. We both empirically and theoretically show that pure asynchronous SGD converges up to heterogeneity level. Besides, we demonstrate that the proposed shuffled job assigning allows to balance more carefully the workers’ contributions.

7 Acknowledgments

The authors thank all anonymous reviewers for their valuable comments and suggestions on how to improve the manuscript. This work was done when Rustem Islamov was a Master’s student at Institut Polytechnique de Paris (IP Paris) and an intern at Institute of

Science and Technology Austria (ISTA). The research of Rustem Islamov was supported by ISTA internship program. Mher Safaryan has received funding from the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 101034413.

References

- Kwangjun Ahn, Chulhee Yun, and Suvrit Sra. Sgd with shuffling: optimal rates without component convexity and large epoch requirements. In *Advances in Neural Information Processing Systems (NeurIPS 2020)*, 2020. (Cited on page 3)
- Dan Alistarh, Christopher De Sa, and Nikola Konstantinov. The convergence of stochastic gradient descent in asynchronous shared memory. In *ACM Symposium on Principles of Distributed Computing, PODC 2018*, 2018. (Cited on pages 1 and 2)
- Yossi Arjevani, Ohad Shamir, and Nathan Srebro. A tight convergence analysis for stochastic gradient descent with delayed updates. In *International Conference on Algorithmic Learning Theory (COLT 2020)*, 2020. (Cited on page 2)
- Rotem Zamir Aviv, Ido Hakimi, Assaf Schuster, and Kfir Yehuda Levy. Learning under delayed feedback: Implicitly adapting to gradient delays. In *International Conference on Machine Learning (ICML 2021)*, 2021. (Cited on page 2)
- Gerard Baudet. Asynchronous iterative methods for multiprocessors. *Journal of the Association for Computing Machinery*, 1978. (Cited on page 2)
- Ron Bekkerman, Mikhail Bilenko, and John Langford. *Scaling up machine learning: Parallel and distributed approaches*. Cambridge University Press, 2011. (Cited on page 1)
- Dimitri P. Bertsekas and John N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Prentice-Hall, 1989. (Cited on page 2)
- Sébastien Bubeck. *Convex optimization: Algorithms*

- and complexity*. Foundations and Trends in Machine Learning, 2015. (Cited on page 6)
- Chih-Chung Chang and Chih-Jen Lin. Libsvm: A library for support vector machines. *ACM transactions on intelligent systems and technology (TIST)*, 2011. (Cited on pages 8 and 15)
- Yae Jee Cho, Pranay Sharma, Gauri Joshi, Zheng Xu, Satyen Kale, and Tong Zhang. On the convergence of federated averaging with cyclic client participation. *arXiv preprint arXiv:2302.03109*, 2023. (Cited on page 3)
- Alon Cohen, Amit Daniely, Yoel Drori, Tomer Koren, and Mariano Schain. Asynchronous stochastic optimization robust to arbitrary delays. In *Advances in Neural Information Processing Systems (NeurIPS 2021)*, 2021. (Cited on page 2)
- Laurent Condat, Ivan Agarský, and Peter Richtárik. Provably doubly accelerated federated learning: The first theoretically successful combination of local training and communication compression. *arXiv preprint arXiv: 2210.13277*, 2023. (Cited on page 1)
- John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. In *Journal of Machine Learning Research*, 2011. (Cited on page 1)
- Mathieu Even, Anastasia Koloskova, and Laurent Massoulié. Asynchronous sgd on graphs: a unified framework for asynchronous decentralized and federated optimization. In *International Conference on Artificial Intelligence and Statistics (AISTATS 2024)*, 2024. (Cited on page 2)
- Yuan Gao, Rustem Islamov, and Sebastian Stich. Econtrol: Fast distributed optimization with compression and error control. In *International Conference on Learning Representations (ICLR 2024)*, 2024. (Cited on page 1)
- Margalit Glasgow and Mary Wootters. Asynchronous distributed optimization with stochastic delays. *arXiv preprint arXiv: 2009.10717*, 2021. (Cited on page 3)
- Eduard Gorbunov, Filip Hanzely, and Peter Richtárik. Local sgd: Unified theory and new efficient methods. In *International Conference on Artificial Intelligence and Statistics (AISTATS 2021)*, 2021. (Cited on page 1)
- Xinran Gu, Kaixuan Huang, Jingzhao Zhang, and Longbo Huang. Fast federated learning in the presence of arbitrary device unavailability. In *Advances in Neural Information Processing Systems (NeurIPS 2021)*, 2021. (Cited on page 3)
- Samuel Horvath, Stefanos Laskaridis, Mario Almeida, Ilias Leontiadis, Stylianos I. Venieris, and Nicholas D. Lane. Fjord: Fair and accurate federated learning under heterogeneous targets with ordered dropout. In *Advances in Neural Information Processing Systems (NeurIPS 2021)*, 2022. (Cited on page 2)
- Peter Kairouz et al. Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 2021. (Cited on page 1)
- Sarit Khirirat, Hamid Reza Feyzmahdavian, and Mikael Johansson. Distributed learning with compressed gradients. *arXiv preprint arXiv:1806.06573*, 2018. (Cited on page 1)
- Diederik P. Kingma and Jimmy Lei Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR 2015)*, 2015. (Cited on page 1)
- Anastasia Koloskova, Sebastian U. Stich, and Martin Jaggi. Decentralized stochastic optimization and gossip algorithms with compressed communication. In *International Conference on Machine Learning (ICML 2019)*, 2019. (Cited on page 1)
- 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 *International Conference on Machine Learning (ICML 2020)*, 2020. (Cited on page 1)
- Anastasia Koloskova, Sebastian Stich, and Martin Jaggi. Sharper convergence guarantees for asynchronous SGD for distributed and federated learning. In *Advances in Neural Information Processing Systems (NeurIPS 2022)*, 2022. (Cited on pages 2, 3, 4, 5, 6, 22, 32, 42, 46, and 47)
- Anastasia Koloskova, Nikita Doikov, Sebastian U. Stich, and Martin Jaggi. Shuffle sgd is always better than sgd: Improved analysis of sgd with arbitrary data orders. *arXiv preprint arXiv:2305.19259*, 2023a. (Cited on pages 3, 6, and 7)
- Anastasia Koloskova, Ryan McKenna, Zachary Charles, Keith Rush, and Brendan McMahan. Convergence of gradient descent with linearly correlated noise and applications to differentially private learning. *Advances in Neural Information Processing Systems (NeurIPS 2023)*, 2023b. (Cited on page 2)
- Jakub Konečný, H. Brendan McMahan, Felix Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. Federated learning: strategies for improving communication efficiency. In *NIPS Private Multi-Party Machine Learning Workshop*, 2016. (Cited on page 1)
- Dmitry Kovalev, Elnur Gasanov, Alexander Gasnikov, and Peter Richtárik. Lower bounds and optimal al-

- gorithms for smooth and strongly convex decentralized optimization over time-varying networks. In *Advances in Neural Information Processing Systems (NeurIPS 2021)*, 2021. (Cited on page 1)
- Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *arXiv preprint arXiv:1812.06127*, 2018. (Cited on page 8)
- Yucheng Lu, Si Yi Meng, and Christopher De Sa. A general analysis of example-selection for stochastic gradient descent. In *International Conference on Learning Representations (ICLR 2022)*, 2022. (Cited on page 2)
- Grigory Malinovsky, Konstantin Mishchenko, and Peter Richtárik. Server-side stepsizes and sampling without replacement provably help in federated optimization. *arXiv preprint arXiv:2201.11066*, 2022. (Cited on page 3)
- 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*, 2017. (Cited on pages 3 and 7)
- Konstantin Mishchenko, Eduard Gorbunov, Martin Takáč, and Peter Richtárik. Distributed learning with compressed gradient differences. *arXiv preprint arXiv:1901.09269*, 2019. (Cited on page 1)
- Konstantin Mishchenko, Ahmed Khaled, and Peter Richtarik. Random reshuffling: Simple analysis with vast improvements. In *Advances in Neural Information Processing Systems (NeurIPS 2022)*, 2020. (Cited on pages 2 and 3)
- Konstantin Mishchenko, Francis Bach, Mathieu Even, and Blake Woodworth. Asynchronous SGD beats minibatch SGD under arbitrary delays. In *Advances in Neural Information Processing Systems (NeurIPS 2022)*, 2022a. (Cited on pages 2, 3, 4, and 6)
- Konstantin Mishchenko, Grigory Malinovsky, Sebastian Stich, and Peter Richtárik. ProxSkip: Yes! local gradient steps provably lead to communication acceleration! finally! In *International Conference on Machine Learning (ICML 2022)*, 2022b. (Cited on page 1)
- Amirkeivan Mohtashami, Sebastian Stich, and Martin Jaggi. Characterizing and finding good data orderings for fast convergence of sequential gradient methods. *arXiv preprint arXiv: 2202.01838*, 2022. (Cited on page 2)
- Giorgi Nadiradze, Ilia Markov, Bapi Chatterjee, Vyacheslav Kungurtsev, and Dan Alistarh. Elastic consistency: A general consistency model for distributed stochastic gradient descent. In *AAAI Conference on Artificial Intelligence*, 2021. (Cited on page 2)
- Angela Nedic and Dimitri Bertsekas. Convergence rate of incremental subgradient algorithms. *Stochastic optimization: algorithms and applications*, 2001. (Cited on page 5)
- John Nguyen, Kshitiz Malik, Hongyuan Zhan, Ashkan Yousefpour, Mike Rabbat, Mani Malek, and Dzmitry Huba. Federated learning with buffered asynchronous aggregation. In *International Conference on Artificial Intelligence and Statistics (AISTATS 2022)*, 2022. (Cited on pages 2, 3, 4, 5, 8, and 44)
- Lam M. Nguyen, Quoc Tran-Dinh, Dzung T. Phan, Phuong Ha Nguyen, and Marten van Dijk. A unified convergence analysis for shuffling-type gradient methods. *Journal of Machine Learning Research*, 2021. (Cited on pages 2 and 3)
- Alexander Rakhlin, Ohad Shamir, and Karthik Sridharan. Making gradient descent optimal for strongly convex stochastic optimization. *International Conference on International Conference on Machine Learning (ICML 2012)*, 2012. (Cited on page 3)
- Herbert Robbins and Sutton Monro. A stochastic approximation method. In *The Annals of Mathematical Statistics*, 1951. (Cited on page 1)
- Abdurakhmon Sadiev, Grigory Malinovsky, Eduard Gorbunov, Igor Sokolov, Ahmed Khaled, Konstantin Burlachenko, and Peter Richtárik. Federated optimization algorithms with random reshuffling and gradient compression. *arXiv preprint arXiv:2206.07021*, 2022. (Cited on page 3)
- Mher Safaryan, Rustem Islamov, Xun Qian, and Peter Richtarik. FedNL: Making Newton-type methods applicable to federated learning. In *International Conference on Machine Learning (ICML 2022)*, 2022. (Cited on page 8)
- Sebastian Stich, Amirkeivan Mohtashami, and Martin Jaggi. Critical parameters for scalable distributed learning with large batches and asynchronous updates. In *International Conference on Artificial Intelligence and Statistics (AISTATS 2021)*, 2021. (Cited on page 2)
- Sebastian U. Stich and Sai Praneeth Karimireddy. The error-feedback framework: Better rates for sgd with delayed gradients and compressed communication. *Journal of Machine Learning Research*, 2020. (Cited on page 2)
- Alexander Tyurin and Peter Richtarik. Optimal time complexities of parallel stochastic optimization

methods under a fixed computation model. *arXiv preprint arXiv: 2305.12387*, 2023. (Cited on page 2)

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. *arXiv preprint arXiv: 2002.07399*, 2020. (Cited on page 3)

Chulhee Yun, Suvrit Sra, and Ali Jadbabaie. Can single-shuffle sgd be better than reshuffling sgd and gd? *arXiv preprint arXiv:2103.07079*, 2021. (Cited on page 3)

Hossein Zakerinia, Shayan Talaei, Giorgi Nadiradze, and Dan Alistarh. Communication-efficient federated learning with data and client heterogeneity. *arXiv preprint arXiv: 2206.10032*, 2022. (Cited on pages 1 and 2)

Wei Zhang, Suyog Gupta, Xiangru Lian, and Ji Liu. Staleness-aware async-sgd for distributed deep learning. In *International Joint Conference on Artificial Intelligence (IJCAI 2016)*, 2016. (Cited on page 3)

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 *International Conference on Machine Learning (ICML 2017)*, 2017. (Cited on page 3)

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. **[Yes]**
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. **[Yes]**
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. **[Yes]**
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. **[Yes]**
 - (b) Complete proofs of all theoretical results. **[Yes]**
 - (c) Clear explanations of any assumptions. **[Yes]**
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). **[Yes]**
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). **[Yes]**
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). **[Yes]**
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). **[Yes]**
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. **[Yes]**
 - (b) The license information of the assets, if applicable. **[Not Applicable]**
 - (c) New assets either in the supplemental material or as a URL, if applicable. **[Not Applicable]**
 - (d) Information about consent from data providers/curators. **[Not Applicable]**
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. **[Not Applicable]**
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots. **[Not Applicable]**
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. **[Not Applicable]**
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. **[Not Applicable]**

Contents

| | | |
|----------|--|-----------|
| 1 | INTRODUCTION | 1 |
| 1.1 | Contributions | 2 |
| 2 | RELATED WORKS | 2 |
| 2.1 | SGD with Arbitrary Data Ordering | 3 |
| 3 | SETUP | 3 |
| 3.1 | Proposed Algorithmic Framework | 3 |
| 3.2 | Special Cases | 4 |
| 4 | CONVERGENCE THEORY | 5 |
| 4.1 | Theoretical assumptions | 5 |
| 4.2 | Notation | 6 |
| 4.3 | Analysis of Gradient Receiving Process | 7 |
| 4.4 | Analysis of Gradient Assigning Process | 7 |
| 5 | EXPERIMENTS | 8 |
| 6 | CONCLUSION | 9 |
| 7 | Acknowledgments | 9 |
| A | ADDITIONAL EXPERIMENTS | 15 |
| A.1 | Synthetic Dataset Generation | 15 |
| A.2 | Comparison on phishing Dataset | 15 |
| A.3 | Comparison on Synthetic Dataset | 16 |
| B | USEFUL LEMMAS | 16 |
| C | PROOFS FOR ANALYSIS OF GRADIENT RECEIVING PROCESS | 18 |
| C.1 | Key Lemmas | 18 |
| C.2 | Proof of Theorem 1 | 22 |
| C.3 | Convergence Guarantees in Special Cases | 26 |
| C.3.1 | Pure Asynchronous SGD | 26 |
| C.3.2 | Mini-batch SGD | 28 |
| C.3.3 | Pure Asynchronous SGD with waiting | 30 |
| C.3.4 | SGD with Random Reshuffling | 31 |
| D | PROOFS FOR ANALYSIS OF GRADIENT ASSIGNING PROCESS | 32 |
| D.1 | Key Lemmas | 33 |

| | | |
|-------|--|----|
| D.2 | Proof of Theorem 2 | 39 |
| D.3 | Convergence Guarantees in Special Cases | 42 |
| D.3.1 | Random Asynchronous SGD | 42 |
| D.3.2 | Random asynchronous SGD with waiting | 44 |
| D.3.3 | Shuffled Asynchronous SGD [NEW] | 45 |
| D.3.4 | Pure Asynchronous SGD | 46 |
| D.3.5 | Pure Asynchronous SGD with waiting | 47 |

A ADDITIONAL EXPERIMENTS




Figure 3: Comparison of pure, random, and shuffled asynchronous SGD with tuned stepsizes and full gradient computation on phishing dataset with various delay patterns. Here $d = 68$, $m = 1105$.

In this section, we provide additional numerical results as well as the detailed experiment setup. We consider logistic regression problem with nonconvex regularization, namely,

$$f(x) = \frac{1}{n} \sum_{i=1}^n f_i(x), \quad \text{where} \quad f_i(x) = \frac{1}{m} \sum_{j=1}^m \log \left(1 + e^{-b_{ij} a_{ij}^\top x} \right) + \lambda \sum_{j=1}^d \frac{x_j^2}{1 + x_j^2}.$$

For synthetic dataset, we set $n = 10$, $m = 200$, $d = 300$, and $\lambda = 0.1$ in all experiments. The initial point $x_0 \in \mathbb{R}^d$ is sampled from standard Gaussian distribution.

We compare the practical performance of FedBuff, pure asynchronous SGD, random asynchronous SGD, and shuffled asynchronous SGD. We assume that all workers are active from the beginning which implies that $\tau_C = n$. We perform a grid search for the stepsize over the set $\{0.005, 0.004, 0.003, 0.002, 0.001, 0.0005, 0.0001\}$, and illustrate the performance with the fine-tuned stepsize only choosing the best one that achieves smaller gradient norm and fluctuations. For FedBuff we perform 20 local steps and wait for one worker only. We set $n = 10$, $m = 200$, $d = 300$, and $\lambda = 0.1$ in all experiments. The initial point $x_0 \in \mathbb{R}^d$ is sampled from standard Gaussian distribution.

A.1 Synthetic Dataset Generation

The generation function of datasets has two parameters α and β . Below we detail the rules of dataset generation:

1. Sample n numbers $B_i \sim \mathcal{N}(0, \beta)$.
2. Sample n vectors $v_i \in \mathbb{R}^d$ such that each component $[v_i]_j \sim \mathcal{N}(B_i, 1)$.
3. For all $i \in [n]$ sample m vectors $a_{ij} \in \mathbb{R}^d$ such that $a_{ij} \sim \mathcal{N}(v_i, \Sigma)$ where Σ is a diagonal matrix with $\Sigma_{jj} = j^{-1.2}$. A set $\{a_{ij}\}_{j=1}^m$ will be utilized feature vectors of i -th worker.
4. Sample n pairs of numbers $u_i \sim \mathcal{N}(0, \alpha)$ and $c_i \sim \mathcal{N}(u_i, 1)$.
5. Sample n vectors $w_i \in \mathbb{R}^d$ such that each component $[w_i]_j \sim \mathcal{N}(u_i, 1)$.
6. Let $p_{ij} = \sigma(w_i^\top a_{ij} + c_i)$ where $\sigma(\cdot)$ is a sigmoid function.
7. For all $i \in [n]$ sample m numbers b_{ij} such that b_{ij} equals to -1 with probability p_{ij} and to 1 otherwise.

The larger α and β are, the more heterogeneous local datasets are. In our experiments, we choose these parameters from the set $\{0.5, 1, 1.5\}$.

A.2 Comparison on phishing Dataset

We consider the same as in Figure 1, but we test the performance of algorithms on phishing dataset [Chang and Lin \(2011\)](#); see Figure 3. We observe that in all cases random and shuffled asynchronous SGD outperform their pure counterpart as asynchronous SGD with randomized job assigning finds better stationary points. Besides, shuffled asynchronous SGD finds a more accurate solution than random asynchronous which happens because of more accurate job assigning.

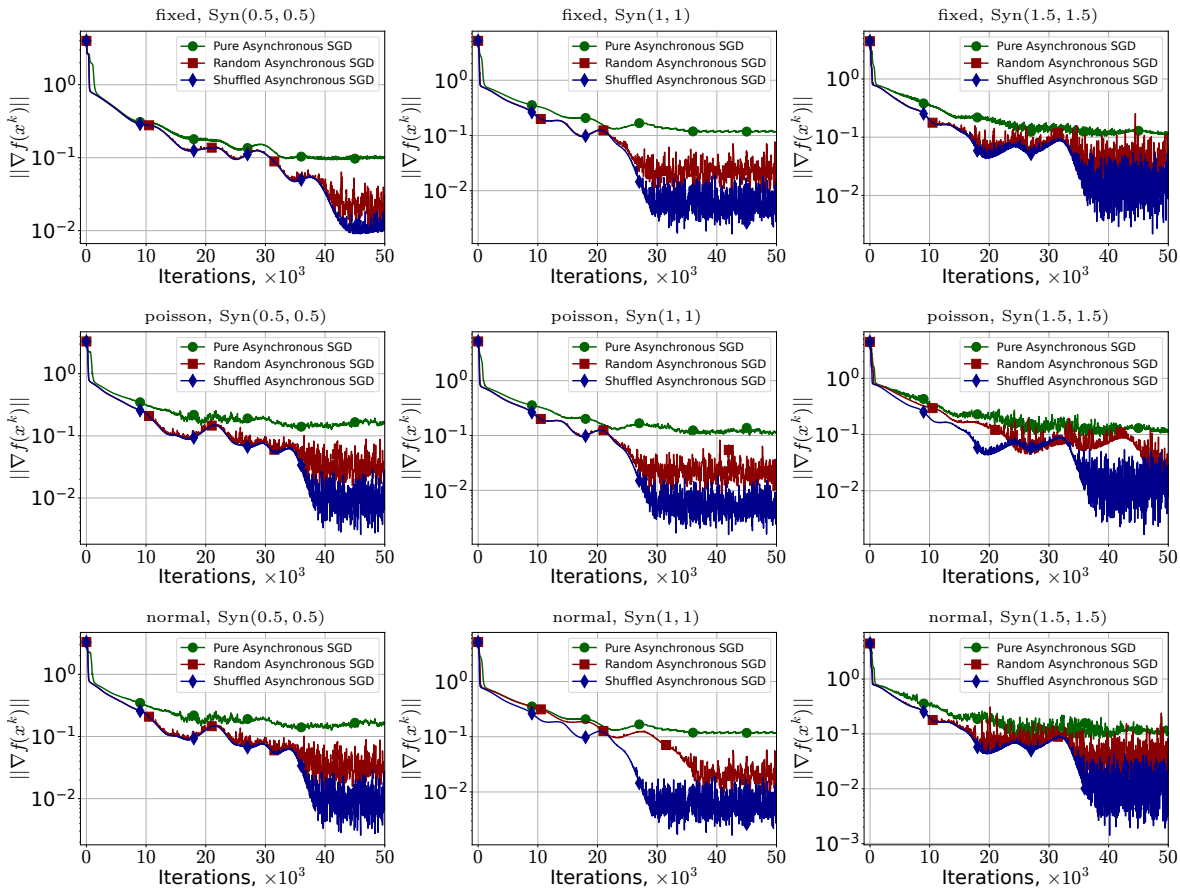


Figure 4: Comparison of pure, random, and shuffled asynchronous SGD with tuned stepsizes and full gradient computation on synthetic datasets with different levels of heterogeneity and various delay patterns.

A.3 Comparison on Synthetic Dataset

First, we test the aforementioned algorithms in the case of full local gradient computation in order to avoid the effect of stochasticity and illustrate the impact of delay patterns only; see Figure 4. We change the heterogeneity level controlling parameters α and β , and delay patterns. We observe that in all cases pure asynchronous SGD gets stuck much earlier than its counterparts regardless of the delay pattern. However, due to that, we can choose a larger stepsize which leads to faster convergence.

Besides, we notice that the shuffled asynchronous SGD method finds a better stationary point than the random version which happens because of the more careful job assigning. Moreover, the shuffled version sometimes can be even faster (e.g., normal Syn(1,1) and poison Syn(1.5, 1.5)).

Next, we conduct experiments in the same setup, but with stochastic gradients where the mini-batch size is 10% of the local dataset. Figure 5 shows the same performance as in the full gradient case: shuffled asynchronous SGD finds a more accurate stationary point than its counterparts.

These numerical results support our theoretical findings. First, pure asynchronous SGD does not converge to the optimum even when choosing small stepsize—it always gets stuck at the heterogeneity level. Moreover, random job assigning allows to get rid of this issue, and converge to a better stationary point. Finally, more accurate job assigning via worker shuffling allows for improving further the performance

B USEFUL LEMMAS

We frequently use several well-known results.

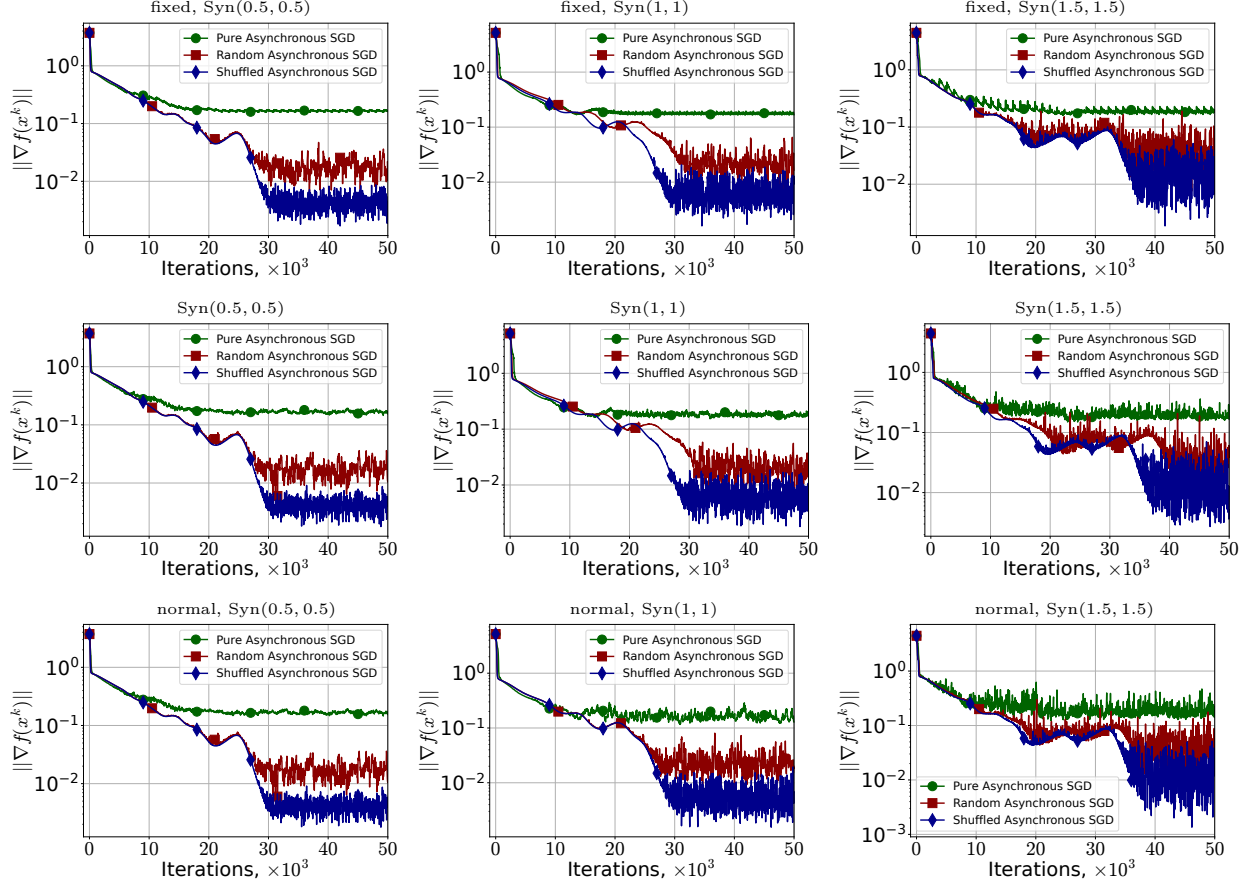


Figure 5: Comparison of pure, random, and shuffled asynchronous SGD with tuned stepsizes and stochastic gradient computation on synthetic datasets with different levels of heterogeneity and various delay pattern. Here $d = 300$, $m = 200$.

Lemma B.1. Let $\{v_i\}_{i=1}^m$ be a finite set of vectors in \mathbb{R}^d . Then we have

$$\left\| \sum_{i=1}^m v_i \right\|^2 \leq m \sum_{i=1}^m \|v_i\|^2. \quad (16)$$

Lemma B.2. Let a, b be any two vectors in \mathbb{R}^d , and α be any positive number. Then we have

$$\langle a, b \rangle \leq \frac{\alpha}{2} \|a\|^2 + \frac{1}{2\alpha} \|b\|^2. \quad (17)$$

The next lemma is useful when we work we double sums in the proofs. Let $\beta(t, \pi_t)$ be any non-negative function with two iterate-dependent arguments t and π_t . For example, $\beta(t, \pi_t) = \|\nabla f_{t_t}(x_{\pi_t})\|^2$ or $\beta(t, \pi_t) = \|x_t - x_{\pi_t}\|^2$.

Lemma B.3. Let $\beta(t, \pi_t)$ be defined as above. If τ_C is defined in (2), then we have

$$\sum_{t=0}^{T-1} \sum_{j=\pi_t}^{t-1} \beta(j, \pi_j) \leq (\tau_C - 1) \sum_{t=0}^{T-1} \beta(t, \pi_t). \quad (18)$$

Proof. Each term $\beta(t, \pi_t)$ appears for all t' such that $t \in [\pi_{t'}, t')$. We need to understand how many such t' might exist at most. Indeed, such term might appear only for t' before t , thus, that job should have been started somewhere in the past but still not been applied. The number of such it is bounded by $|\mathcal{A}_t \setminus \mathcal{R}_t| = |\mathcal{A}_{t+1} \setminus \mathcal{R}_t| - 1 \leq \tau_C - 1$. \square

C PROOFS FOR ANALYSIS OF GRADIENT RECEIVING PROCESS

For $r(t) \leq m < r(t) + \tau$, denote

$$\phi_t^m(x) := \mathbb{E} \left[\left\| \sum_{j=r(t)}^m (\nabla f(x) - \nabla f_{i_j}(x)) \right\|^2 \right].$$

Moreover, for shortness we use the following notation

$$\begin{aligned} A &:= \sum_{t=0}^{T-1} \mathbb{E} [\|x_t - x_{\pi_t}\|^2], \quad B := \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(x_t)\|^2], \quad \Phi := \sum_{t=0}^{T-1} \phi_t^{t-1}(x_{r(t)}) \\ \Psi &:= \sum_{t=0}^{T-1} \mathbb{E} \left[\left\| \sum_{j=\pi_t}^{t-1} \nabla f_{i_j}(x_{\pi_j}) - \nabla f(x_{\pi_j}) \right\|^2 \right]. \end{aligned}$$

The key part of the proof is bounding the distance between virtual and real iterates. Based on (2) and (11), we have

$$x_t = x_{r(t)} - \gamma \sum_{j=r(t)}^{t-1} g_{i_j}(x_{\pi_j}), \quad \tilde{x}_t = x_{r(t)} - \gamma \sum_{j=r(t)}^{t-1} \nabla f(x_j).$$

Also, denote

$$\Delta_t^m := \sum_{j=r(t)}^m (\nabla f(x_j) - g_{i_j}(x_{\pi_j})). \quad (19)$$

C.1 Key Lemmas

Lemma C.1. If $20\gamma L\tau \leq 1$, then

$$\begin{aligned} \mathbb{E} [\|\Delta_t^m\|^2] &\leq 4\mathbb{E} [\phi_t^m(x_{r(t)})] + \frac{1}{6\tau} \sum_{j=r(t)}^m \mathbb{E} [\phi_t^{j-1}(x_{r(t)})] + \frac{25}{6} L^2 \tau \sum_{j=r(t)}^m \mathbb{E} [\|x_j - x_{\pi_j}\|^2] \\ &\quad + \frac{\tau}{24} \sum_{j=r(t)}^m \mathbb{E} [\|\nabla f(x_j)\|^2] + \frac{\tau}{24} \sigma^2. \end{aligned}$$

Proof. Fix any m such that $r(t) \leq m < r(t) + \tau$, so that the iterates are within one block.

$$\begin{aligned} \mathbb{E} [\|\Delta_t^m\|^2] &\leq \mathbb{E} \left[\left\| \sum_{j=r(t)}^m \nabla f_{i_j}(x_{\pi_j}) - \nabla f(x_j) \right\|^2 \right] + \tau \sigma^2 \\ &\leq 4\mathbb{E} \left[\left\| \sum_{j=r(t)}^m \nabla f_{i_j}(x_{r(t)}) - \nabla f(x_{r(t)}) \right\|^2 \right] + 4\mathbb{E} \left[\left\| \sum_{j=r(t)}^m \nabla f_{i_j}(x_{r(t)}) - \nabla f_{i_j}(x_j) \right\|^2 \right] \\ &\quad + 4\mathbb{E} \left[\left\| \sum_{j=r(t)}^m \nabla f_{i_j}(x_j) - \nabla f_{i_j}(x_{\pi_j}) \right\|^2 \right] + 4\mathbb{E} \left[\left\| \sum_{j=r(t)}^m \nabla f(x_{r(t)}) - \nabla f(x_j) \right\|^2 \right] + \tau \sigma^2 \\ &\leq 4\mathbb{E} [\phi_t^m(x_{r(t)})] + 4L^2\tau \sum_{j=r(t)}^m \mathbb{E} [\|x_j - x_{\pi_j}\|^2] + 8L^2\tau \sum_{j=r(t)}^m \mathbb{E} [\|x_j - x_{r(t)}\|^2] + \tau \sigma^2. \end{aligned} \quad (20)$$

Next, we bound the third term involving $\mathbb{E} [\|x_j - x_{r(t)}\|^2]$.

$$\begin{aligned}
 & \sum_{j=r(t)}^m \mathbb{E} [\|x_j - x_{r(t)}\|^2] = \gamma^2 \sum_{j=r(t)}^m \mathbb{E} \left[\left\| \sum_{l=r(t)}^{j-1} g_{i_l}(x_{\pi_l}) \right\|^2 \right] \\
 & \leq 2\gamma^2 \sum_{j=r(t)}^m \mathbb{E} \left[\left\| \sum_{l=r(t)}^{j-1} g_{i_l}(x_{\pi_l}) - \nabla f(x_l) \right\|^2 \right] + 2\gamma^2 \sum_{j=r(t)}^m \mathbb{E} \left[\left\| \sum_{l=r(t)}^{j-1} \nabla f(x_l) \right\|^2 \right] \\
 & = 2\gamma^2 \sum_{j=r(t)}^m \mathbb{E} \left[\|\Delta_t^{j-1}\|^2 \right] + 2\gamma^2 \sum_{j=r(t)}^m \mathbb{E} \left[\left\| \sum_{l=r(t)}^{j-1} \nabla f(x_l) \right\|^2 \right] \\
 & \stackrel{(20)}{\leq} 8\gamma^2 \sum_{j=r(t)}^m \mathbb{E} \left[\phi_t^{j-1}(x_{r(t)}) \right] + 16\gamma^2 L^2 \tau \sum_{j=r(t)}^m \sum_{l=r(t)}^{j-1} \mathbb{E} [\|x_l - x_{r(t)}\|^2] \\
 & \quad + 8\gamma^2 L^2 \tau \sum_{j=r(t)}^m \sum_{l=r(t)}^{j-1} \mathbb{E} [\|x_l - x_{\pi_l}\|^2] + 2\gamma^2 \tau \sum_{j=r(t)}^m \sum_{l=r(t)}^{j-1} \mathbb{E} [\|\nabla f(x_l)\|^2] + 2\gamma^2 \tau^2 \sigma^2 \\
 & \leq 8\gamma^2 \sum_{j=r(t)}^m \mathbb{E} \left[\phi_t^{j-1}(x_{r(t)}) \right] + 16\gamma^2 L^2 \tau^2 \sum_{j=r(t)}^m \mathbb{E} [\|x_j - x_{r(t)}\|^2] + 8\gamma^2 L^2 \tau^2 \sum_{j=r(t)}^m \mathbb{E} [\|x_j - x_{\pi_j}\|^2] \\
 & \quad + 2\gamma^2 \tau^2 \sum_{j=r(t)}^m \mathbb{E} [\|\nabla f(x_j)\|^2] + 2\gamma^2 \tau^2 \sigma^2,
 \end{aligned}$$

where we use the fact that $\tau_C \geq 1$. Choosing $\gamma L \tau \leq \frac{1}{20}$, we cancel the term $\|x_j - x_{r(t)}\|^2$ from the right hand side:

$$\begin{aligned}
 \sum_{j=r(t)}^m \mathbb{E} [\|x_j - x_{r(t)}\|^2] & \leq \frac{25}{3} \gamma^2 \sum_{j=r(t)}^m \mathbb{E} \left[\phi_t^{j-1}(x_{r(t)}) \right] + \frac{25}{3} \gamma^2 L^2 \tau^2 \sum_{j=r(t)}^m \mathbb{E} [\|x_j - x_{\pi_j}\|^2] \\
 & \quad + \frac{25}{12} \gamma^2 \tau^2 \sum_{j=r(t)}^m \mathbb{E} [\|\nabla f(x_j)\|^2] + \frac{25}{12} \gamma^2 \tau^2 \sigma^2.
 \end{aligned}$$

Plugging this inequality back to (20), we get

$$\begin{aligned}
 \mathbb{E} [\|\Delta_t^m\|^2] & \leq 4\mathbb{E} [\phi_t^m(x_{r(t)})] + 4L^2 \tau \sum_{j=r(t)}^m \mathbb{E} [\|x_j - x_{\pi_j}\|^2] + 8L^2 \tau \sum_{j=r(t)}^m \mathbb{E} [\|x_j - x_{r(t)}\|^2] \\
 & \leq 4\mathbb{E} [\phi_t^m(x_{r(t)})] + \frac{200}{3} \gamma^2 L^2 \tau \sum_{j=r(t)}^m \mathbb{E} \left[\phi_t^{j-1}(x_{r(t)}) \right] \\
 & \quad + \frac{25}{6} L^2 \tau \sum_{j=r(t)}^m \mathbb{E} [\|x_j - x_{\pi_j}\|^2] + \frac{50}{3} \gamma^2 L^2 \tau^3 \sum_{j=r(t)}^m \mathbb{E} [\|\nabla f(x_j)\|^2] + \frac{50}{3} \gamma^2 L^2 \tau^3 \sigma^2.
 \end{aligned}$$

Using the relation $\gamma^2 L^2 \tau^2 \leq \frac{1}{400}$ further, we can simplify the above bound as

$$\begin{aligned}
 \mathbb{E} [\|\Delta_t^m\|^2] & \leq 4\mathbb{E} [\phi_t^m(x_{r(t)})] + \frac{1}{6\tau} \sum_{j=r(t)}^m \mathbb{E} \left[\phi_t^{j-1}(x_{r(t)}) \right] + \frac{25}{6} L^2 \tau \sum_{j=r(t)}^m \mathbb{E} [\|x_j - x_{\pi_j}\|^2] \\
 & \quad + \frac{\tau}{24} \sum_{j=r(t)}^m \mathbb{E} [\|\nabla f(x_j)\|^2] + \frac{\tau}{24} \sigma^2.
 \end{aligned}$$

Since $\frac{1}{16} > \frac{1}{24}$ we conclude the proof. \square

Lemma C.2. If $20\gamma L\tau \leq 1$, then

$$\begin{aligned}
 \sum_{t=0}^{T-1} \mathbb{E} [\|x_t - \tilde{x}_t\|^2] &\leq \underbrace{\frac{25}{6}\gamma^2 \sum_{t=0}^{T-1} \mathbb{E} [\phi_t^{t-1}(x_{r(t)})]}_{:=\Phi} + \underbrace{\frac{1}{96} \sum_{t=0}^{T-1} \mathbb{E} [\|x_t - x_{\pi_t}\|^2]}_{:=A} \\
 &\quad + \underbrace{\frac{1}{9600L^2} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(x_t)\|^2]}_{:=B} + \frac{\gamma^2\tau}{24}\sigma^2T. \tag{21}
 \end{aligned}$$

Proof. Notice that for $r(t) < j \leq t$ it holds $r(t) = r(j)$. Hence, we can replace $\phi_t^{j-1}(x_{r(t)})$ by $\phi_j^{j-1}(x_{r(j)})$ in the summation below. If $j \leq r(t)$, then $\phi_t^{j-1}(x_{r(t)}) = 0$.

$$\begin{aligned}
 \sum_{t=0}^{T-1} \mathbb{E} [\|x_t - \tilde{x}_t\|^2] &= \gamma^2 \sum_{t=0}^{T-1} \mathbb{E} \left[\left\| \sum_{j=r(t)}^{t-1} g_{i_j}(x_{\pi_j}) - \nabla f(x_j) \right\|^2 \right] \\
 &= \gamma^2 \sum_{t=0}^{T-1} \mathbb{E} [\|\Delta_t^{t-1}\|^2] \\
 &\stackrel{\text{Lemma C.1}}{\leq} 4\gamma^2 \sum_{t=0}^{T-1} \mathbb{E} [\phi_t^{t-1}(x_{r(t)})] + \frac{\gamma^2}{6\tau} \sum_{t=0}^{T-1} \sum_{j=r(t)}^{t-1} \mathbb{E} [\phi_t^{j-1}(x_{r(t)})] \\
 &\quad + \frac{1}{96\tau} \sum_{t=0}^{T-1} \sum_{j=r(t)}^{t-1} \mathbb{E} [\|x_j - x_{\pi_j}\|^2] + \frac{\gamma^2\tau}{24} \sum_{t=0}^{T-1} \sum_{j=r(t)}^{t-1} \mathbb{E} [\|\nabla f(x_j)\|^2] \\
 &\quad + \frac{\gamma^2\tau}{24}\sigma^2T \\
 &\leq \frac{25}{6}\gamma^2 \sum_{t=0}^{T-1} \mathbb{E} [\phi_t^{t-1}(x_{r(t)})] + \frac{1}{96} \sum_{t=0}^{T-1} \mathbb{E} [\|x_t - x_{\pi_t}\|^2] \\
 &\quad + \frac{\gamma^2\tau^2}{24} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(x_t)\|^2] + \frac{\gamma^2\tau}{24}\sigma^2T.
 \end{aligned}$$

Using the bound on the learning rate we conclude the proof. \square

Lemma C.3. If $20\gamma L\sqrt{\tau_{\max}T_C} \leq 1$, then

$$\begin{aligned}
 \underbrace{\sum_{t=0}^{T-1} \mathbb{E} [\|x_t - x_{\pi_t}\|^2]}_{:=A} &\leq \underbrace{\frac{1}{132L^2} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(x_t)\|^2]}_{:=B} + \frac{\gamma}{5L}T\sigma^2 \\
 &\quad + \underbrace{\frac{100\gamma^2}{33} \sum_{t=0}^{T-1} \mathbb{E} \left[\left\| \sum_{j=\pi_t}^{t-1} \nabla f_{i_j}(x_{\pi_j}) - \nabla f(x_{\pi_j}) \right\|^2 \right]}_{:=\Psi}.
 \end{aligned}$$

Proof. Recall that $\pi_t = t - \tau_t$.

$$\begin{aligned}
 \mathbb{E} [\|x_t - x_{\pi_t}\|^2] &= \gamma^2 \mathbb{E} \left[\left\| \sum_{j=\pi_t}^{t-1} g_{i_j}(x_{\pi_j}) \right\|^2 \right] \\
 &\leq \gamma^2 \mathbb{E} \left[\left\| \sum_{j=\pi_t}^{t-1} \nabla f_{i_j}(x_{\pi_j}) \right\|^2 \right] + \tau_t \gamma^2 \sigma^2 \\
 &\leq 3\gamma^2 \mathbb{E} \left[\left\| \sum_{j=\pi_t}^{t-1} \nabla f_{i_j}(x_{\pi_j}) - \nabla f(x_{\pi_j}) \right\|^2 \right] + 3\gamma^2 \mathbb{E} \left[\left\| \sum_{j=\pi_t}^{t-1} \nabla f(x_{\pi_j}) - \nabla f(x_j) \right\|^2 \right] \\
 &\quad + 3\gamma^2 \mathbb{E} \left[\left\| \sum_{j=\pi_t}^{t-1} \nabla f(x_j) \right\|^2 \right] + \tau_t \gamma^2 \sigma^2 \\
 &\leq 3\gamma^2 L^2 \tau_t \sum_{j=\pi_t}^{t-1} \mathbb{E} [\|x_{\pi_j} - x_j\|^2] + 3\gamma^2 \mathbb{E} \left[\left\| \sum_{j=\pi_t}^{t-1} \nabla f_{i_j}(x_{\pi_j}) - \nabla f(x_{\pi_j}) \right\|^2 \right] \\
 &\quad + 3\gamma^2 \tau_t \sum_{j=\pi_t}^{t-1} \|\nabla f(x_j)\|^2 + \tau_t \gamma^2 \sigma^2.
 \end{aligned}$$

Then we add summation over the entire iterates and count the number of times (which is $\leq \tau_C$) each term appears in the sum:

$$\begin{aligned}
 \sum_{t=0}^{T-1} \mathbb{E} [\|x_{\pi_t} - x_t\|^2] &\leq 3\gamma^2 L^2 \tau_{\max} \sum_{t=0}^{T-1} \sum_{j=\pi_t}^{t-1} \mathbb{E} [\|x_{\pi_j} - x_j\|^2] + 3\gamma^2 \tau_{\max} \sum_{t=0}^{T-1} \sum_{j=\pi_t}^{t-1} \mathbb{E} [\|\nabla f(x_j)\|^2] \\
 &\quad + 3\gamma^2 \sum_{t=0}^{T-1} \mathbb{E} \left[\left\| \sum_{j=\pi_t}^{t-1} \nabla f_{i_j}(x_{\pi_j}) - \nabla f(x_{\pi_j}) \right\|^2 \right] + \tau_{\text{avg}} T \gamma^2 \sigma^2 \\
 &\leq 3\gamma^2 L^2 \tau_{\max} \tau_C \sum_{t=0}^{T-1} \mathbb{E} [\|x_{\pi_t} - x_t\|^2] + 3\gamma^2 \tau_{\max} \tau_C \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(x_t)\|^2] \\
 &\quad + 3\gamma^2 \sum_{t=0}^{T-1} \mathbb{E} \left[\left\| \sum_{j=\pi_t}^{t-1} \nabla f_{i_j}(x_j) - \nabla f(x_j) \right\|^2 \right] + \tau_{\text{avg}} T \gamma^2 \sigma^2 \\
 &\leq \frac{3}{400} \sum_{t=0}^{T-1} \mathbb{E} [\|x_{\pi_t} - x_t\|^2] + \frac{3}{400L^2} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(x_t)\|^2] \\
 &\quad + 3\gamma^2 \sum_{t=0}^{T-1} \mathbb{E} \left[\left\| \sum_{j=\pi_t}^{t-1} \nabla f_{i_j}(x_{\pi_j}) - \nabla f(x_{\pi_j}) \right\|^2 \right] + \tau_{\text{avg}} T \gamma^2 \sigma^2,
 \end{aligned}$$

provided that $\gamma^2 L^2 \tau_{\max} \tau_C \leq \frac{1}{400}$. After cancellation, we get

$$\begin{aligned}
 \sum_{t=0}^{T-1} \mathbb{E} [\|x_t - x_{\pi_t}\|^2] &\leq \frac{1}{132L^2} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(x_t)\|^2] + 2\tau_{\text{avg}} T \gamma^2 \sigma^2 \\
 &\quad + \frac{100\gamma^2}{33} \sum_{t=0}^{T-1} \mathbb{E} \left[\left\| \sum_{j=\pi_t}^{t-1} \nabla f_{i_j}(x_{\pi_j}) - \nabla f(x_{\pi_j}) \right\|^2 \right].
 \end{aligned}$$

Since $\tau_{\text{avg}} \leq \tau_{\text{max}}$ and $\tau_{\text{avg}} \leq 2\tau_C$ (the latter holds because of the Remark 5 in [Koloskova et al. \(2022\)](#)), then we finally derive

$$\begin{aligned}
 \sum_{t=0}^{T-1} \mathbb{E} [\|x_t - x_{\pi_t}\|^2] &\leq \frac{1}{132L^2} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(x_t)\|^2] + \frac{2\tau_{\text{avg}}}{20L\sqrt{\tau_{\text{max}}\tau_C}} T\gamma\sigma^2 \\
 &\quad + \frac{100\gamma^2}{33} \sum_{t=0}^{T-1} \mathbb{E} \left[\left\| \sum_{j=\pi_t}^{t-1} \nabla f_{i_j}(x_{\pi_j}) - \nabla f(x_{\pi_j}) \right\|^2 \right] \\
 &\leq \frac{1}{132L^2} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(x_t)\|^2] + \frac{\gamma}{5L} T\sigma^2 \\
 &\quad + \frac{100\gamma^2}{33} \sum_{t=0}^{T-1} \mathbb{E} \left[\left\| \sum_{j=\pi_t}^{t-1} \nabla f_{i_j}(x_{\pi_j}) - \nabla f(x_{\pi_j}) \right\|^2 \right].
 \end{aligned}$$

□

We can combine all the previous lemmas and bound $\mathbb{E} [\|x_t - \tilde{x}_t\|^2]$ in the following lemma:

Lemma C.4. If $20\gamma L\tau \leq 1$ and $20\gamma L\sqrt{\tau_{\text{max}}\tau_C} \leq 1$ hold, then

$$\sum_{t=0}^{T-1} \mathbb{E} [\|x_t - \tilde{x}_t\|^2] \leq 5\gamma^2\Phi + \gamma^2\Psi + \frac{1}{5460L^2}B + \frac{1}{240L}\gamma\sigma^2T. \quad (22)$$

Proof. Summary of obtained inequalities

$$\begin{aligned}
 \sum_{t=0}^{T-1} \mathbb{E} [\|x_t - \tilde{x}_t\|^2] &\stackrel{\text{Lemma C.2}}{\leq} \frac{25}{6}\gamma^2\Phi + \frac{1}{96}A + \frac{1}{9600L^2}B + \frac{\gamma^2\tau}{24}\sigma^2T \\
 A &\stackrel{\text{Lemma C.3}}{\leq} \frac{1}{132L^2}B + \frac{100}{33}\gamma^2\Psi + \frac{\gamma}{5L}T\sigma^2.
 \end{aligned}$$

Hence, using the restriction $20\gamma L\tau \leq 1$ we get

$$\begin{aligned}
 \sum_{t=0}^{T-1} \mathbb{E} [\|x_t - \tilde{x}_t\|^2] &\leq \frac{25}{6}\gamma^2\Phi + \frac{1}{96}A + \frac{1}{9600L^2}B + \frac{\gamma}{480L}\sigma^2T \\
 &\leq \frac{25}{6}\gamma^2\Phi + \frac{1}{96 \cdot 132L^2}B + \frac{1}{9600L^2}B + \frac{1}{30}\gamma^2\Psi + \frac{11\gamma T\sigma^2}{100L} \\
 &= 5\gamma^2\Phi + \frac{1}{5460L^2}B + \gamma^2\Psi + \frac{\gamma}{240L}T\sigma^2.
 \end{aligned}$$

□

C.2 Proof of Theorem 1

Theorem 1. Let Assumptions 1 and 2 hold. Let the stepsize γ satisfy inequalities $6L\gamma \leq 1$ and $20L\gamma\sqrt{\tau_{\text{max}}\tau_C} \leq 1$, the correlation period $\tau = \left\lfloor \frac{1}{20L\gamma} \right\rfloor$, and quantities $\{\sigma_{k,\tau}^2\}_{k=0}^{\lfloor T/\tau \rfloor}$ and ν^2 are finite. Then

$$\mathbb{E} [\|\nabla f(\hat{x}_T)\|^2] \leq \mathcal{O} \left(\frac{F_0}{\gamma T} + L\gamma\sigma^2 + L^2\gamma^2\Phi \right), \quad (12)$$

where $F_0 := f(x_0) - f^*$, $\Phi := \frac{1}{\lfloor T/\tau \rfloor} \sum_{k=0}^{\lfloor T/\tau \rfloor} \sigma_{k,\tau}^2 + \frac{1}{T}\nu^2$, and \hat{x}_T is chosen uniformly at random from $\{x_0, \dots, x_{T-1}\}$.

Proof. The analysis starts from deriving a descent inequality for the virtual iterates \tilde{x}_t defined in (11). Consider two cases for virtual iterates: depending on whether restarts happen or not.

Iterations without restart: If restarts do not happen, namely $(t+1) \bmod \tau \neq 0$ then from the smoothness assumption of f and (11), we have

$$\begin{aligned}
 \mathbb{E}[f(\tilde{x}_{t+1})] &\leq \mathbb{E}[f(\tilde{x}_t)] - \gamma \mathbb{E}[\langle \nabla f(\tilde{x}_t), \nabla f(x_t) \rangle] + \frac{L\gamma^2}{2} \mathbb{E}[\|\nabla f(x_t)\|^2] \\
 &= \mathbb{E}[f(\tilde{x}_t)] - \frac{\gamma}{2} \mathbb{E}[\|\nabla f(\tilde{x}_t)\|^2] - \mathbb{E}\left[\frac{\gamma}{2}\|\nabla f(x_t)\|^2\right] + \frac{\gamma}{2} \mathbb{E}[\|\nabla f(\tilde{x}_t) - \nabla f(x_t)\|^2] \\
 &\quad + \frac{L\gamma^2}{2} \mathbb{E}[\|\nabla f(x_t)\|^2] \\
 &\leq \mathbb{E}[f(\tilde{x}_t)] - \frac{\gamma}{2} \mathbb{E}[\|\nabla f(\tilde{x}_t)\|^2] - \frac{\gamma}{2} \mathbb{E}[\|\nabla f(x_t)\|^2] + \frac{L^2\gamma}{2} \mathbb{E}[\|\tilde{x}_t - x_t\|^2] \\
 &\quad + \frac{L\gamma^2}{2} \mathbb{E}[\|\nabla f(x_t)\|^2] \\
 &\leq \mathbb{E}[f(\tilde{x}_t)] - \frac{\gamma}{2} \mathbb{E}[\|\nabla f(\tilde{x}_t)\|^2] - \frac{\gamma}{3} \mathbb{E}[\|\nabla f(x_t)\|^2] + \frac{L^2\gamma}{2} \mathbb{E}[\|\tilde{x}_t - x_t\|^2].
 \end{aligned}$$

Iterations with restart: If a restart happens, namely $(t+1) \bmod \tau = 0$ then

$$\begin{aligned}
 \tilde{x}_{t+1} &= x_{t+1} = x_t - \gamma g_{i_t}(x_{\pi_t}) \\
 &= \tilde{x}_t + (x_t - \tilde{x}_t) - \gamma \nabla f(x_t) + (\gamma \nabla f(x_t) - \gamma g_{i_t}(x_{\pi_t})) \\
 &= \tilde{x}_t - \gamma \nabla f(x_t) + \underbrace{\gamma \sum_{j=r(t)}^t \nabla f(x_j) - g_{i_j}(x_{\pi_j})}_{=\Delta_t^t}.
 \end{aligned}$$

Then we use smoothness of f to get

$$\begin{aligned}
 &\mathbb{E}[f(\tilde{x}_{t+1})] \\
 &\leq \mathbb{E}[f(\tilde{x}_t)] - \gamma \mathbb{E}[\langle \nabla f(\tilde{x}_t), \nabla f(x_t) - \tilde{\Delta}_t^t \rangle] + \frac{L\gamma^2}{2} \mathbb{E}[\|\nabla f(x_t) - \tilde{\Delta}_t^t\|^2] \\
 &\leq \mathbb{E}[f(\tilde{x}_t)] - \gamma \mathbb{E}[\langle \nabla f(\tilde{x}_t), \nabla f(x_t) \rangle] + \gamma \mathbb{E}[\langle \nabla f(\tilde{x}_t), \tilde{\Delta}_t^t \rangle] + L\gamma^2 \mathbb{E}[\|\nabla f(x_t)\|^2] + L\gamma^2 \mathbb{E}[\|\tilde{\Delta}_t^t\|^2] \\
 &\leq \mathbb{E}[f(\tilde{x}_t)] - \frac{\gamma}{2} \mathbb{E}[\|\nabla f(\tilde{x}_t)\|^2] - \frac{\gamma}{2} \mathbb{E}[\|\nabla f(x_t)\|^2] + \frac{\gamma}{2} \mathbb{E}[\|\nabla f(\tilde{x}_t) - \nabla f(x_t)\|^2] \\
 &\quad + \frac{1}{160L} \mathbb{E}[\|\nabla f(\tilde{x}_t)\|^2] + 40L\gamma^2 \mathbb{E}[\|\tilde{\Delta}_t^t\|^2] + L\gamma^2 \mathbb{E}[\|\nabla f(x_t)\|^2] + L\gamma^2 \mathbb{E}[\|\tilde{\Delta}_t^t\|^2] \\
 &\leq \mathbb{E}[f(\tilde{x}_t)] - \frac{\gamma}{2} \mathbb{E}[\|\nabla f(\tilde{x}_t)\|^2] - \frac{\gamma}{3} \mathbb{E}[\|\nabla f(x_t)\|^2] + \frac{L^2\gamma}{2} \mathbb{E}[\|\tilde{x}_t - x_t\|^2] + \frac{1}{160L} \mathbb{E}[\|\nabla f(\tilde{x}_t)\|^2],
 \end{aligned}$$

where in the third inequality Young's inequality is used.

If we denote by ξ_t the indicator function of restart event at $t+1$, namely $\xi_{k\tau-1} = 1$ for all $k \geq 1$ and is 0 otherwise, then we can unify the descent inequality for both cases as follows:

$$\begin{aligned}
 \mathbb{E}[f(\tilde{x}_{t+1})] &\leq \mathbb{E}[f(\tilde{x}_t)] - \frac{\gamma}{2} \mathbb{E}[\|\nabla f(\tilde{x}_t)\|^2] - \frac{\gamma}{3} \mathbb{E}[\|\nabla f(x_t)\|^2] + \frac{L^2\gamma}{2} \mathbb{E}[\|\tilde{x}_t - x_t\|^2] \\
 &\quad + \left(\frac{1}{160L} \mathbb{E}[\|\nabla f(\tilde{x}_t)\|^2] + 41L\gamma^2 \mathbb{E}[\|\Delta_t^t\|^2] \right) \xi_t, \quad \forall t \geq 0.
 \end{aligned} \tag{23}$$

Next, we apply summation over the entire iterates and bound the terms that appear only in every τ iteration.

$$\begin{aligned}
 \frac{1}{L} \mathbb{E} [\|\nabla f(\tilde{x}_t)\|^2] &= \frac{1}{L\tau} \sum_{j=0}^{\tau-1} \mathbb{E} [\|\nabla f(\tilde{x}_t)\|^2] \\
 &\leq \frac{2}{L\tau} \sum_{j=0}^{\tau-1} \mathbb{E} [\|\nabla f(\tilde{x}_t) - \nabla f(\tilde{x}_{t-j})\|^2] + \frac{2}{L\tau} \sum_{j=0}^{\tau-1} \mathbb{E} [\|\nabla f(\tilde{x}_{t-j})\|^2] \\
 &\leq \frac{2L}{\tau} \sum_{j=0}^{\tau-1} \mathbb{E} [\|\tilde{x}_t - \tilde{x}_{t-j}\|^2] + \frac{2}{L\tau} \sum_{j=0}^{\tau-1} \mathbb{E} [\|\nabla f(\tilde{x}_{t-j})\|^2] \\
 &\leq \frac{2L\gamma^2}{\tau} \sum_{j=0}^{\tau-1} \mathbb{E} \left[\left\| \sum_{l=t-j}^{t-1} \nabla f(x_l) \right\|^2 \right] + \frac{2}{L\tau} \sum_{j=0}^{\tau-1} \mathbb{E} [\|\nabla f(\tilde{x}_{t-j})\|^2] \\
 &\leq 2L\gamma^2 \sum_{j=0}^{\tau-1} \sum_{l=t-j}^{t-1} \mathbb{E} [\|\nabla f(x_l)\|^2] + \frac{2}{L\tau} \sum_{j=0}^{\tau-1} \mathbb{E} [\|\nabla f(\tilde{x}_{t-j})\|^2] \\
 &\leq 2L\gamma^2 \tau \sum_{j=0}^{\tau-1} \mathbb{E} [\|\nabla f(x_{t-j})\|^2] + \frac{2}{L\tau} \sum_{j=0}^{\tau-1} \mathbb{E} [\|\nabla f(\tilde{x}_{t-j})\|^2] \\
 &\leq \frac{\gamma}{10} \sum_{j=0}^{\tau-1} \mathbb{E} [\|\nabla f(x_{t-j})\|^2] + 80\gamma \sum_{j=0}^{\tau-1} \mathbb{E} [\|\nabla f(\tilde{x}_{t-j})\|^2]
 \end{aligned}$$

provided that $\frac{1}{40} \leq L\gamma\tau \leq \frac{1}{20}$ (e.g., $\tau = \lfloor \frac{1}{20L\gamma} \rfloor$). Then we can use this bound to derive

$$\sum_{t=0}^{T-1} \frac{1}{160L} \mathbb{E} [\|\nabla f(\tilde{x}_t)\|^2] \xi_t \leq \frac{\gamma}{1600} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(x_t)\|^2] + \frac{\gamma}{2} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(\tilde{x}_t)\|^2]. \quad (24)$$

Then we use Lemma C.1 to bound Δ_t^t :

$$\begin{aligned}
 L\gamma^2 \sum_{t=0}^{T-1} \mathbb{E} [\|\Delta_t^t\|^2] \xi_t &\leq 4L\gamma^2 \sum_{t=0}^{T-1} \mathbb{E} [\phi_t^t(x_{r(t)})] \xi_t + \frac{L\gamma^2}{6\tau} \sum_{t=0}^{T-1} \sum_{j=r(t)}^t \mathbb{E} [\phi_t^{j-1}(x_{r(t)})] \xi_t \\
 &\quad + \frac{25}{6} L^3 \gamma^2 \tau \sum_{t=0}^{T-1} \sum_{j=r(t)}^t \mathbb{E} [\|x_j - x_{\pi_j}\|^2] \xi_t \\
 &\quad + \frac{L\gamma^2 \tau}{24} \sum_{t=0}^{T-1} \sum_{j=r(t)}^t \mathbb{E} [\|\nabla f(x_j)\|^2] \xi_t + \frac{L\gamma^2 \tau}{24} \sigma^2 \sum_{t=0}^{T-1} \xi_t.
 \end{aligned}$$

Notice that summation over the entire iterates with weights ξ_t is equivalent to division by τ .

$$\sum_{t=0}^{T-1} \sum_{j=r(t)}^t \phi_t^{j-1}(x_{r(t)}) \xi_t = \sum_{t=0}^{T-1} \sum_{j=r(t)}^{t-1} \phi_t^j(x_{r(t)}) \xi_t = \sum_{t=0}^{T-1} \sum_{j=r(t)}^{t-1} \phi_j^j(x_{r(j)}) \xi_t \leq \sum_{t=0}^{T-1} \phi_t^t(x_{r(t)}).$$

Hence

$$\begin{aligned}
 41L\gamma^2 \sum_{t=0}^{T-1} \mathbb{E} [\|\tilde{\Delta}_t^t\|^2] \xi_t &\leq 164L\gamma^2 \sum_{t=0}^{T-1} \mathbb{E} [\phi_t^t(x_{r(t)})] \xi_t + \frac{41L\gamma^2}{6\tau} \sum_{t=0}^{T-1} \mathbb{E} [\phi_t^t(x_{r(t)})] \\
 &\quad + \underbrace{171L^3 \gamma^2 \tau \sum_{t=0}^{T-1} \mathbb{E} [\|x_t - x_{\pi_t}\|^2]}_{=A} \\
 &\quad + \frac{41}{24} L\gamma^2 \tau \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(x_t)\|^2] + \frac{41}{24} L\gamma^2 \sigma^2 T.
 \end{aligned}$$

where $A \leq \frac{1}{132L^2}B + \frac{100}{33}\gamma^2\Psi + \frac{\gamma T\sigma^2}{5L}$ due to Lemma C.3. Therefore,

$$\begin{aligned}
 & 41L\gamma^2 \sum_{t=0}^{T-1} \mathbb{E} [\|\Delta_t^t\|^2] \xi_t \leq 164L\gamma^2 \sum_{t=0}^{T-1} \mathbb{E} [\phi_t^t(x_{r(t)})] \xi_t + \frac{41L\gamma^2}{6\tau} \sum_{t=0}^{T-1} \mathbb{E} [\phi_t^t(x_{r(t)})] \\
 & + \frac{171}{132}L\gamma^2\tau B + 520L^3\gamma^4\tau\Psi + \frac{171}{100}L\gamma^2\sigma^2T \\
 & + \frac{41}{24}L\gamma^2\tau \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(x_t)\|^2] + \frac{41}{24}L\gamma^2\sigma^2T \\
 & \leq 164L\gamma^2 \sum_{t=0}^{T-1} \mathbb{E} [\phi_t^t(x_{r(t)})] \xi_t + \frac{41L\gamma^2}{6\tau} \sum_{t=0}^{T-1} \mathbb{E} [\phi_t^t(x_{r(t)})] + \frac{171}{132}L\gamma^2\tau B + 26L^2\gamma^3\Psi \\
 & + \frac{171}{100}L\gamma^2\sigma^2T + \frac{41}{24}L\gamma^2\tau B + \frac{41}{16}L\gamma^2\sigma^2T \\
 & \leq 6950L^2\gamma^3 \sum_{t=0}^{T-1} \mathbb{E} [\phi_t^t(x_{r(t)})] + \frac{793}{5280}\gamma B + 26L^2\gamma^3\Psi + 4L\gamma^2\sigma^2T. \tag{25}
 \end{aligned}$$

Plugging (24) and (25) (and lemmas as well) into (23) and adding summation, we have

$$\begin{aligned}
 & \mathbb{E}[f(\tilde{x}_T) - f(\tilde{x}_0)] \leq -\frac{\gamma}{2} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(\tilde{x}_t)\|^2] - \frac{\gamma}{3} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(x_t)\|^2] + \frac{L^2\gamma}{2} \sum_{t=0}^{T-1} \mathbb{E} [\|\tilde{x}_t - x_t\|^2] \\
 & + \frac{1}{160L} \sum_{t=0}^{T-1} \xi_t \mathbb{E} [\|\nabla f(\tilde{x}_t)\|^2] + 41L\gamma^2 \sum_{t=0}^{T-1} \xi_t \mathbb{E} [\|\Delta_t^t\|^2] \\
 & \leq -\frac{\gamma}{2} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(\tilde{x}_t)\|^2] - \frac{\gamma}{3}B \\
 & + \frac{L^2\gamma}{2} \left(5\gamma^2 \sum_{t=0}^{T-1} \mathbb{E} [\phi_t^{t-1}(x_{r(t)})] + \gamma^2\Psi + \frac{1}{5460L^2}B + \frac{1}{240L}\gamma\sigma^2T \right) \\
 & + \frac{\gamma}{1600}B + \frac{\gamma}{2} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(\tilde{x}_t)\|^2] \\
 & + 6950L^2\gamma^3 \sum_{t=0}^{T-1} \mathbb{E} [\phi_t^t(x_{r(t)})] + \frac{793}{5280}\gamma B + 26L^2\gamma^3\Psi + 4L\gamma^2\sigma^2T \\
 & \leq \left(-\frac{\gamma}{2} + \frac{\gamma}{2} \right) \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(\tilde{x}_t)\|^2] + \left(-\frac{\gamma}{3} + \frac{\gamma}{10920} + \frac{\gamma}{1600} + \frac{793}{5280}\gamma \right) B \\
 & + \frac{5}{2}L^2\gamma^3 \sum_{t=0}^{T-1} \mathbb{E} [\phi_t^{t-1}(x_{r(t)})] + 6950L^2\gamma^3 \sum_{t=0}^{T-1} \mathbb{E} [\phi_t^t(x_{r(t)})] + 27L^2\gamma^3\Psi \\
 & + 5L\gamma^2\sigma^2T \\
 & \leq -\frac{\gamma}{5}B + \frac{5}{2}L^2\gamma^3 \sum_{t=0}^{T-1} \mathbb{E} [\phi_t^{t-1}(x_{r(t)})] + 6950L^2\gamma^3 \sum_{t=0}^{T-1} \mathbb{E} [\phi_t^t(x_{r(t)})] + 27L^2\gamma^3\Psi \\
 & + 5L\gamma^2\sigma^2T
 \end{aligned}$$

According to the statement of the theorem, we assume that

$$\sigma_{k,\tau}^2 := \sup_{x \in \mathbb{R}^d} \max_{0 \leq j < \tau} \mathbb{E} \left[\left\| \sum_{t=k\tau}^{\min\{k\tau+j, T\}} (\nabla f_{i_t}(x_{k\tau}) - \nabla f(x_{k\tau})) \right\|^2 \right],$$

Algorithm 2 Pure Asynchronous SGD

- 1: **Input:** $x_0 \in \mathbb{R}^d$, stepsize $\gamma > 0$, set of assigned jobs $\mathcal{A}_0 = \emptyset$, set of received jobs $\mathcal{R}_0 = \emptyset$
 - 2: **Initialization:** for all jobs $(i, 0) \in \mathcal{A}_1$, server assigns worker i to compute a stochastic gradient $g_i(x_0)$
 - 3: **for** $t = 0, 1, 2, \dots, T - 1$ **do**
 - 4: once worker i_t finishes a job $(i_t, \pi_t) \in \mathcal{A}_{t+1}$, it sends $g_{i_t}(x_{\pi_t})$ to the server
 - 5: server updates the current model $x_{t+1} = x_t - \gamma g_{i_t}(x_{\pi_t})$ and the set $\mathcal{R}_{t+1} = \mathcal{R}_t \cup \{(i_t, \pi_t)\}$
 - 6: server assigns worker i_t to compute a gradient $g_{i_t}(x_{t+1})$
 - 7: server updates the set $\mathcal{A}_{t+2} = \mathcal{A}_{t+1} \cup \{(i_t, t + 1)\}$
 - 8: **end for**
-

and

$$\nu^2 := \sum_{t=0}^{T-1} \mathbb{E} \left[\left\| \sum_{j=\pi_t}^{t-1} (\nabla f_{i_j}(x_{\pi_j}) - \nabla f(x_{\pi_j})) \right\|^2 \right]$$

are bounded. So ν^2 upper bounds Ψ term in the above inequality while $\sigma_{k,\tau}^2$ upper bounds ϕ_t^{t-1} and ϕ_t^t terms. Then, by averaging we get

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(x_t)\|^2] \leq \frac{5(f(x_0) - f^*)}{\gamma T} + PL^2\gamma^2 \left(\frac{1}{\lfloor T/\tau \rfloor} \sum_{k=0}^{\lfloor \frac{T}{\tau} \rfloor - 1} \sigma_{k,\tau}^2 + \frac{1}{T} \nu^2 \right) + 25L\gamma\sigma^2, \quad (26)$$

where $P := 35000$, and we use the fact that $\tilde{x}_0 = x_0$ by the definition (11). □

C.3 Convergence Guarantees in Special Cases

We emphasize that in the choice of stepsizes, we omit numerical constants for simplicity, and show the dependency on the parameters of the problem only.

C.3.1 Pure Asynchronous SGD

At the beginning of the algorithm, the server assigns each worker to compute a gradient at x_0 , i.e. $\mathcal{A}_1 = \{(i, 0) : i \in [n]\}$. As soon as one of them (denoted by i_t) finishes its job, the server assigns him back new job at the updated point, i.e. $k_t \equiv i_t$. We present the convergence for pure asynchronous SGD.

Proposition C.1. Let Assumptions 1, 2, and 3 hold. Let the stepsize γ satisfy inequalities $20L\gamma\sqrt{\tau_{\max}\tau_C} \leq 1$ and $6L\gamma \leq 1$. Let $\tau = \lfloor \frac{1}{20L\gamma} \rfloor$. Then the iterates of Algorithm 2 satisfy

$$\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] \leq \mathcal{O} \left(\frac{F_0}{\gamma T} + L\gamma\sigma^2 + \zeta^2 \right), \quad (27)$$

where \hat{x}_t is chosen uniformly at random from $\{x_0, \dots, x_{T-1}\}$ and $F_0 := f(x_0) - f^*$. Moreover, if we tune the stepsize, then the iterates of pure asynchronous SGD satisfy

$$\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] \leq \mathcal{O} \left(\frac{LF_0\sqrt{\tau_{\max}\tau_C}}{T} + \left(\frac{LF_0\sigma^2}{T} \right)^{1/2} + \zeta^2 \right), \quad (28)$$

Proof. We need to bound quantities $\sigma_{k,\tau}^2$ and ν^2 . If within interval $[k\tau, k\tau + j]$ all indices from $[n]$ appear, then we do not have to consider them since

$$\sum_{i=1}^n (\nabla f_i(x) - \nabla f(x)) = 0.$$

Thus, we have

$$\sum_{t=k\tau}^{k\tau+j} (\nabla f_{i_t}(x) - \nabla f(x)) = \sum_{i \in \mathcal{S}} (\nabla f_i(x) - \nabla f(x)),$$

where $|\mathcal{S}| \leq \tau$, and hence, we continue as follows

$$\begin{aligned} \left\| \sum_{t=k\tau}^{k\tau+j} (\nabla f_{i_t}(x) - \nabla f(x)) \right\|^2 &= \left\| \sum_{i \in \mathcal{S}} (\nabla f_i(x) - \nabla f(x)) \right\|^2 \\ &\leq |\mathcal{S}| \sum_{i \in \mathcal{S}} \|\nabla f_i(x) - \nabla f(x)\|^2 \\ &\stackrel{\text{As.3}}{\leq} |\mathcal{S}|^2 \zeta^2 \leq \tau^2 \zeta^2. \end{aligned}$$

Next, we bound the term with ν^2 . We have

$$\begin{aligned} \nu^2 &= \sum_{t=0}^{T-1} \left\| \sum_{j=\pi_t}^{t-1} (\nabla f_{i_j}(x_{\pi_t}) - \nabla f(x_{\pi_t})) \right\|^2 \\ &\stackrel{\text{Lemma B.1}}{\leq} \sum_{t=0}^{T-1} \tau_t \sum_{j=\pi_t}^{t-1} \|\nabla f_{i_j}(x_{\pi_t}) - \nabla f(x_{\pi_t})\|^2 \\ &\stackrel{\text{Lemma B.3}}{\leq} \tau_C \tau_{\max} \zeta^2 T. \end{aligned}$$

Thus, the final rate is

$$\begin{aligned} \mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] &\leq \mathcal{O} \left(\frac{F_0}{\gamma T} + L\gamma\sigma^2 + L^2\gamma^2 \left(\frac{1}{\lfloor T/\tau \rfloor} \sum_{k=1}^{\lfloor T/\tau \rfloor} \tau^2 \zeta^2 + \frac{1}{T} \tau_C \tau_{\max} \zeta^2 T \right) \right) \\ &= \mathcal{O} \left(\frac{F_0}{\gamma T} + L\gamma\sigma^2 + L^2\gamma^2 \tau^2 \zeta^2 + L^2\gamma^2 \tau_C \tau_{\max} \zeta^2 \right). \end{aligned}$$

Using the stepsize restriction and the value of τ , we get

$$\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] \leq \mathcal{O} \left(\frac{F_0}{\gamma T} + L\gamma\sigma^2 + \zeta^2 \right).$$

Now we need to tune the stepsize. If we choose $\gamma = \Theta \left(\min \left\{ \frac{1}{L\sqrt{\tau_{\max}\tau_C}}, \left(\frac{F_0}{L\sigma^2 T} \right)^{1/2} \right\} \right)$, then we have two cases

- if $\gamma = \Theta \left(\frac{1}{L\sqrt{\tau_{\max}\tau_C}} \right)$, then

$$\begin{aligned} \mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] &\leq \mathcal{O} \left(\frac{F_0}{T} \sqrt{L\tau_{\max}\tau_C} + L\sigma^2 \left(\frac{F_0}{L\sigma^2 T} \right)^{1/2} + \zeta^2 \right) \\ &= \mathcal{O} \left(\frac{LF_0\sqrt{\tau_{\max}\tau_C}}{T} + \left(\frac{LF_0\sigma^2}{T} \right)^{1/2} + \zeta^2 \right). \end{aligned}$$

- $\gamma = \Theta \left(\left(\frac{F_0}{L\sigma^2 T} \right)^{1/2} \right)$, then

$$\begin{aligned} \mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] &\leq \mathcal{O} \left(\frac{F_0}{T} \left(\frac{L\sigma^2 T}{F_0} \right)^{1/2} + L\sigma^2 \left(\frac{F_0}{L\sigma^2 T} \right)^{1/2} + \zeta^2 \right) \\ &= \mathcal{O} \left(\left(\frac{LF_0\sigma^2}{T} \right)^{1/2} + \zeta^2 \right). \end{aligned}$$

It is left to choose the stepsize to be the minimum over two cases. \square

C.3.2 Mini-batch SGD

First, we show that the iterates of mini-batch SGD suit the update rule (2). The standard iteration of mini-batch SGD has the following form

$$z_{q+1} = z_q - \frac{\gamma}{b} \sum_{i \in B_q} \nabla f_i(z_q).^7$$

Let $B_q = \{i_{q,1}, \dots, i_{q,b}\}$ be the batch at iteration q , i.e. $i_{q,l} \in [n]$ for all q and l , and B_q is sampled from $[n]$ uniformly at random without replacement. We initialize $x_0 = z_0$, and the stepsize $\tilde{\gamma} = \frac{\gamma}{b}$. Then let us consider the following chain of updates

$$x_{qb+l} = x_{qb+l-1} - \tilde{\gamma} \nabla f_{i_{q,l}}(x_{qb}) \quad \text{for } l = 1, \dots, b.$$

Then,

$$\begin{aligned} x_{qb+b} &= x_{qb+b-1} - \tilde{\gamma} \nabla f_{i_{q,b-1}}(x_{qb}) \\ &\vdots \\ &= x_{qb} - \tilde{\gamma} \sum_{i \in B_q} \nabla f_i(x_{qb}), \end{aligned} \tag{29}$$

which is exactly the mini-batch update. Besides, we have that $x_{qb} \equiv z_q$.

Let us give now the values for τ_{\max} and τ_C . At the beginning of each step of mini-batch SGD the server selects uniformly at random b workers and sends them current model x_{qb} . Then, workers send one-by-one evaluated gradients while the server does not assign new jobs until the last worker finishes its job. Observing (29) we conclude that the maximum delay is $\tau_{\max} = b - 1 < b$ for the slowest worker while $\tau_C = b - 1 < b$ which is attained at the beginning of the batch.

Now we apply Theorem 1. For that, we need τ to be a multiple of b . This is achieved by carefully choosing the stepsize and τ restrictions.

Proposition C.2. Let Assumptions 1 and 3 hold. Let the stepsize γ satisfy $20L\gamma \leq 1$. Let $\tau = b \lfloor \frac{1}{20L\gamma} \rfloor$. Then the iterates of mini-batch SGD with batch size b satisfy

$$\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] \leq \mathcal{O} \left(\frac{F_0}{\gamma T} + \frac{L\gamma\zeta^2}{b} \right), \tag{30}$$

where \hat{x}_t is chosen uniformly at random from $\{x_0, \dots, x_{T-1}\}$ and $F_0 := f(x_0) - f^*$. Moreover, if we tune the stepsize, then the iterates of mini-batch SGD with batch size b satisfy

$$\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] \leq \mathcal{O} \left(\frac{LF_0}{T} + \sqrt{\frac{LF_0\zeta^2}{Tb}} \right). \tag{31}$$

Proof. We run mini-batch SGD from Algorithm 1 point of view with stepsize $\tilde{\gamma} = \frac{\gamma}{b}$. Then we should choose τ such that $20L\tilde{\gamma}\tau \leq 1$. So we can choose $\tau = b \lfloor \frac{1}{20L\tilde{\gamma}} \rfloor$. Now we have two restrictions on $\tilde{\gamma}$

$$\begin{aligned} \tilde{\gamma} &\leq \frac{1}{6L} \Rightarrow \gamma \leq \frac{b}{6L}, \\ \tilde{\gamma} &\leq \frac{1}{20L\sqrt{\tau_{\max}\tau_C}} \Rightarrow \gamma \leq \frac{1}{20L}. \end{aligned}$$

That is why we should assume that $\gamma \leq \frac{1}{20L}$ with previously chosen τ .

With the choice τ we have, every chunk of size τ consists of several full batches (each chunk has the same number of batches). This choice of τ is needed in order to apply conditional expectation correctly. Indeed, let us take the conditional expectation depending on all the events before iteration $(k-1)\tau$. This means that $x_{(k-1)\tau}$ is fixed as well since it is computed using the gradients before iteration $(k-1)\tau$.

⁷Here we denote iterates by z_q in order not to confuse them with the notation of the paper

Let us assume that each period of size τ consists of $m \geq 1$ batches. Then we can compute the conditional expectation from the definition of $\sigma_{k,\tau}^2$ splitting it into m independent conditional expectations for each batch. Note that in this case, the maximum over j in (10) is always attained for $j = \tau - 1$. Thus, we have

$$\begin{aligned}\sigma_{k,\tau}^2 &= \mathbb{E}_{(k-1)\tau} \left\| \sum_{j=k\tau}^{k\tau+\tau-1} (\nabla f_{i_j}(x_{(k-1)\tau}) - \nabla f(x_{(k-1)\tau})) \right\|^2 \\ &= m \cdot b^2 \frac{\zeta^2}{b} = mb\zeta^2 = \tau\zeta^2.\end{aligned}$$

where in the last step we use the variance bound of the mini-batch estimator. Note that in one node setting ζ^2 plays the role of standard variance bound of stochastic gradient estimator.

Taking conditional expectations one by one we can bound each $\sigma_{k,\tau}^2$ term by $\tau\zeta^2$.

Now we switch to the ν^2 term. Note that in this case, the delayed iterate for gradients within one batch is the same, i.e. for $j \in [qb, (q+1)b)$ we have $\pi_j = qb$, thus we can compute expectations correctly. We split ν^2 is split into $\frac{T}{b}$ terms ν_q^2 of the form

$$\begin{aligned}\nu_q^2 &:= \sum_{t=qb}^{(q+1)b-1} \mathbb{E}_{qb} \left[\left\| \sum_{j=\pi_t}^{t-1} (\nabla f_{i_j}(x_{\pi_j}) - \nabla f(x_{\pi_j})) \right\|^2 \right] \\ &= \sum_{t=qb}^{(q+1)b-1} \mathbb{E}_{qb} \left[\left\| \sum_{j=qb}^{t-1} (\nabla f_{i_j}(x_{qb}) - \nabla f(x_{qb})) \right\|^2 \right] \\ &\leq \sum_{t=qb}^{(q+1)b-1} (t - qb)^2 \frac{\zeta^2}{(t - qb)} \\ &= \sum_{t=qb}^{(q+1)b-1} (t - qb)\zeta^2 \\ &\leq \frac{b^2}{2}\zeta^2.\end{aligned}$$

Here we again need to take conditional expectation w.r.t x_{qb} , and afterwards compute bounds for ν_q^2 one by one. Saying that, we have $\nu^2 \leq \frac{T}{b} \frac{b^2}{2} \zeta^2 \leq Tb\zeta^2$. Now we apply Theorem 1.

$$\begin{aligned}\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] &\leq \mathcal{O} \left(\frac{F_0}{\tilde{\gamma}T} + L^2\tilde{\gamma}^2\tau\zeta^2 + L^2\tilde{\gamma}^2b\zeta^2 \right) \\ &\leq \mathcal{O} \left(\frac{F_0}{\gamma\frac{T}{b}} + \frac{L\gamma\zeta^2}{b} \right),\end{aligned}\tag{32}$$

where we use $L\tilde{\gamma}\tau = \Theta(1)$, $\tilde{\gamma} = \frac{\gamma}{b}$, and $\gamma \leq \frac{1}{20L}$. Note that T iterations of Algorithm 1 are equivalent to $\frac{T}{b}$ iterations of mini-batch SGD. That is why the right-hand side of (32) is a standard rate of mini-batch SGD. However, we observe that the left-hand side of (32) is slightly different from what we expect; there we get convergence for all intermediate iterates as well. In order to get the standard rate we need to modify restarting virtual iterates in the following way

$$\tilde{x}_{t+1} = \begin{cases} \tilde{x}_t - \gamma\nabla f(x_{\pi_t}) & \text{if } t+1 \neq \tau k \text{ for any } k \geq 1, \\ x_{t+1} & \text{if } t+1 = \tau k \text{ for some } k \geq 1, \end{cases}$$

then the left-hand side of (32) will be transformed to

$$\frac{1}{T} \sum_{q=0}^{T/b-1} \sum_{l=0}^{b-1} \mathbb{E} [\|\nabla f(x_{\pi_{qb+l}})\|^2] = \frac{1}{T} \sum_{q=0}^{T/b-1} \sum_{l=0}^{b-1} \mathbb{E} [\|\nabla f(x_{qb})\|^2] = \frac{1}{T/b} \sum_{q=0}^{T/b-1} \mathbb{E} [\|\nabla f(x_{qb})\|^2],$$

Algorithm 3 Pure Asynchronous SGD with waiting

1: **Input:** $x_0 \in \mathbb{R}^d$, stepsize $\gamma > 0$, set of assigned jobs $\mathcal{A}_0 = \emptyset$, set of received jobs $\mathcal{R}_0 = \emptyset$, batch size $b \geq 1$, gradient estimator $g = 0$ and number of received gradients $r = 0$
 2: **Initialization:** for all jobs $(i, 0) \in \mathcal{A}_1$, server assigns worker i to compute a stochastic gradient $g_i(x_0)$
 3: **for** $t = 0, 1, 2, \dots, T - 1$ **do**
 4: server sets $g_{t,0} = 0$ and $\mathcal{R}_{t,0} = \mathcal{R}_t$
 5: **for** $j = 1, \dots, b$ **do**
 6: once worker $i_{t,j}$ finishes a job $(i_{t,j}, \pi_{t,j}) \in \mathcal{A}_{t+1}$, it sends $g_{i_{t,j}}(x_{\pi_{t,j}})$ to the server
 7: server updates $g_{t,j} = g_{t,j-1} + g_{i_{t,j}}(x_{\pi_{t,j}})$
 8: server updates the set $\mathcal{R}_{t,j} = \mathcal{R}_{t,j-1} \cup \{(i_{t,j}, \pi_{t,j})\}$
 9: **end for**
 10: server updates the current model $x_{t+1} = x_t - \frac{\gamma}{b} g_{t,b}$ and set $\mathcal{R}_{t+1} = \mathcal{R}_{t,b}$
 11: server assigns worker $i_{t,j}$ to compute a gradient $g_{i_{t,j}}(x_{t+1})$ for all $j \in [b]$
 12: server updates the set $\mathcal{A}_{t+2} = \mathcal{A}_{t+1} \cup \{(i_{t,1}, t+1)\} \cup \dots \cup \{(i_{t,b}, t+1)\}$
 13: **end for**

which is exactly what we want. We only need to change T to Tb if we want to present the rate w.r.t. the number of mini-batch SGD steps.

Now, if we choose stepsize $\gamma = \Theta\left(\min\left\{\frac{1}{L}, \sqrt{\frac{F_0 b}{LT\zeta^2}}\right\}\right)$, then we have two possible options

- if $\gamma = \Theta\left(\frac{1}{L}\right)$ we have

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(x_t)\|^2] &\leq \mathcal{O}\left(\frac{F_0}{\frac{1}{L}T} + \frac{L\zeta^2}{b} \sqrt{\frac{F_0 b}{LT\zeta^2}}\right) \\ &\leq \mathcal{O}\left(\frac{LF_0}{T} + \sqrt{\frac{LF_0\zeta^2}{Tb}}\right). \end{aligned}$$

- if $\gamma = \Theta\left(\sqrt{\frac{F_0 b}{LT\zeta^2}}\right)$ we have

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(x_t)\|^2] &\leq \mathcal{O}\left(\frac{F_0}{T} \sqrt{\frac{LT\zeta^2}{F_0 b}} + \frac{L\zeta^2}{b} \sqrt{\frac{F_0 b}{LT\zeta^2}}\right) \\ &\leq \mathcal{O}\left(\sqrt{\frac{LF_0\zeta^2}{Tb}}\right). \end{aligned}$$

It is left to set the stepsize to be the minimum over two cases. □

C.3.3 Pure Asynchronous SGD with waiting

This case is similar to the previous one; the only change is that the server waits for the first b fastest workers and assigns them new jobs back. Formally, the update has the following form

$$z_{q+1} = z_q - \gamma \sum_{i \in B_q} g_i(z_{\pi_{q,i}}), \quad (33)$$

where $\pi_{q,i}$ is the iteration counter where worker i evaluated its gradient (possible delayed)⁸. Let $B_q := \{i_{q,1}, \dots, i_{q,b}\}$ be the set of b fastest workers at iteration q . Then, we can rewrite this update in the following form with $x_0 = z_0$.

⁸Here we again use z_t iterates notation in order not to confound with the updates of Algorithm 1

$$\begin{aligned}
 x_{qb+b} &= x_{qb+b-1} - \gamma g_{i_{q,b-1}}(x_{qb-\tau_{q,b-1}}) \\
 &\vdots \\
 &= x_{qb} - \gamma \sum_{l=0}^{b-1} g_{i_{q,l}}(x_{qb-\tau_{q,l}}).
 \end{aligned}$$

We extend one iteration of (33) to b intermediate iterations of Algorithm 1, i.e. the total number of iterations increases by a factor b . We also highlight the fact that the server does not assign new jobs before all the workers from the batch send the gradients. This means that workers always send gradients at points x_{qb} . Note that τ_C remains the same while τ_{\max} might increase b times in the worst case. We are ready to apply Theorem 1.

Proposition C.3. Let Assumptions 1, 2, and 3 hold. Let the stepsize γ satisfies $20L\sqrt{b\tau_{\max}\tau_C}\gamma \leq b$ and $6L\gamma \leq 1$. Then the iterates of Algorithm 3 satisfy

$$\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] \leq \mathcal{O} \left(\frac{F_0}{\gamma T} + L\gamma \frac{\sigma^2}{b} + \zeta^2 \right), \quad (34)$$

where \hat{x}_t is chosen uniformly at random from $\{x_0, \dots, x_{b-1}, x_b, \dots, x_{2b-1}, \dots, x_{Tb-1}\}$ and $F_0 = f(x_0) - f^*$. Moreover, if we tune the stepsize, then the iterates of pure asynchronous SGD with waiting satisfy

$$\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] \leq \mathcal{O} \left(\frac{F_0\sqrt{\tau_{\max}\tau_C}}{T\sqrt{b}} + \left(\frac{LF_0\sigma^2}{Tb} \right)^{1/2} + \zeta^2 \right). \quad (35)$$

Proof. We again need to find restrictions on the stepsize. We have

$$\begin{aligned}
 6L\tilde{\gamma} \leq 1 &\Rightarrow \gamma \leq \frac{b}{6L}, \\
 20L\tilde{\gamma}\sqrt{b\tau_{\max}\tau_C} \leq 1 &\Rightarrow \gamma \leq \frac{b}{20L\sqrt{b\tau_{\max}\tau_C}}.
 \end{aligned}$$

Then we choose $\tau = \lfloor \frac{1}{20L\tilde{\gamma}} \rfloor$, and similarly to the case of pure asynchronous SGD we bound $\sigma_{k,\tau}^2 \leq \tau^2\zeta^2$ while $\nu^2 \leq b\tau_C\tau_{\max}\zeta^2T$. Then the rate is

$$\begin{aligned}
 \frac{1}{Tb} \sum_{q=0}^{T-1} \sum_{l=0}^{b-1} \mathbb{E} [\|\nabla f(x_{qb+l})\|^2] &\leq \mathcal{O} \left(\frac{F_0}{\tilde{\gamma}Tb} + L\tilde{\gamma}\sigma^2 + L^2\tilde{\gamma}^2\tau^2\zeta^2 + L^2\tilde{\gamma}^2b\tau_C\tau_{\max}\zeta^2 \right) \\
 &\leq \mathcal{O} \left(\frac{F_0}{\gamma T} + L\gamma \frac{\sigma^2}{b} + \zeta^2 \right).
 \end{aligned}$$

It is left to set the stepsize choice in a similar way as for pure asynchronous SGD. \square

C.3.4 SGD with Random Reshuffling

In the beginning of each epoch the server sample a random permutation χ_q of $[n]$, and then gradients come following that order. We initialize $x_0 = z_0$. The chain of updates in this case has the form as follows

$$\begin{aligned}
 x_{qn+n} &= x_{qn+n-1} - \gamma \nabla f_{\chi_q(n-1)}(x_{qn+n-1}) \\
 &\vdots \\
 &= x_{qn} - \gamma \sum_{l=0}^{n-1} \nabla f_{\chi_q(l)}(x_{qn+l}).
 \end{aligned}$$

Note that in this case, the server receives gradients from the workers one by one without delays, i.e. $\tau_{\max} = 0$. This means that we have the only stepsize restriction $\gamma \leq \frac{1}{6L}$ since the other one $20L\gamma\sqrt{\tau_{\max}\tau_C} \leq 1$ holds for any choice of γ .

Moreover, since we take expectation only once at the end of the proof, we cover the case of SGD shuffle as well without any additional changes. Now we apply Theorem 1 to derive convergence guarantees in this case.

Proposition C.4. Let Assumptions 1, 3 hold. Let the stepsize γ satisfy $20nL\gamma \leq 1$. Let $\tau = n \lfloor \frac{1}{20Ln\gamma} \rfloor$. Then the iterates of SGD with random reshuffling satisfy

$$\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] \leq \mathcal{O} \left(\frac{F_0}{\gamma T} + L^2 n \gamma^2 \zeta^2 \right), \quad (36)$$

where \hat{x}_t is chosen uniformly at random from $\{x_0, \dots, x_{n-1}, x_n, \dots, x_{2n-1}, \dots, x_{T-1}\}$ and $F_0 := f(x_0) - f^*$, and T is the total number of gradient evaluations. If we tune the stepsize, then the iterates of SGD with random reshuffling satisfy

$$\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] \leq \mathcal{O} \left(\frac{LF_0 n}{T} + \left(\frac{LF_0 \sqrt{n} \zeta}{T} \right)^{2/3} \right). \quad (37)$$

Proof. We need τ to be a multiple of n in order to be able to correctly apply conditional expectations (the same trick as for mini-batch SGD). We force the stepsize $\gamma \leq \frac{1}{20Ln}$. Then we automatically satisfy the restriction $6L\gamma \leq 1$. Now we need to choose τ . We need $20L\gamma\tau \leq 1$, so let us choose $\tau = n \lfloor \frac{1}{20Ln\gamma} \rfloor$, then $20L\gamma\tau \leq 20L\gamma n \frac{1}{20Ln\gamma} \leq 1$.

We use the computations from [Koloskova et al. \(2022\)](#), Section C.2. Thus, we have $\sigma_{k,\tau}^2 \leq \min\{\tau, n\}\zeta^2$. The fact that $\tau \geq n$ implies $\sigma_{k,\tau}^2 \leq n\zeta^2$. Since we do not have delays in this case, then $\nu^2 = 0$. This gives us the following result applying Theorem 1.

$$\frac{1}{T} \sum_{q=0}^{T/n-1} \sum_{l=0}^{n-1} \mathbb{E} [\|\nabla f(x_{qn+l})\|^2] \leq \mathcal{O} \left(\frac{F_0}{\gamma T} + L^2 n \gamma^2 \zeta^2 \right).$$

Now, if we need to tune the stepsize. Similarly to mini-batch SGD we should consider two cases.

- if $\gamma = \Theta(\frac{1}{Ln})$ we have

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(x_t)\|^2] &\leq \mathcal{O} \left(\frac{F_0}{\frac{1}{Ln} T} + L^2 n \zeta^2 \left(\frac{F_0}{L^2 n T \zeta^2} \right)^{2/3} \right) \\ &\leq \mathcal{O} \left(\frac{LF_0 n}{T} + \left(\frac{L^2 F_0^2 n \zeta^2}{T^2} \right)^{1/3} \right). \end{aligned}$$

- if $\gamma = \Theta \left(\left(\frac{F_0}{L^2 n T \zeta^2} \right)^{1/3} \right)$ we have

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(x_t)\|^2] &\leq \mathcal{O} \left(\frac{F_0}{T} \left(\frac{L^2 n T \zeta^2}{F_0} \right)^{1/3} + L^2 n \zeta^2 \left(\frac{F_0}{L^2 n T \zeta^2} \right)^{2/3} \right) \\ &\leq \mathcal{O} \left(\left(\frac{L^2 F_0^2 n \zeta^2}{T^2} \right)^{1/3} \right). \end{aligned}$$

To get the final rate after tuning we need to set the stepsize as the minimum over two cases. \square

D PROOFS FOR ANALYSIS OF GRADIENT ASSIGNING PROCESS

In this section the iteration counter starts with 1, i.e. we consider y_1 as an initial point. Following (13) and (14), we have

$$y_t = y_{r(t)} - \gamma \sum_{j=r(t)}^{t-1} g_{k_j}(x_{\alpha_j}), \quad \tilde{y}_t = y_{r(t)} - \gamma \sum_{j=r(t)}^{t-1} \nabla f(x_j),$$

where $g_{k_j}(x_{\alpha_j})$ is an unbiased stochastic estimator of $\nabla f_{k_j}(x_{\alpha_j})$ with variance bounded by σ^2 .

Analogously to the proofs of Theorem 2 we define

$$\tilde{\Delta}_t^m = \sum_{j=r(t)}^m \nabla f(x_j) - g_{k_j}(x_{\alpha_j}), \quad \tilde{\phi}_t^m(x) = \left\| \sum_{j=r(t)}^m \nabla f_{k_j}(x) - \nabla f(x) \right\|^2.$$

In this section we use shortcuts

$$\begin{aligned} \tilde{A} &:= \sum_{t=0}^{T-1} \mathbb{E} [\|\tilde{y}_t - y_{\alpha_t}\|^2], \quad B := \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(x_t)\|^2], \quad \tilde{\Phi} := \sum_{t=0}^{T-1} \mathbb{E} [\tilde{\phi}_t^{t-1}(x_{r(t)})] \\ \tilde{\Psi} &:= \sum_{t=0}^{T-1} \mathbb{E} \left[\left\| \sum_{j=\pi_t}^{t-1} \nabla f_{k_j}(x_{\alpha_j}) - \nabla f(x_{\alpha_j}) \right\|^2 \right]. \end{aligned}$$

D.1 Key Lemmas

Lemma D.1. Let real $\{x_t\}_{t=0}^T$ and virtual $\{y_t\}_{t=0}^T$ iterates be defined in (2) and (13) respectively. Then

$$x_t - y_t = \gamma \sum_{(i,j) \in \mathcal{A}_t \setminus \mathcal{R}_t} g_i(x_j).$$

Proof. We prove the statement by induction. The base is trivial $x_0 - y_0 = 0$ by definition while $\mathcal{A}_0 = \mathcal{R}_0 = \emptyset$ as well. Let the statement hold at iteration t , now we show that it is true at iteration $t+1$ as well. Indeed,

$$\begin{aligned} x_{t+1} - y_{t+1} &= (x_t - \nabla f_{i_t}(x_{\pi_t})) - (y_t - \gamma \nabla f_{k_t}(x_{\alpha_t})) \\ &\stackrel{\text{Base}}{=} \gamma \sum_{(i,j) \in \mathcal{A}_t \setminus \mathcal{R}_t} \nabla f_i(x_j) + \gamma (\nabla f_{k_t}(x_{\alpha_t}) - f_{i_t}(x_{\pi_t})). \end{aligned} \quad (38)$$

Now we note that the connection between sets is updating as follows

$$\begin{aligned} \mathcal{A}_{t+1} \setminus \mathcal{R}_{t+1} &= (\mathcal{A}_t \cup \{(k_t, \alpha_t)\}) \setminus (\mathcal{R}_t \cup \{(i_t, \pi_t)\}) \\ &= (\mathcal{A}_t \setminus \mathcal{R}_t) \cup \{(k_t, \alpha_t)\} \setminus \{(i_t, \pi_t)\}. \end{aligned}$$

Plugging this into (38) we get $x_{t+1} - y_{t+1} = \gamma \sum_{(i,j) \in \mathcal{A}_{t+1} \setminus \mathcal{R}_{t+1}} \nabla f_i(x_j)$. \square

Lemma D.2. At time step t the following inequality holds

$$\mathbb{E} [\|y_t - x_t\|^2] \leq \gamma^2 (\tau_C - 1)^2 G^2 + (\tau_C - 1) \gamma^2 \sigma^2. \quad (39)$$

Proof. Loaded with Lemma D.1

$$\begin{aligned} \mathbb{E} [\|y_t - x_t\|^2] &= \gamma^2 \mathbb{E} \left[\left\| \sum_{(i,j) \in \mathcal{A}_t \setminus \mathcal{R}_t} g_i(x_j) \right\|^2 \right] \leq \gamma^2 \mathbb{E} \left[\left\| \sum_{(i,j) \in \mathcal{A}_t \setminus \mathcal{R}_t} \nabla f_i(x_j) \right\|^2 \right] + \gamma^2 \sigma^2 (\tau_C - 1) \\ &\leq \gamma^2 (\tau_C - 1) \sum_{(i,j) \in \mathcal{A}_t \setminus \mathcal{R}_t} \mathbb{E} [\|\nabla f_i(x_j)\|^2] + \gamma^2 \sigma^2 (\tau_C - 1) \\ &\leq \gamma^2 (\tau_C - 1)^2 G^2 + \gamma^2 \sigma^2 (\tau_C - 1), \end{aligned}$$

since $|\mathcal{A}_{t+1}| = |\mathcal{A}_t| + 1$. \square

Lemma D.3. If $30L\gamma\tau \leq 1$, then

$$\begin{aligned} \mathbb{E} [\|\tilde{\Delta}_t^m\|^2] &\leq 4\mathbb{E} [\tilde{\phi}_t^m(x_{r(t)})] + \frac{48}{7} L\gamma \sum_{j=r(t)}^m \mathbb{E} [\tilde{\phi}_t^{j-1}(x_{r(t)})] + \frac{6}{71} (\tau_C - 1)^2 G^2 + \frac{6}{71} (\tau_C - 1) \sigma^2 \\ &\quad + \frac{90}{7} L^2 \tau \sum_{j=r(t)}^m \mathbb{E} [\|y_j - y_{\alpha_j}\|^2] + \frac{4}{71} \tau \sum_{j=r(t)}^m \mathbb{E} [\|\nabla f(x_j)\|^2] + \frac{2}{35} \tau \sigma^2. \end{aligned} \quad (40)$$

Proof. Using Young's inequality we have

$$\begin{aligned}
 \mathbb{E} \left[\|\tilde{\Delta}_t^m\|^2 \right] &\leq \mathbb{E} \left[\left\| \sum_{j=r(t)}^m \nabla f_{k_j}(x_{\alpha_j}) - \nabla f(x_j) \right\|^2 \right] + \tau\sigma^2 \\
 &= \mathbb{E} \left[\left\| \sum_{j=r(t)}^m \nabla f_{k_j}(x_{\alpha_j}) \pm \nabla f_{k_j}(x_j) \pm \nabla f_{k_j}(x_{r(t)}) \pm \nabla f(x_{r(t)}) - \nabla f(x_j) \right\|^2 \right] + \tau\sigma^2 \\
 &\leq 4\mathbb{E} \left[\left\| \sum_{j=r(t)}^m \nabla f_{k_j}(x_{\alpha_j}) - \nabla f_{k_j}(x_j) \right\|^2 \right] + 4\mathbb{E} \left[\left\| \sum_{j=r(t)}^m \nabla f_{k_j}(x_j) - \nabla f_{k_j}(x_{r(t)}) \right\|^2 \right] \\
 &\quad + 4\mathbb{E} \left[\left\| \sum_{j=r(t)}^m \nabla f_{k_j}(x_{r(t)}) - \nabla f(x_{r(t)}) \right\|^2 \right] + 4\mathbb{E} \left[\left\| \sum_{j=r(t)}^m \nabla f(x_{r(t)}) - \nabla f(x_j) \right\|^2 \right] + \tau\sigma^2
 \end{aligned}$$

We continue as follows

$$\begin{aligned}
 \mathbb{E} \left[\|\tilde{\Delta}_t^m\|^2 \right] &\leq 4\mathbb{E} \left[\tilde{\phi}_t^t(x_{r(t)}) \right] + 4L^2\tau \sum_{j=r(t)}^m \mathbb{E} [\|x_j - x_{\alpha_j}\|^2] + 8L^2\tau \sum_{j=r(t)}^m \mathbb{E} [\|x_j - x_{r(t)}\|^2] \\
 &\quad + \tau\sigma^2 \\
 &\leq 4\mathbb{E} \left[\tilde{\phi}_t^m(x_{r(t)}) \right] + 12L^2\tau \sum_{j=r(t)}^m \mathbb{E} [\|x_j - y_j\|^2 + \|y_j - y_{\alpha_j}\|^2 + \|y_{\alpha_j} - x_{\alpha_j}\|^2] \\
 &\quad + 24L^2\tau \sum_{j=r(t)}^m \mathbb{E} [\|x_j - y_j\|^2 + \|y_j - y_{r(t)}\|^2 + \|y_{r(t)} - x_{r(t)}\|^2] + \tau\sigma^2.
 \end{aligned}$$

For the terms of the form $\|x_a - y_a\|^2$ for some a we use Lemma C.1. Using the above and the stepsize restriction, we have

$$\begin{aligned}
 \mathbb{E} \left[\|\tilde{\Delta}_t^m\|^2 \right] &\leq 4\mathbb{E} \left[\tilde{\phi}_t^m(x_{r(t)}) \right] + \frac{2}{25}(\tau_C - 1)^2 G^2 + 12L^2\tau \sum_{j=r(t)}^m \mathbb{E} [\|y_j - y_{\alpha_j}\|^2] + \tau\sigma^2 \\
 &\quad + 24L^2\tau \sum_{j=r(t)}^m \mathbb{E} [\|y_{r(t)} - y_j\|^2] + \frac{2}{25}(\tau_C - 1)\sigma^2, \tag{41}
 \end{aligned}$$

For the last term, we have

$$\begin{aligned}
 \sum_{j=r(t)}^m \mathbb{E} [\|y_{r(t)} - y_j\|^2] &= \sum_{j=r(t)}^m \gamma^2 \mathbb{E} \left[\left\| \sum_{l=r(t)}^{j-1} g_{k_l}(x_{\alpha_l}) \right\|^2 \right] \\
 &\leq 2\gamma^2 \sum_{j=r(t)}^m \mathbb{E} \left[\left\| \sum_{l=r(t)}^{j-1} g_{k_l}(x_{\alpha_l}) - \nabla f(x_l) \right\|^2 \right] + 2\gamma^2 \sum_{j=r(t)}^m \mathbb{E} \left[\left\| \sum_{l=r(t)}^{j-1} \nabla f(x_l) \right\|^2 \right] \\
 &= 2\gamma^2 \sum_{j=r(t)}^m \mathbb{E} [\|\tilde{\Delta}_t^{j-1}\|^2] + 2\gamma^2 \tau \sum_{j=r(t)}^m \sum_{l=r(t)}^{j-1} \mathbb{E} [\|\nabla f(x_l)\|^2] \\
 &\stackrel{(41)}{\leq} 8\gamma^2 \sum_{j=r(t)}^m \mathbb{E} [\tilde{\phi}_t^{j-1}(x_{r(t)})] + \frac{4}{25}\gamma^2(\tau_C - 1)^2 G^2 \tau + 24\gamma^2 L^2 \tau \sum_{j=r(t)}^m \mathbb{E} \left[\sum_{l=r(t)}^{j-1} \|y_l - y_{\alpha_l}\|^2 \right] \\
 &\quad + 48\gamma^2 L^2 \tau \sum_{j=r(t)}^m \sum_{l=r(t)}^{j-1} \mathbb{E} [\|y_{r(t)} - y_l\|^2] + 2\gamma^2 \tau \sum_{j=r(t)}^m \sum_{l=r(t)}^{j-1} \mathbb{E} [\|\nabla f(x_l)\|^2] + 2\gamma^2 \tau^2 \sigma^2 \\
 &\quad + \frac{4}{25}\gamma^2(\tau_C - 1)\sigma^2 \\
 &\leq 8\gamma^2 \sum_{j=r(t)}^m \tilde{\phi}_t^{j-1}(x_{r(t)}) + \frac{4}{25}\gamma^2(\tau_C - 1)^2 G^2 \tau + 24\gamma^2 L^2 \tau^2 \sum_{j=r(t)}^m \|y_j - y_{\alpha_j}\|^2 \\
 &\quad + 48\gamma^2 L^2 \tau^2 \sum_{j=r(t)}^m \|y_{r(t)} - y_j\|^2 + 2\gamma^2 \tau^2 \sum_{j=r(t)}^m \|\nabla f(x_j)\|^2 + 2\gamma^2 \tau^2 \sigma^2 + \frac{4}{25}\gamma^2(\tau_C - 1)\sigma^2 \\
 &\leq 8\gamma^2 \sum_{j=r(t)}^m \tilde{\phi}_t^{j-1}(x_{r(t)}) + \frac{4}{25}\gamma^2(\tau_C - 1)^2 G^2 \tau + \frac{2}{75} \sum_{j=r(t)}^m \|y_j - y_{\alpha_j}\|^2 \\
 &\quad + \frac{4}{75} \sum_{j=r(t)}^m \|y_{r(t)} - y_j\|^2 + 2\gamma^2 \tau^2 \sum_{j=r(t)}^m \|\nabla f(x_j)\|^2 + 2\gamma^2 \tau^2 \sigma^2 + \frac{4}{25}\gamma^2(\tau_C - 1)\sigma^2.
 \end{aligned}$$

Hence,

$$\begin{aligned}
 \sum_{j=r(t)}^m \|y_{r(t)} - y_j\|^2 &\leq \frac{60}{7}\gamma^2 \sum_{j=r(t)}^m \tilde{\phi}_t^{j-1}(x_{r(t)}) + \frac{12}{71}\gamma^2(\tau_C - 1)^2 G^2 \tau + \frac{2}{71} \sum_{j=r(t)}^m \|y_j - y_{\alpha_j}\|^2 \\
 &\quad + \frac{150}{71}\gamma^2 \tau^2 \sum_{j=r(t)}^m \|\nabla f(x_j)\|^2 + \frac{150}{71}\gamma^2 \tau^2 \sigma^2 + \frac{12}{71}\gamma^2(\tau_C - 1)\sigma^2. \tag{42}
 \end{aligned}$$

Plugging (42) in (41) we get

$$\begin{aligned}
 \mathbb{E} \left[\|\tilde{\Delta}_t^m\|^2 \right] &\leq 4\mathbb{E} \left[\tilde{\phi}_t^m(x_{r(t)}) \right] + \frac{2}{25}(\tau_C - 1)^2 G^2 + 12L^2\tau \sum_{j=r(t)}^m \mathbb{E} [\|y_j - y_{\alpha_j}\|^2] \\
 &\quad + 24L^2\tau \left(\frac{60}{7}\gamma^2 \sum_{j=r(t)}^m \mathbb{E} [\tilde{\phi}_t^{j-1}(x_{r(t)})] + \frac{12}{71}\gamma^2(\tau_C - 1)^2 G^2\tau + \frac{2}{71} \sum_{j=r(t)}^m \mathbb{E} [\|y_j - y_{\alpha_j}\|^2] \right. \\
 &\quad \left. + \frac{150}{71}\gamma^2\tau^2 \sum_{j=r(t)}^m \mathbb{E} [\|\nabla f(x_j)\|^2] + \frac{150}{71}\gamma^2\tau^2\sigma^2 + \frac{12}{71}\gamma^2(\tau_C - 1)\sigma^2 \right) \\
 &= 4\mathbb{E} \left[\tilde{\phi}_t^m(x_{r(t)}) \right] + \frac{48}{7}L\gamma \sum_{j=r(t)}^m \mathbb{E} [\tilde{\phi}_t^{j-1}(x_{r(t)})] + \frac{6}{71}(\tau_C - 1)^2 G^2 \\
 &\quad + \frac{90}{7}L^2\tau \sum_{j=r(t)}^m \mathbb{E} [\|y_j - y_{\alpha_j}\|^2] + \frac{4}{71}\tau \sum_{j=r(t)}^m \|\nabla f(x_j)\|^2 + \frac{2}{35}\tau\sigma^2 + \frac{6}{71}(\tau_C - 1)\sigma^2,
 \end{aligned}$$

where we again use the fact that $L\gamma\tau \leq \frac{1}{30}$. □

Lemma D.4. If $30L\gamma\tau \leq 1$ and $30L\gamma\tau_C \leq 1$, then

$$\begin{aligned}
 \sum_{t=1}^T \mathbb{E} [\|y_t - \tilde{y}_t\|^2] &\leq \frac{44}{7}\gamma^2 \sum_{t=1}^T \mathbb{E} [\tilde{\phi}_t^{t-1}(x_{r(t)})] + \frac{6}{71}\gamma^2(\tau_C - 1)^2 G^2 T + \frac{1}{70} \sum_{t=1}^T \mathbb{E} [\|y_t - y_{\alpha_t}\|^2] \\
 &\quad + \frac{1}{15975L^2} \sum_{t=1}^T \mathbb{E} [\|\nabla f(x_t)\|^2] + \frac{2}{35}\gamma^2\tau\sigma^2 T + \frac{6}{71}\gamma^2(\tau_C - 1)\sigma^2. \tag{43}
 \end{aligned}$$

Proof. Using $L\gamma\tau \leq \frac{1}{30}$

$$\begin{aligned}
 \sum_{t=1}^T \mathbb{E} [\|y_t - \tilde{y}_t\|^2] &= \gamma^2 \sum_{t=1}^T \mathbb{E} \left[\left\| \sum_{j=r(t)}^{t-1} g_{k_j}(x_{\alpha_j}) - \nabla f(x_j) \right\|^2 \right] = \gamma^2 \sum_{t=1}^T \mathbb{E} [\|\tilde{\Delta}_t^{t-1}\|^2] \\
 &\stackrel{(40)}{\leq} 4\gamma^2 \sum_{t=1}^T \mathbb{E} [\tilde{\phi}_t^{t-1}(x_{r(t)})] + \frac{6}{71}\gamma^2(\tau_C - 1)^2 G^2 T + \frac{48}{7}L\gamma^3 \sum_{t=1}^T \sum_{j=r(t)}^{t-1} \mathbb{E} [\tilde{\phi}_t^{j-1}(x_{r(t)})] \\
 &\quad + \frac{90}{7}L^2\gamma^2\tau \sum_{t=1}^T \sum_{j=r(t)}^m \mathbb{E} [\|y_j - y_{\alpha_j}\|^2] + \frac{4\gamma^2\tau}{71} \sum_{t=1}^T \sum_{j=r(t)}^{t-1} \mathbb{E} [\|\nabla f(x_j)\|^2] + \frac{2}{35}\gamma^2\tau\sigma^2 T \\
 &\quad + \frac{6}{71}\gamma^2(\tau_C - 1)\sigma^2 \\
 &\leq \frac{44}{7}\gamma^2 \sum_{t=1}^T \mathbb{E} [\tilde{\phi}_t^{t-1}(x_{r(t)})] + \frac{6}{71}\gamma^2(\tau_C - 1)^2 G^2 T + \frac{90}{7}L^2\gamma^2\tau \sum_{t=1}^T \sum_{j=r(t)}^m \mathbb{E} [\|y_j - y_{\alpha_j}\|^2] \\
 &\quad + \frac{4\gamma^2\tau}{71} \sum_{t=1}^T \sum_{j=r(t)}^{t-1} \mathbb{E} [\|\nabla f(x_j)\|^2] + \frac{2}{35}\gamma^2\tau\sigma^2 T + \frac{6}{71}\gamma^2(\tau_C - 1)\sigma^2.
 \end{aligned}$$

We continue as follows

$$\begin{aligned}
 \sum_{t=1}^T \mathbb{E} [\|y_t - \tilde{y}_t\|^2] &\leq \frac{44}{7}\gamma^2 \sum_{t=1}^T \mathbb{E} [\tilde{\phi}_t^{t-1}(x_{r(t)})] + \frac{6}{71}\gamma^2(\tau_C - 1)^2 G^2 T + \frac{1}{70} \sum_{t=1}^T \mathbb{E} [\|y_t - y_{\alpha_t}\|^2] \\
 &\quad + \frac{1}{15975L^2} \sum_{t=1}^T \mathbb{E} [\|\nabla f(x_t)\|^2] + \frac{2}{35}\gamma^2\tau\sigma^2 T + \frac{6}{71}\gamma^2(\tau_C - 1)\sigma^2.
 \end{aligned}$$

□

Lemma D.5. If $30L\gamma\tilde{\tau}_{\max} \leq 1$, then

$$\begin{aligned} \sum_{t=1}^T \mathbb{E} [\|y_t - y_{\alpha_t}\|^2] &\leq \frac{2}{99}\gamma^2(\tau_C - 1)^2 G^2 T + \frac{100}{33}\gamma^2 \sum_{t=1}^T \mathbb{E} \left[\left\| \sum_{j=\alpha_t}^{t-1} \nabla f_{k_t}(x_{\alpha_j}) - \nabla f(x_{\alpha_j}) \right\|^2 \right] \\ &\quad + \frac{1}{297L^2} \sum_{t=1}^T \mathbb{E} [\|\nabla f(x_t)\|^2] + \frac{\gamma}{15L}\sigma^2 T + \frac{2}{99}\gamma^2(\tau_C - 1)\sigma^2 T. \end{aligned} \quad (44)$$

Proof. We have

$$\begin{aligned} \mathbb{E} [\|y_t - y_{\alpha_t}\|^2] &= \gamma^2 \mathbb{E} \left[\left\| \sum_{j=\alpha_t}^{t-1} g_{k_j}(x_{\alpha_j}) \right\|^2 \right] \\ &\leq \gamma^2 \mathbb{E} \left[\left\| \sum_{j=\alpha_t}^{t-1} \nabla f_{k_j}(x_{\alpha_j}) \pm \nabla f(x_{\alpha_j}) \pm \nabla f(x_j) \right\|^2 \right] + \gamma^2 \tau_t \sigma^2 \\ &\leq 3\gamma^2 \mathbb{E} \left[\left\| \sum_{j=\alpha_t}^{t-1} \nabla f_{k_j}(x_{\alpha_j}) - \nabla f(x_{\alpha_j}) \right\|^2 \right] + 3\gamma^2 \mathbb{E} \left[\left\| \sum_{j=\alpha_t}^{t-1} \nabla f(x_{\alpha_j}) - \nabla f(x_j) \right\|^2 \right] \\ &\quad + 3\gamma^2 \mathbb{E} \left[\left\| \sum_{j=\alpha_t}^{t-1} \nabla f(x_j) \right\|^2 \right] + \gamma^2 \tau_t \sigma^2 \\ &\leq 3L^2\gamma^2\tilde{\tau}_t \sum_{j=\alpha_t}^{t-1} \mathbb{E} [\|x_{\alpha_j} - x_j\|^2] + 3\gamma^2 \mathbb{E} \left[\left\| \sum_{j=\alpha_t}^{t-1} \nabla f_{k_j}(x_{\alpha_j}) - \nabla f(x_{\alpha_j}) \right\|^2 \right] \\ &\quad + 3\gamma^2\tilde{\tau}_t \sum_{j=\alpha_t}^{t-1} \mathbb{E} [\|\nabla f(x_j)\|^2] + \gamma^2 \tau_t \sigma^2. \end{aligned}$$

Then we add summation over the entire iterates and count the number of times each term appears

$$\begin{aligned} \sum_{t=1}^T \mathbb{E} [\|y_t - y_{\alpha_t}\|^2] &\leq 3L^2\gamma^2\tilde{\tau}_{\max} \sum_{t=1}^T \sum_{j=\alpha_t}^{t-1} \mathbb{E} [\|x_{\alpha_j} - x_j\|^2] + 3\gamma^2 \sum_{t=1}^T \mathbb{E} \left[\left\| \sum_{j=\alpha_t}^{t-1} \nabla f_{k_t}(x_{\alpha_j}) - \nabla f(x_{\alpha_j}) \right\|^2 \right] \\ &\quad + 3\gamma^2\tilde{\tau}_{\max} \sum_{t=1}^T \sum_{j=\alpha_t}^{t-1} \mathbb{E} [\|\nabla f(x_j)\|^2] + \tilde{\tau}_{\text{avg}}\gamma^2\sigma^2 T \\ &\leq 3L^2\gamma^2\tilde{\tau}_{\max}^2 \sum_{t=1}^T \mathbb{E} [\|x_{\alpha_t} - x_t\|^2] + 3\gamma^2 \sum_{t=1}^T \mathbb{E} \left[\left\| \sum_{j=\alpha_t}^{t-1} \nabla f_{k_t}(x_{\alpha_j}) - \nabla f(x_{\alpha_j}) \right\|^2 \right] \\ &\quad + 3\gamma^2\tilde{\tau}_{\max}^2 \sum_{t=1}^T \mathbb{E} [\|\nabla f(x_t)\|^2] + \tilde{\tau}_{\text{avg}}\gamma^2\sigma^2 T \\ &\leq \frac{1}{300} \sum_{t=1}^T \mathbb{E} [\|x_{\alpha_t} - x_t\|^2] + 3\gamma^2 \sum_{t=1}^T \mathbb{E} \left[\left\| \sum_{j=\alpha_t}^{t-1} \nabla f_{k_t}(x_{\alpha_j}) - \nabla f(x_{\alpha_j}) \right\|^2 \right] \\ &\quad + \frac{1}{300L^2} \sum_{t=1}^T \mathbb{E} [\|\nabla f(x_t)\|^2] + \tilde{\tau}_{\text{avg}}\gamma^2\sigma^2 T, \end{aligned}$$

where in the second inequality each term in the double sum appears at most $\tilde{\tau}_{\max}$ times. Using Young's inequality we continue as follows

$$\begin{aligned}
 \sum_{t=1}^T \mathbb{E} [\|y_t - y_{\alpha_t}\|^2] &\leq \frac{1}{100} \sum_{t=1}^T \mathbb{E} [\|x_{\alpha_t} - y_{\alpha_t}\|^2 + \|y_{\alpha_t} - y_t\|^2 + \|x_t - y_t\|^2] + \tilde{\tau}_{\text{avg}} \gamma^2 \sigma^2 T \\
 &\quad + 3\gamma^2 \sum_{t=1}^T \mathbb{E} \left[\left\| \sum_{j=\alpha_t}^{t-1} \nabla f_{k_t}(x_{\alpha_j}) - \nabla f(x_{\alpha_j}) \right\|^2 \right] + \frac{1}{300L^2} \sum_{t=1}^T \mathbb{E} [\|\nabla f(x_t)\|^2] \\
 &\leq \frac{1}{50} \gamma^2 T ((\tau_C - 1)^2 G^2 T + (\tau_C - 1) \sigma^2) + \frac{1}{100} \sum_{t=1}^T \mathbb{E} [\|y_{\alpha_t} - y_t\|^2] + \tilde{\tau}_{\text{avg}} \gamma^2 \sigma^2 T \\
 &\quad + 3\gamma^2 \sum_{t=1}^T \mathbb{E} \left[\left\| \sum_{j=\alpha_t}^{t-1} \nabla f_{k_t}(x_{\alpha_j}) - \nabla f(x_{\alpha_j}) \right\|^2 \right] + \frac{1}{300L^2} \sum_{t=1}^T \mathbb{E} [\|\nabla f(x_t)\|^2],
 \end{aligned}$$

where we again use Lemma D.2. After cancellation, we get

$$\begin{aligned}
 \sum_{t=1}^T \mathbb{E} [\|y_t - y_{\alpha_t}\|^2] &\leq \frac{2}{99} \gamma^2 (\tau_C - 1)^2 G^2 T + \frac{100}{33} \gamma^2 \sum_{t=1}^T \mathbb{E} \left[\left\| \sum_{j=\alpha_t}^{t-1} \nabla f_{k_t}(x_{\alpha_j}) - \nabla f(x_{\alpha_j}) \right\|^2 \right] \\
 &\quad + \frac{1}{297L^2} \sum_{t=1}^T \mathbb{E} [\|\nabla f(x_t)\|^2] + \frac{\gamma}{15L} \sigma^2 T + \frac{2}{99} \gamma^2 (\tau_C - 1) \sigma^2 T,
 \end{aligned}$$

where we use the stepsize restriction. \square

We can combine all previous lemmas into one.

Lemma D.6. If $30\gamma L\tau \leq 1$ and $30\gamma L \max\{\tilde{\tau}_{\max}, \tau_C\} \leq 1$ hold, then

$$\begin{aligned}
 \sum_{t=1}^T \mathbb{E} [\|y_t - \tilde{y}_t\|^2] &\leq 8\gamma^2 \tilde{\Phi} + \gamma^2 \tilde{\Psi} + \frac{1}{10} \gamma^2 (\tau_C - 1)^2 G^2 T + \frac{1}{10} \gamma^2 (\tau_C - 1)^2 G^2 T + \frac{1}{9033L^2} B \\
 &\quad + \frac{1}{525L} \gamma \sigma^2 T.
 \end{aligned} \tag{45}$$

Proof. Summary of obtained inequalities

$$\begin{aligned}
 \sum_{t=1}^T \mathbb{E} [\|y_t - \tilde{y}_t\|^2] &\stackrel{\text{Lemma D.4}}{\leq} \frac{44}{7} \gamma^2 \tilde{\Phi} + \frac{6}{71} \gamma^2 (\tau_C - 1)^2 G^2 T + \frac{1}{70} \tilde{A} + \frac{1}{15975L^2} B + \frac{2}{35} \gamma^2 \tau \sigma^2 T \\
 &\quad + \frac{6}{71} \gamma^2 (\tau_C - 1) \sigma^2 T, \\
 \tilde{A} &\stackrel{\text{Lemma D.5}}{\leq} \frac{2}{99} \gamma^2 (\tau_C - 1)^2 G^2 T + \frac{2}{99} \gamma^2 (\tau_C - 1) \sigma^2 T + \frac{100}{33} \gamma^2 \tilde{\Psi} + \frac{1}{297L^2} B \\
 &\quad + \frac{\gamma}{15L} \sigma^2 T.
 \end{aligned}$$

Hence, using the stepsize restriction $30L\gamma L\tau_C \leq 1$

$$\begin{aligned}
 \sum_{t=1}^T \mathbb{E} [\|y_t - \tilde{y}_t\|^2] &\leq \frac{44}{7} \gamma^2 \tilde{\Phi} + \frac{6}{71} \gamma^2 (\tau_C - 1)^2 G^2 T + \frac{1}{15975L^2} B + \frac{1}{525L} \gamma \sigma^2 T + \frac{6}{71} \gamma^2 (\tau_C - 1) \sigma^2 T \\
 &\quad + \frac{1}{70} \left(\frac{2}{99} \gamma^2 (\tau_C - 1)^2 G^2 T + \frac{2}{99} \gamma^2 (\tau_C - 1) \sigma^2 T + \frac{100}{33} \gamma^2 \tilde{\Psi} + \frac{1}{297L^2} B + \frac{\gamma}{15L} \sigma^2 T \right) \\
 &\leq 8\gamma^2 \tilde{\Phi} + \frac{1}{10} \gamma^2 (\tau_C - 1)^2 G^2 T + \frac{1}{10} \gamma^2 (\tau_C - 1) \sigma^2 T + \frac{1}{9033L^2} B + \frac{1}{525L} \gamma \sigma^2 T + \frac{10}{231} \gamma^2 \tilde{\Psi}.
 \end{aligned}$$

\square

D.2 Proof of Theorem 2

Theorem 2. *Let Assumptions 1, 2, and 4 hold. Let the stepsize γ satisfies inequalities $6L\gamma \leq 1$ and $30L\gamma \max\{\tilde{\tau}_{\max}, \tau_C\} \leq 1$, the correlation period $\tau = \lfloor \frac{1}{30L\gamma} \rfloor$, quantities $\{\tilde{\sigma}_{k,\tau}^2\}_{k=0}^{\lfloor T/\tau \rfloor}$ and $\tilde{\nu}^2$ are finite. Then*

$$\begin{aligned} \mathbb{E} [\|\nabla f(\hat{x}_T)\|^2] &\leq \mathcal{O}\left(\frac{F_1}{\gamma T} + L\gamma\sigma^2 + L^2\gamma^2\tilde{\Phi}\right. \\ &\quad \left.+ L^2\gamma^2(\tau_C - 1)^2G^2\right), \end{aligned} \quad (15)$$

where $F_1 := f(y_1) - f^*$, $\tilde{\Phi} := \frac{1}{\lfloor T/\tau \rfloor} \sum_{k=0}^{\lfloor T/\tau \rfloor} \tilde{\sigma}_{k,\tau}^2 + \frac{1}{T}\tilde{\nu}^2$, and \hat{x}_T is chosen uniformly at random from $\{x_1, \dots, x_T\}$.

Proof. Again, we analyze separately iterations of restarts and without.

Iterations without restart: If restarts do not happen, namely $(t+1) \bmod \tau \neq 0$ then from the smoothness assumption of f and \tilde{y}_t update rule, we have

$$\begin{aligned} f(\tilde{y}_{t+1}) &\leq f(\tilde{y}_t) - \gamma \langle \nabla f(\tilde{y}_t), \nabla f(x_t) \rangle + \frac{L\gamma^2}{2} \|\nabla f(x_t)\|^2 \\ &= f(\tilde{y}_t) - \frac{\gamma}{2} \|\nabla f(\tilde{y}_t)\|^2 - \frac{\gamma}{2} \|\nabla f(x_t)\|^2 + \frac{\gamma}{2} \|\nabla f(\tilde{y}_t) - \nabla f(x_t)\|^2 + \frac{L\gamma^2}{2} \|\nabla f(x_t)\|^2 \\ &\leq f(\tilde{y}_t) - \frac{\gamma}{2} \|\nabla f(\tilde{y}_t)\|^2 - \frac{\gamma}{2} \|\nabla f(x_t)\|^2 + \frac{L^2\gamma}{2} \|\tilde{y}_t - x_t\|^2 + \frac{L\gamma^2}{2} \|\nabla f(x_t)\|^2 \\ &\leq f(\tilde{y}_t) - \frac{\gamma}{2} \|\nabla f(\tilde{y}_t)\|^2 - \frac{\gamma}{3} \|\nabla f(x_t)\|^2 + L^2\gamma \|\tilde{y}_t - y_t\|^2 + L^2\gamma \|x_t - y_t\|^2. \end{aligned}$$

Iterations with restart: If a restart happens, namely $(t+1) \bmod \tau = 0$ then

$$\begin{aligned} \tilde{y}_{t+1} &= y_{t+1} = y_t - \gamma g_{k_t}(x_{\alpha_t}) \\ &= \tilde{y}_t + (y_t - \tilde{y}_t) - \gamma \nabla f(x_t) + (\gamma \nabla f(x_t) - \gamma g_{k_t}(x_{\alpha_t})) \\ &= \tilde{y}_t - \gamma \nabla f(x_t) + \underbrace{\gamma \sum_{j=r(t)}^t \nabla f(x_j) - g_{k_j}(x_{\alpha_j})}_{=\tilde{\Delta}_t^t}. \end{aligned}$$

Then we use smoothness of f to get

$$\begin{aligned} &f(\tilde{y}_{t+1}) \\ &\leq f(\tilde{y}_t) - \gamma \langle \nabla f(\tilde{y}_t), \nabla f(x_t) - \tilde{\Delta}_t^t \rangle + \frac{L\gamma^2}{2} \|\nabla f(x_t) - \tilde{\Delta}_t^t\|^2 \\ &\leq f(\tilde{y}_t) - \gamma \langle \nabla f(\tilde{y}_t), \nabla f(x_t) \rangle + \gamma \langle \nabla f(\tilde{y}_t), \tilde{\Delta}_t^t \rangle + L\gamma^2 \|\nabla f(x_t)\|^2 + L\gamma^2 \|\tilde{\Delta}_t^t\|^2 \\ &\leq f(\tilde{y}_t) - \frac{\gamma}{2} \|\nabla f(\tilde{y}_t)\|^2 - \frac{\gamma}{2} \|\nabla f(x_t)\|^2 + \frac{\gamma}{2} \|\nabla f(\tilde{y}_t) - \nabla f(x_t)\|^2 \\ &\quad + \frac{1}{240L} \|\nabla f(\tilde{y}_t)\|^2 + 60L\gamma^2 \|\tilde{\Delta}_t^t\|^2 + L\gamma^2 \|\nabla f(x_t)\|^2 + L\gamma^2 \|\tilde{\Delta}_t^t\|^2 \\ &\leq f(\tilde{y}_t) - \frac{\gamma}{2} \|\nabla f(\tilde{y}_t)\|^2 - \frac{\gamma}{3} \|\nabla f(x_t)\|^2 + \frac{L^2\gamma}{2} \|\tilde{y}_t - x_t\|^2 + \frac{1}{240L} \|\nabla f(\tilde{y}_t)\|^2 + 61L\gamma^2 \|\tilde{\Delta}_t^t\|^2 \\ &\leq f(\tilde{y}_t) - \frac{\gamma}{2} \|\nabla f(\tilde{y}_t)\|^2 - \frac{\gamma}{3} \|\nabla f(x_t)\|^2 + L^2\gamma \|\tilde{y}_t - y_t\|^2 + L^2\gamma \|x_t - y_t\|^2 + \frac{1}{160L} \|\nabla f(\tilde{y}_t)\|^2 \\ &\quad + 61L\gamma^2 \|\tilde{\Delta}_t^t\|^2, \end{aligned}$$

where in the third and fourth inequality Young's inequality is used.

If we denote by ξ_t the indicator function of restart event at $t+1$, namely $\xi_{k\tau-1} = 1$ for all $k \geq 1$ and is 0

otherwise, then we can take expectation and unify the descent inequality for both cases as follows:

$$\begin{aligned} \mathbb{E}[f(\tilde{y}_{t+1})] &\leq \mathbb{E}[f(\tilde{y}_t)] - \frac{\gamma}{2}\mathbb{E}[\|\nabla f(\tilde{y}_t)\|^2] - \frac{\gamma}{3}\mathbb{E}[\|\nabla f(x_t)\|^2] + L^2\gamma\mathbb{E}[\|\tilde{y}_t - y_t\|^2] \\ &\quad + L^2\gamma\mathbb{E}[\|y_t - x_t\|^2] + \left(\frac{1}{240L}\mathbb{E}[\|\nabla f(\tilde{y}_t)\|^2] + 61L\gamma^2\mathbb{E}[\|\tilde{\Delta}_t^t\|^2]\right)\xi_t \end{aligned} \quad (46)$$

for all $t \geq 0$.

Next, we apply summation over the entire iterates and bound the terms that appear only in every τ iteration.

$$\begin{aligned} \frac{1}{L}\mathbb{E}[\|\nabla f(\tilde{y}_t)\|^2] &= \frac{1}{L\tau}\sum_{j=0}^{\tau-1}\mathbb{E}[\|\nabla f(\tilde{y}_t)\|^2] \\ &\leq \frac{2}{L\tau}\sum_{j=0}^{\tau-1}\mathbb{E}[\|\nabla f(\tilde{y}_t) - \nabla f(\tilde{y}_{t-j})\|^2] + \frac{2}{L\tau}\sum_{j=0}^{\tau-1}\mathbb{E}[\|\nabla f(\tilde{y}_{t-j})\|^2] \\ &\leq \frac{2L}{\tau}\sum_{j=0}^{\tau-1}\mathbb{E}[\|\tilde{y}_t - \tilde{y}_{t-j}\|^2] + \frac{2}{L\tau}\sum_{j=0}^{\tau-1}\mathbb{E}[\|\nabla f(\tilde{y}_{t-j})\|^2] \\ &\leq \frac{2L\gamma^2}{\tau}\sum_{j=0}^{\tau-1}\mathbb{E}\left[\left\|\sum_{l=t-j}^{t-1}\nabla f(x_l)\right\|^2\right] + \frac{2}{L\tau}\sum_{j=0}^{\tau-1}\mathbb{E}[\|\nabla f(\tilde{y}_{t-j})\|^2] \\ &\leq 2L\gamma^2\sum_{j=0}^{\tau-1}\sum_{l=t-j}^{t-1}\mathbb{E}[\|\nabla f(x_l)\|^2] + \frac{2}{L\tau}\sum_{j=0}^{\tau-1}\mathbb{E}[\|\nabla f(\tilde{y}_{t-j})\|^2] \\ &\leq 2L\gamma^2\tau\sum_{j=0}^{\tau-1}\mathbb{E}[\|\nabla f(x_{t-j})\|^2] + \frac{2}{L\tau}\sum_{j=0}^{\tau-1}\mathbb{E}[\|\nabla f(\tilde{y}_{t-j})\|^2] \\ &\leq \frac{\gamma}{15}\sum_{j=0}^{\tau-1}\mathbb{E}[\|\nabla f(x_{t-j})\|^2] + 120\gamma\sum_{j=0}^{\tau-1}\mathbb{E}[\|\nabla f(\tilde{y}_{t-j})\|^2] \end{aligned}$$

provided that $\frac{1}{60} \leq L\gamma\tau \leq \frac{1}{30}$ (e.g., $\tau = \lfloor \frac{1}{30L\gamma} \rfloor$). Then we can use this bound to derive

$$\sum_{t=1}^T \frac{1}{240L}\mathbb{E}[\|\nabla f(\tilde{y}_t)\|^2]\xi_t \leq \frac{\gamma}{3600}\sum_{t=1}^T\mathbb{E}[\|\nabla f(x_t)\|^2] + \frac{\gamma}{2}\sum_{t=1}^T\mathbb{E}[\|\nabla f(\tilde{y}_t)\|^2]. \quad (47)$$

Then we use (40) to bound $\mathbb{E}[\|\Delta_t^t\|^2]$:

$$\begin{aligned} L\gamma^2\sum_{t=1}^T\mathbb{E}[\|\tilde{\Delta}_t^t\|^2]\xi_t &\leq 4L\gamma^2\sum_{t=1}^T\mathbb{E}[\tilde{\phi}_t^t(x_{r(t)})]\xi_t + \sum_{t=1}^T\frac{6L\gamma^2}{71}(\tau_C - 1)^2G^2\xi_t \\ &\quad + \sum_{t=1}^T\frac{6L\gamma^2}{71}(\tau_C - 1)\sigma^2\xi_t + \frac{90}{7}L^3\gamma^2\tau\sum_{t=1}^T\sum_{j=r(t)}^t\mathbb{E}[\|y_j - y_{\alpha_j}\|^2]\xi_t \\ &\quad + \frac{48}{7}L^2\gamma^3\sum_{t=1}^T\sum_{j=r(t)}^{t-1}\mathbb{E}[\tilde{\phi}_t^{j-1}(x_{r(t)})]\xi_t \\ &\quad + \frac{4L\gamma^2\tau}{71}\sum_{t=1}^T\sum_{j=r(t)}^t\mathbb{E}[\|\nabla f(x_t)\|^2]\xi_t + \frac{2}{35}L\gamma^2\tau\sigma^2\frac{T}{\tau}. \end{aligned}$$

Notice that summation over the entire iterates with weights ξ_t is equivalent to division by τ .

$$\begin{aligned} \sum_{t=1}^T \sum_{j=r(t)}^t \mathbb{E} \left[\tilde{\phi}_t^{j-1}(x_{r(t)}) \right] \xi_t &= \sum_{t=1}^T \sum_{j=r(t)}^{t-1} \mathbb{E} \left[\tilde{\phi}_t^j(x_{r(t)}) \right] \xi_t = \sum_{t=1}^T \sum_{j=r(t)}^{t-1} \mathbb{E} \left[\tilde{\phi}_j^j(x_{r(j)}) \right] \xi_t \\ &\leq \sum_{t=1}^T \mathbb{E} \left[\tilde{\phi}_t^t(x_{r(t)}) \right]. \end{aligned}$$

Hence,

$$\begin{aligned} &61L\gamma^2 \sum_{t=1}^T \mathbb{E} \left[\|\tilde{\Delta}_t^t\|^2 \right] \xi_t \\ &\leq 15100L^2\gamma^3 \sum_{t=1}^T \mathbb{E} \left[\tilde{\phi}_t^t(x_{r(t)}) \right] + 6L\gamma^2 \frac{T}{\tau} (\tau_C - 1)^2 G^2 + 6L\gamma^2 \frac{T}{\tau} (\tau_C - 1) \sigma^2 \\ &\quad + \frac{2}{35} L\gamma^2 \sigma^2 T + \frac{183}{7} L^2\gamma \sum_{t=1}^T \mathbb{E} [\|y_t - y_{\alpha_t}\|^2] + \frac{244}{71} L\gamma^2 \tau \sum_{t=1}^T \mathbb{E} [\|\nabla f(x_t)\|^2] \\ \stackrel{\text{Lemma D.5}}{\leq} &15100L^2\gamma^3 \sum_{t=1}^T \mathbb{E} \left[\tilde{\phi}_t^t(x_{r(t)}) \right] + 360L^2\gamma^3 (\tau_C - 1)^2 G^2 T + 360L^2\gamma^3 (\tau_C - 1) \sigma^2 T \\ &\quad + \frac{2}{35} L\gamma^2 \sigma^2 T + \frac{183}{7} L^2\gamma \left(\frac{2}{99} \gamma^2 (\tau_C - 1)^2 G^2 T + \frac{2}{99} \gamma^2 (\tau_C - 1) \sigma^2 T + \frac{100}{33} \gamma^2 \tilde{\Psi} \right. \\ &\quad \left. + \frac{1}{297L^2} B + \frac{\gamma}{15L} \sigma^2 T \right) + \frac{122}{1065} \gamma B \\ &\leq 15100L^2\gamma^3 \sum_{t=1}^T \mathbb{E} \left[\tilde{\phi}_t^t(x_{r(t)}) \right] + 361L^2\gamma^3 (\tau_C - 1)^2 G^2 T + 361L^2\gamma^3 (\tau_C - 1) \sigma^2 T \\ &\quad + \frac{9}{5} L\gamma^2 \sigma^2 T + \frac{49837}{246015} \gamma B + 80L^2\gamma^3 \tilde{\Psi}. \end{aligned} \tag{48}$$

Plugging (47) and (48), lemmas D.2 and D.4 into (46) and adding summation, we have

$$\begin{aligned} \mathbb{E} [f(\tilde{y}_{T+1}) - f(\tilde{y}_1)] &\leq -\frac{\gamma}{2} \sum_{t=1}^T \mathbb{E} [\|\nabla f(\tilde{y}_t)\|^2] - \frac{\gamma}{3} \sum_{t=1}^T \mathbb{E} [\|\nabla f(x_t)\|^2] \\ &\quad + L^2\gamma \sum_{t=1}^T \mathbb{E} [\|\tilde{y}_t - y_t\|^2] + L^2\gamma \sum_{t=1}^T \mathbb{E} [\|x_t - y_t\|^2] \\ &\quad + \frac{1}{240L} \sum_{t=1}^T \xi_t \mathbb{E} [\|\nabla f(\tilde{x}_t)\|^2] + 61L\gamma^2 \sum_{t=1}^T \xi_t \mathbb{E} [\|\tilde{\Delta}_t^t\|^2] \\ &\leq -\frac{\gamma}{2} \sum_{t=1}^T \mathbb{E} [\|\nabla f(\tilde{y}_t)\|^2] - \frac{\gamma}{3} B + \frac{\gamma}{3600} B + \frac{\gamma}{2} \sum_{t=1}^T \mathbb{E} [\|\nabla f(\tilde{y}_t)\|^2] \\ &\quad + L^2\gamma \left(\frac{44}{7} \gamma^2 \tilde{\Phi} + \frac{6}{71} \gamma^2 (\tau_C - 1)^2 G^2 T + \frac{1}{70} \tilde{A} + \frac{1}{15975L^2} B + \frac{2}{35} \gamma^2 \tau \sigma^2 T \right. \\ &\quad \left. + \frac{6}{71} \gamma^2 (\tau_C - 1) \sigma^2 T \right) + L^2\gamma^3 (\tau_C - 1)^2 G^2 T + L^2\gamma^3 (\tau_C - 1) \sigma^2 T \\ &\quad + 15100L^2\gamma^3 \sum_{t=1}^T \mathbb{E} \left[\tilde{\phi}_t^t(x_{r(t)}) \right] + 361L^2\gamma^3 (\tau_C - 1)^2 G^2 T \\ &\quad + 361L^2\gamma^3 (\tau_C - 1) \sigma^2 T + 2L\gamma^2 \sigma^2 T + 80L^2\gamma^3 \tilde{\Psi} + \frac{49837}{246015} \gamma B. \end{aligned}$$

Algorithm 4 Random Asynchronous SGD

- 1: **Input:** $x_0 \in \mathbb{R}^d$, stepsize $\gamma > 0$, set of assigned jobs $\mathcal{A}_0 = \emptyset$, set of received jobs $\mathcal{R}_0 = \emptyset$
 - 2: **Initialization:** for all jobs $(i, 0) \in \mathcal{A}_1$, server assigns worker i to compute a stochastic gradient $g_i(x_0)$
 - 3: **for** $t = 0, 1, 2, \dots, T - 1$ **do**
 - 4: once worker i_t finishes a job $(i_t, \pi_t) \in \mathcal{A}_{t+1}$, it sends $g_{i_t}(x_{\pi_t})$ to the server
 - 5: server updates the current model $x_{t+1} = x_t - \gamma g_{i_t}(x_{\pi_t})$ and the set $\mathcal{R}_{t+1} = \mathcal{R}_t \cup \{(i_t, \pi_t)\}$
 - 6: server assigns worker $k_{t+1} \sim \text{Uni}[1, \dots, n]$ to compute a gradient $g_{k_{t+1}}(x_{t+1})$ \triangleleft i.e., $\alpha_t \equiv t$
 - 7: server updates the set $\mathcal{A}_{t+2} = \mathcal{A}_{t+1} \cup \{(k_{t+1}, t + 1)\}$
 - 8: **end for**
-

Rearranging terms we have

$$\begin{aligned}
 \mathbb{E}[f(\tilde{y}_{T+1}) - f(\tilde{y}_1)] &\leq -\frac{\gamma}{3}B + \frac{\gamma}{3600}B + \frac{\gamma}{15975}B + \frac{49837\gamma}{246015}B + 4L\gamma^2\sigma^2T \\
 &\quad + 364L^2\gamma^3(\tau_C - 1)^2G^2T + 364L^2\gamma^3(\tau_C - 1)\sigma^2T \\
 &\quad + 8L^2\gamma^3\tilde{\Phi} + 81L^2\gamma^3\tilde{\Psi} + \frac{\gamma}{20790}B + 15100L^2\gamma^3 \sum_{t=1}^T \mathbb{E} \left[\tilde{\phi}_t^t(x_{r(t)}) \right] \\
 &\leq -\frac{\gamma}{7}B + 4L\gamma^2\sigma^2T + 8L^2\gamma^3\tilde{\Phi} + 81L^2\gamma^3\tilde{\Psi} \\
 &\quad + 15100L^2\gamma^3 \sum_{t=1}^T \mathbb{E} \left[\tilde{\phi}_t^t(x_{r(t)}) \right] + 364L^2\gamma^3(\tau_C - 1)^2G^2T \\
 &\quad + 364L^2\gamma^3(\tau_C - 1)\sigma^2T.
 \end{aligned}$$

Using the same argument as in the proof of Theorem 1, the above leads to

$$\begin{aligned}
 \frac{1}{T} \sum_{t=1}^T \mathbb{E} [\|\nabla f(x_t)\|^2] &\leq \frac{7(f(y_1) - f^*)}{\gamma T} + 2600L^2\gamma^2(\tau_C - 1)^2G^2 + 2600L\gamma\sigma^2 \\
 &\quad + 106000L^2\gamma^2 \left(\frac{1}{\lfloor T/\tau \rfloor} \sum_{k=0}^{\lfloor T/\tau \rfloor} \tilde{\sigma}_{k,\tau}^2 + \frac{1}{T} \tilde{\nu}^2 \right).
 \end{aligned}$$

□

D.3 Convergence Guarantees in Special Cases

D.3.1 Random Asynchronous SGD

The algorithm [Koloskova et al. \(2022\)](#) is almost identical to pure asynchronous SGD. The difference in the assigning process; the server chooses a new worker uniformly at random from all the workers regardless it is busy or not. Such choice allows to equalize the contribution from all workers. In this case, we always assign a new job at the last available model, i.e. $\tilde{\tau}_t \equiv 0$.

Now we derive the convergence guarantees for this algorithm.

Proposition D.1. Let Assumptions 1, 2, 3, and 4 hold. Let the stepsize satisfy $30L\tau_C\gamma \leq 1$, and $\tau = \lfloor \frac{1}{30L\gamma} \rfloor$. Then the iterates of Algorithm 4 satisfy

$$\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] \leq \mathcal{O} \left(\frac{F_1}{\gamma T} + L\gamma\sigma^2 + L\gamma\zeta^2 + L^2\tau_C^2\gamma^2G^2 \right), \tag{49}$$

where \hat{x}_t is chosen uniformly at random from $\{x_1, \dots, x_T\}$ and $F_1 = f(y_1) - f^*$. Moreover, if we tune the stepsize, then the iterates of random asynchronous SGD satisfy

$$\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] \leq \mathcal{O} \left(\frac{LF_1\tau_C}{T} + \left(\frac{LF_1\sigma^2}{T} \right)^{1/2} + \left(\frac{LF_1\zeta^2}{T} \right)^{1/2} + \left(\frac{F_1L\tau_C G}{T} \right)^{2/3} \right). \tag{50}$$

Proof. The proof is similar to that of pure asynchronous SGD. In this case we have

$$\sum_{t=k\tau}^{k\tau+j} (\nabla f_{k_t}(x)) = \sum_{i \in \mathcal{S}} (\nabla f_i(x) - \nabla f(x)),$$

where $|\mathcal{S}| \leq \tau$. Therefore, we can continue in the following way using the fact that k_t is always sampled independently

$$\begin{aligned} \mathbb{E} \left[\left\| \sum_{t=k\tau}^{k\tau+j} (\nabla f_{k_t}(x) - \nabla f(x)) \right\|^2 \right] &= \mathbb{E} \left[\left\| \sum_{i \in \mathcal{S}} (\nabla f_i(x) - \nabla f(x)) \right\|^2 \right] \\ &= \sum_{i \in \mathcal{S}} \mathbb{E} [\|\nabla f_i(x) - \nabla f(x)\|^2] \\ &\stackrel{\text{Asm. 3}}{\leq} |\mathcal{S}| \zeta^2 \leq \tau \zeta^2. \end{aligned}$$

This gives $\sigma_{k,\tau}^2 \leq \tau \zeta^2$ bound. Since $\alpha_t \equiv t$ in this case, then $\nu^2 = 0$. This means that the rate is given by

$$\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] \leq \mathcal{O} \left(\frac{F_1}{\gamma T} + L\gamma\sigma^2 + L\gamma\zeta^2 + L^2\tau_C^2\gamma^2 G^2 \right).$$

Now we need to tune the stepsize. Now we have three cases to consider. Indeed,

- if $\gamma = \Theta \left(\frac{1}{L\tau_C} \right)$, then

$$\begin{aligned} \mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] &\leq \mathcal{O} \left(\frac{F_1}{\frac{1}{L\tau_C} T} + L(\sigma^2 + \zeta^2) \left(\frac{F_1}{LT(\sigma^2 + \zeta^2)} \right)^{1/2} \right. \\ &\quad \left. + L^2\tau_C^2 G^2 \left(\frac{F_1}{L^2\tau_C^2 T G^2} \right)^{2/3} \right) \\ &\leq \mathcal{O} \left(\frac{LF_1\tau_C}{T} + \left(\frac{LF_1(\sigma^2 + \zeta^2)}{T} \right)^{1/2} + \left(\frac{F_1 L\tau_C G}{T} \right)^{2/3} \right). \end{aligned}$$

- if $\gamma = \Theta \left(\left(\frac{F_1}{LT\sigma^2} \right)^{1/2} \right)$, then

$$\begin{aligned} \mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] &\leq \mathcal{O} \left(\frac{F_1}{T} \left(\frac{LT\sigma^2}{F_1} \right)^{1/2} + L\sigma^2 \left(\frac{F_1}{LT\sigma^2} \right)^{1/2} \right. \\ &\quad \left. + L\zeta^2 \left(\frac{F_1}{LT\zeta^2} \right)^{1/2} + L^2\tau_C^2 G^2 \left(\frac{F_1}{L^2\tau_C^2 T G^2} \right)^{2/3} \right) \\ &\leq \mathcal{O} \left(\left(\frac{LF_1\sigma^2}{T} \right)^{1/2} + \left(\frac{LF_1\zeta^2}{T} \right)^{1/2} + \left(\frac{F_1 L\tau_C G}{T} \right)^{2/3} \right). \end{aligned}$$

- if $\gamma = \Theta \left(\left(\frac{F_1}{LT\zeta^2} \right)^{1/2} \right)$, then

$$\begin{aligned} \mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] &\leq \mathcal{O} \left(\frac{F_1}{T} \left(\frac{LT\zeta^2}{F_1} \right)^{1/2} + L\sigma^2 \left(\frac{F_1}{LT\sigma^2} \right)^{1/2} \right. \\ &\quad \left. + L\zeta^2 \left(\frac{F_1}{LT\zeta^2} \right)^{1/2} + L^2\tau_C^2 G^2 \left(\frac{F_1}{L^2\tau_C^2 T G^2} \right)^{2/3} \right) \\ &\leq \mathcal{O} \left(\left(\frac{LF_1\sigma^2}{T} \right)^{1/2} + \left(\frac{LF_1\zeta^2}{T} \right)^{1/2} + \left(\frac{F_1 L\tau_C G}{T} \right)^{2/3} \right). \end{aligned}$$

Algorithm 5 Random Asynchronous SGD with waiting

- 1: **Input:** $x_0 \in \mathbb{R}^d$, stepsize $\gamma > 0$, set of assigned jobs $\mathcal{A}_0 = \emptyset$, set of received jobs $\mathcal{R}_0 = \emptyset$, batch size $b \geq 1$, gradient estimator $g = 0$ and number of received gradients $r = 0$
 - 2: **Initialization:** for all jobs $(i, 0) \in \mathcal{A}_1$, server assigns worker i to compute a stochastic gradient $g_i(x_0)$
 - 3: **for** $t = 0, 1, 2, \dots, T - 1$ **do**
 - 4: server sets $g_{t,0} = 0$ and $\mathcal{R}_{t,0} = \mathcal{R}_t$
 - 5: **for** $j = 1, \dots, b$ **do**
 - 6: once worker $i_{t,j}$ finishes a job $(i_{t,j}, \pi_{t,j}) \in \mathcal{A}_{t+1}$, it sends $g_{i_{t,j}}(x_{\pi_{t,j}})$ to the server
 - 7: server updates $g_{t,j} = g_{t,j-1} + g_{i_{t,j}}(x_{\pi_{t,j}})$
 - 8: server updates the set $\mathcal{R}_{t,j} = \mathcal{R}_{t,j-1} \cup \{(i_{t,j}, \pi_{t,j})\}$
 - 9: **end for**
 - 10: server updates the current model $x_{t+1} = x_t - \frac{\gamma}{b}g_{t,b}$ and set $\mathcal{R}_{t+1} = \mathcal{R}_{t,b}$
 - 11: server samples new batch of workers $\{k_{t+1,1}, \dots, k_{t+1,b}\}$ of size b
 - 12: server assigns worker $k_{t+1,j}$ to compute a gradient $g_{k_{t+1,j}}(x_{t+1})$ for all $j \in [b]$ \triangleleft i.e., $\alpha_{t,j} = t$
 - 13: server updates the set $\mathcal{A}_{t+2} = \mathcal{A}_{t+1} \cup \{(k_{t+1,1}, t + 1)\} \cup \dots \cup \{(k_{t+1,b}, t + 1)\}$
 - 14: **end for**
-

- if $\gamma = \Theta\left(\left(\frac{F_1}{L^2\tau_C^2TG^2}\right)^{1/3}\right)$, then

$$\begin{aligned} \mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] &\leq \mathcal{O}\left(\frac{F_1}{T} \left(\frac{L^2\tau_C^2TG^2}{F_1}\right)^{1/3} + L\sigma^2 \left(\frac{F_1}{LT\sigma^2}\right)^{1/2}\right. \\ &\quad \left.+ L\zeta^2 \left(\frac{F_1}{LT\zeta^2}\right)^{1/2} + L^2\tau_C^2G^2 \left(\frac{F_1}{L^2\tau_C^2TG^2}\right)^{2/3}\right) \\ &\leq \mathcal{O}\left(\left(\frac{LF_1\sigma^2}{T}\right)^{1/2} + \left(\frac{LF_1\zeta^2}{T}\right)^{1/2} + \left(\frac{F_1L\tau_C G}{T}\right)^{2/3}\right). \end{aligned}$$

To get the final rate after stepsize tuning we need to set the minimal stepsize over all cases. □

D.3.2 Random asynchronous SGD with waiting

The algorithm we consider in this paragraph is a special case of FedBuff [Nguyen et al. \(2022\)](#) when we perform only one local step (because this is out of the scope of our work). The difference with the previous algorithm is that the server waits for the first b workers before performing one step, i.e. one step has a form

$$z_{q+1} = z_q - \tilde{\gamma} \sum_{k \in B_q} \nabla f_i(z_{\pi_{q,i}}), \quad (51)$$

where $\tilde{\gamma} = \frac{\gamma}{b}$. Next the server uniformly at random chooses new b workers $B_q = \{k_{q,0}, \dots, k_{q,b-1}\}$, and send them new jobs. The update of $\{y_t\}_{t=0}^{T-1}$

$$\begin{aligned} y_{qb+b} &= y_{qb+b-1} - \tilde{\gamma} \nabla f_{k_{q,b-1}}(x_{qb}) \\ &\quad \vdots \\ &= y_{qb} - \tilde{\gamma} \sum_{k \in B_q} \nabla f_k(x_{qb}), \end{aligned}$$

so that now it satisfies (13). As we can see, $\tilde{\tau}_t$ changes from 0 to $b - 1$ within one batch, i.e. $\tilde{\tau}_{\max} \leq b$. That is why we choose stepsize such that $30L\tilde{\gamma} \max\{b, \tau_C\} \leq 1$ and $6L\tilde{\gamma} \leq 1$, then all stepsize constraints are satisfied. Now we set $\tau = b \lfloor \frac{1}{30L\tilde{\gamma}} \rfloor$, so that τ is a multiple of b which is needed to take correctly conditional expectations. This gives the bound $30L\frac{\gamma}{b}\tau \leq 1$ we need in the proofs as well.

Proposition D.2. Let Assumptions 1, 2, 3, and 4 hold. Let the stepsize satisfy $30L\gamma \leq 1$, and $\tau = b \lfloor \frac{1}{30L\gamma} \rfloor$. Then the iterates of Algorithm 5 with waiting satisfy

$$\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] \leq \mathcal{O}\left(\frac{F_1}{\gamma T} + L\gamma \frac{\zeta^2 + \sigma^2}{b} + L^2\gamma^2 G^2 \frac{\tau_C^2}{b^2}\right), \quad (52)$$

Algorithm 6 Shuffled Asynchronous SGD

- 1: **Input:** $x_0 \in \mathbb{R}^d$, stepsize $\gamma > 0$, set of assigned jobs $\mathcal{A}_0 = \emptyset$, set of received jobs $\mathcal{R}_0 = \emptyset$, random perturbation of workers χ , worker counter $r = 0$
 - 2: **Initialization:** for all jobs $(i, 0) \in \mathcal{A}_1$, server assigns worker i to compute a stochastic gradient $g_i(x_0)$
 - 3: **for** $t = 0, 1, 2, \dots, T - 1$ **do**
 - 4: once worker i_t finishes a job $(i_t, \pi_t) \in \mathcal{A}_{t+1}$, it sends $g_{i_t}(x_{\pi_t})$ to the server
 - 5: server updates the current model $x_{t+1} = x_t - \gamma g_{i_t}(x_{\pi_t})$ and the set $\mathcal{R}_{t+1} = \mathcal{R}_t \cup \{(i_t, \pi_t)\}$
 - 6: server assigns worker $\chi(r)$ to compute a gradient $g_{\chi(r)}(x_{t+1})$ and updates $r \leftarrow r + 1$ \triangleleft i.e., $\alpha_t \equiv t$
 - 7: server updates the set $\mathcal{A}_{t+2} = \mathcal{A}_{t+1} \cup \{(\chi(r), t + 1)\}$
 - 8: **if** $r = n$ **then**
 - 9: server samples new perturbation of workers χ and set $r = 0$
 - 10: **end if**
 - 11: **end for**
-

where \hat{x}_t is chosen uniformly at random from $\{x_1, \dots, x_b, \dots, x_{2b}, \dots, x_{Tb}\}$ and $F_1 = f(y_1) - f^*$. Moreover, if we tune the stepsize, then the iterates of random asynchronous SGD with waiting satisfy

$$\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] \leq \mathcal{O} \left(\frac{LF_1\tau_C}{T} + \left(\frac{LF_1\sigma^2}{bT} \right)^{1/2} + \left(\frac{LF_1\zeta^2}{bT} \right)^{1/2} + \left(\frac{F_1L\tau_C G}{Tb} \right)^{2/3} \right). \quad (53)$$

Proof. We again start with stepsize restriction. The effective stepsize $\tilde{\gamma} = \frac{\gamma}{b}$ should satisfy $6L\tilde{\gamma} \leq 1$ and $30L\tilde{\gamma} \max\{b, \tau_C\} \leq 1$. This is implied if γ satisfies $30L\gamma \max\{b, \tau_C\} \leq b$. Next, we choose $\tau = b \lfloor \frac{1}{30L\gamma} \rfloor$. First, this choice allows to apply conditional expectation correctly. Second, it satisfies the restriction $30L\tilde{\gamma}\tau = 30L\frac{\gamma}{b}b \lfloor \frac{1}{30L\gamma} \rfloor \leq 1$.

Now we need to compute $\tilde{\sigma}_{k,\tau}^2$ and $\tilde{\nu}^2$ quantities. We start with the first one. Each chunk of size τ in this case consists of $m \geq 1$ full batches of size b . Every batch is independent of others, therefore we need to compute sequence correlation within one batch and then sum altogether. The sequence correlation within one batch is the same as in the case of mini-batch SGD since a batch is chosen uniformly at random, thus, $\tilde{\sigma}_{k,\tau}^2 \leq \tau\zeta^2$. $\tilde{\nu}^2$ term is similarly bounded by $Tb\zeta^2$ since for all delayed gradients within one batch the argument α_j is the same, and therefore we are able to take conditional expectation correctly. Thus, the rate we derive is

$$\begin{aligned} \frac{1}{Tb} \sum_{q=0}^{T-1} \sum_{l=0}^{b-1} \mathbb{E} [\|\nabla f(x_{qb+l})\|^2] &\leq \mathcal{O} \left(\frac{F_1}{\tilde{\gamma}Tb} + L\tilde{\gamma}\sigma^2 + L^2\tilde{\gamma}^2\tau\zeta^2 + L^2\tilde{\gamma}^2b\zeta^2 + L^2\tau_C^2\tilde{\gamma}^2G^2 \right) \\ &\leq \mathcal{O} \left(\frac{F_1}{\gamma T} + L\gamma \frac{\sigma^2 + \zeta^2}{b} + L^2\gamma^2G^2 \frac{\tau_C^2}{b^2} \right). \end{aligned}$$

It is left to tune the stepsize. We have three cases similar to the previous case, therefore we skip computations for that part. □

Remark. We observe that waiting for b workers improves the second and third terms. However, we pay for that by waiting time. In practice, there is a trade-off between the number of workers we need to wait and the convergence speed.

D.3.3 Shuffled Asynchronous SGD [NEW]

This method is inspired by the superiority of SGD with random reshuffling and shuffle once. In random asynchronous SGD workers might have different amounts of work depending on the random seed. However, we want to ensure that all workers have the same number of jobs in order to have a balance between all of them, but we still want to do it in random order. Therefore, we sample a permutation χ_q of $[n]$ at epoch q and then assign new jobs according to χ_q . We can sample permutation before every epoch, or just once in the beginning. Here we want to utilize all available resources, i.e. $\tau_C = n$.

Now we present the method formally. Before epoch q we sample a permutation χ_q of $[n]$. Then one epoch of shuffled asynchronous SGD has the following form

$$x_{qn+l+1} = x_{qn+l} - \gamma \nabla f_{i_{qn+l}}(x_{\pi_{qn+l}}) \quad \forall l \in [0, n-1], \quad (54)$$

where π_{qn+l} is the iteration counter of the model where the corresponding gradient has started to be computed. The sequence of virtual iterates $\{y_t\}_{t=0}^T$ is computed as follows

$$y_{qn+l+1} = y_{qn+l} - \gamma \nabla f_{\chi_q(l)}(x_{qn}). \quad (55)$$

Again, $\tilde{\tau}_t \equiv 0$, therefore, the stepsize restrictions are $30L\gamma n \leq 1$ and $6L\gamma \leq 1$. We choose $\tau = n \lfloor \frac{1}{30Ln\gamma} \rfloor$, so that τ is a multiple of n in order to take conditional expectations correctly. Note that we satisfy the condition $30L\gamma\tau \leq 30L\gamma n \frac{1}{30Ln\gamma} = 1$ as well. We are ready to apply Theorem 2.

Proposition D.3. Let Assumptions 1, 2, 3, and 4 hold. Let the stepsize satisfy $30L\gamma n \leq 1$, and $\tau = n \lfloor \frac{1}{30Ln\gamma} \rfloor$. Then the iterates of shuffled asynchronous SGD satisfy

$$\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] \leq \mathcal{O} \left(\frac{F_1}{\gamma T} + L\gamma\sigma^2 + L^2\gamma^2 n\zeta^2 + L^2\gamma^2 G^2\tau_C^2 \right), \quad (56)$$

where \hat{x}_t is chosen uniformly at random from $\{x_1, \dots, x_n, \dots, x_{2n}, \dots, x_{Tn}\}$ and $F_1 = f(y_1) - f^*$. Moreover, if the stepsize is set as $\gamma = \min \left\{ \frac{1}{30Ln}, \left(\frac{F_1}{L^2 n \zeta^2 T} \right)^{1/3}, \left(\frac{F_1}{L^2 n^2 G^2 \zeta^2 T} \right)^{1/3} \right\}$, then

$$\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] \leq \mathcal{O} \left(\frac{LnF_1}{T} + \left(\frac{LF_1\sigma^2}{T} \right)^{1/2} + \left(\frac{LF_1\sqrt{n}\zeta}{T} \right)^{2/3} + \left(\frac{LGnF_1}{T} \right)^{2/3} \right). \quad (57)$$

Proof. We have already shown that the choice of stepsize γ and period τ satisfy the conditions of Theorem 2. Now we need to bound $\tilde{\sigma}_{k,\tau}^2$ and $\tilde{\nu}^2$. This is done analogously to SGD with random reshuffling. We have that $\tilde{\sigma}_{k,\tau}^2 \leq n\zeta^2$, and $\tilde{\nu}^2 = 0$ since we do not have delays during assigning. Thus the convergence guarantees for this method are

$$\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] \leq \mathcal{O} \left(\frac{F_1}{\gamma T} + L\gamma\sigma^2 + L^2\gamma^2 n\zeta^2 + L^2\gamma^2 G^2 n^2 \right),$$

since $\tau_C = n$ in this case. We observe that both third term depends on γ^2 while the convergence of random asynchronous SGD the term with ζ^2 depends on γ only.

Similarly to the previous two cases, we can tune the stepsize. We skip this part as it is almost exactly the same. \square

Remark 1. Let us set $\tau_C = n$ in random asynchronous SGD [Koloskova et al. \(2022\)](#) in order to compare it with shuffled asynchronous SGD proposed in this work. The difference in the rates comes from $\left(\frac{LF_1\zeta^2}{T} \right)^{1/2}$ for random asynchronous SGD and $\left(\frac{LF_1\sqrt{n}\zeta}{T} \right)^{2/3}$ for shuffled asynchronous SGD. Both terms become dominate in highly heterogeneous regime, i.e. when ζ^2 is large. If we want to achieve $\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] \leq \varepsilon$, then random asynchronous SGD requires $\mathcal{O} \left(\frac{LF_1\zeta^2}{\varepsilon^2} \right)$ iterations while it is $\mathcal{O} \left(\frac{LF_1\sqrt{n}\zeta}{\varepsilon^{3/2}} \right)$ for the shuffled asynchronous SGD. This means that shuffled asynchronous SGD needs less iterations if $\zeta \geq \sqrt{n}\varepsilon^{1/2}$, i.e. if we consider strongly heterogeneous regime which is typically the case in Federated Learning.

D.3.4 Pure Asynchronous SGD

We consider exactly the same algorithm that was covered in Proposition C.1 but with an additional assumption of bounded gradients. In this case $\alpha_t \equiv t$, i.e. $\tilde{\nu}^2 = 0$. Similarly, we bound $\tilde{\sigma}_{k,\tau}^2 \leq \tau^2\zeta^2$. Applying Theorem 2 we derive the following convergence guarantees.

Proposition D.4. Let Assumption 1, 2, 3, and 4 hold. Let the stepsize γ satisfy conditions $30L\gamma\tau_C \leq 1$ and $6L\gamma \leq 1$. Let $\tau = \lfloor \frac{1}{30L\gamma} \rfloor$. Then the iterates of Algorithm 2 satisfy

$$\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] \leq \mathcal{O} \left(\frac{F_1}{\gamma T} + L^2\gamma^2\tau_C^2G^2 + \zeta^2 \right), \quad (58)$$

where \hat{x}_t is chosen uniformly at random from $\{x_1, \dots, x_T\}$. Moreover, if we tune the stepsize, then the convergence rate is

$$\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] \leq \mathcal{O} \left(\frac{L\tau_C F_1}{T} + \left(\frac{L\tau_C G F_1}{T} \right)^{2/3} + \zeta^2 \right). \quad (59)$$

Proof. Since $\tilde{\tau}_t \equiv 0$, then we have to satisfy two stepsize conditions that are indicated in the proposition statement. In fact, we only need to satisfy $30L\gamma\tau_C \leq 1$ as then $6L\gamma \leq 1$ will be automatically satisfied. Since we have $\tilde{\sigma}_{k,\tau}^2 \leq \tau^2\zeta^2$ and $\tilde{\nu}^2 = 0$, then applying Theorem 2 we get

$$\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] \leq \mathcal{O} \left(\frac{F_1}{\gamma T} + L^2\gamma^2\tau_C^2G^2 + L^2\gamma^2\tau^2\zeta^2 \right) \quad (60)$$

$$\leq \mathcal{O} \left(\frac{F_1}{\gamma T} + L^2\gamma^2\tau_C^2G^2 + \zeta^2 \right). \quad (61)$$

Now we switch to the tuning of the stepsize. We have two cases.

- if $\gamma = \Theta \left(\frac{1}{L\tau_C} \right)$, then

$$\begin{aligned} \mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] &\leq \mathcal{O} \left(\frac{F_1}{\frac{1}{30L\tau_C}T} + L^2\tau_C^2G^2 \left(\frac{F_1}{L^2\tau_C^2G^2T} \right)^{2/3} + \zeta^2 \right) \\ &\leq \mathcal{O} \left(\frac{L\tau_C F_1}{T} + \left(\frac{L\tau_C G F_1}{T} \right)^{2/3} + \zeta^2 \right). \end{aligned}$$

- if $\gamma = \Theta \left(\left(\frac{F_1}{L^2\tau_C^2G^2T} \right)^{1/3} \right)$, then

$$\begin{aligned} \mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] &\leq \mathcal{O} \left(\frac{F_1}{T} \left(\frac{L^2\tau_C^2G^2T}{F_1} \right)^{1/3} + L^2\tau_C^2G^2 \left(\frac{F_1}{L^2\tau_C^2G^2T} \right)^{2/3} + \zeta^2 \right) \\ &\leq \mathcal{O} \left(\left(\frac{L\tau_C G F_1}{T} \right)^{2/3} + \zeta^2 \right). \end{aligned}$$

It is left to take the minimum over two cases. We highlight that the rate does not depend on the maximum delay, it only depends on τ_C which is proportional to τ_{avg} (Koloskova et al., 2022). \square

D.3.5 Pure Asynchronous SGD with waiting

We consider the same algorithm from Proposition C.3, but with additional gradient bound assumption 4. If we apply Theorem 2 we get the following statement.

Proposition D.5. Let Assumption 1, 2, 3, and 4 hold. Let the stepsize γ satisfy conditions $30L\gamma\tau_C \leq 1$ and $6L\gamma \leq 1$. Let $\tau = \lfloor \frac{1}{30L\gamma} \rfloor$. Then the iterates of Algorithm 3 satisfy

$$\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] \leq \mathcal{O} \left(\frac{F_1}{\gamma T b} + L^2\gamma^2\tau_C^2G^2 + \zeta^2 \right), \quad (62)$$

where \hat{x}_t is chosen uniformly at random from $\{x_1, \dots, x_b, x_{b+1}, \dots, x_{2b}, \dots, x_{Tb}\}$. Moreover, if we tune the stepsize, then the convergence rate is

$$\mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] \leq \mathcal{O} \left(\frac{L\tau_C F_1}{T b} + \left(\frac{L\tau_C G F_1}{T b} \right)^{2/3} + \zeta^2 \right). \quad (63)$$

Proof. First, we start with the stepsize conditions. Note that when we rewrite the iterations of pure asynchronous SGD with waiting so that they suit (2), then we increase τ_t by b at most, but τ_C remains unchanged. This gives us even better improvement w.r.t. b . Indeed, we still need to satisfy $30L\gamma\tau_C \leq 1$ and $6L\gamma \leq 1$. If the first one holds, then the second one as well. Now we apply Theorem 2. Note that T iterations of pure asynchronous SGD with waiting are Tb iterations of Algorithm 1. Then we have

$$\begin{aligned} \mathbb{E} [\|\nabla f(\hat{x}_t)\|^2] &\leq \mathcal{O} \left(\frac{F_1}{\gamma Tb} + L^2\gamma^2\tau_C^2G^2 + L^2\gamma^2\tau^2\zeta^2 \right) \\ &\leq \mathcal{O} \left(\frac{F_1}{\gamma Tb} + L^2\gamma^2\tau_C^2G^2 + \zeta^2 \right). \end{aligned}$$

If we choose the stepsize $\gamma = \Theta \left(\min \left\{ \frac{1}{L\tau_C}, \left(\frac{F_1}{L^2\tau_C^2G^2Tb} \right)^{1/3} \right\} \right)$, then we get the second statement of the proposition. □