

A Unified Convergence Analysis for Shuffling-Type Gradient Methods

Lam M. Nguyen

*IBM Research, Thomas J. Watson Research Center
Yorktown Heights, NY10598, USA.*

LAMNGUYEN.MLTD@IBM.COM

Quoc Tran-Dinh

*Department of Statistics and Operations Research
The University of North Carolina at Chapel Hill
Chapel Hill, NC27599, USA.*

QUOCTD@EMAIL.UNC.EDU

Dzung T. Phan

*IBM Research, Thomas J. Watson Research Center
Yorktown Heights, NY10598, USA.*

PHANDU@US.IBM.COM

Phuong Ha Nguyen

eBay Inc., San Jose, CA 95125.

PHUONGHA.NTU@GMAIL.COM

Marten van Dijk

Centrum Wiskunde & Informatica, Amsterdam, Netherlands.

MARTEN.VAN.DIJK@CWI.NL

Editor: Prateek Jain

Abstract

In this paper, we propose a unified convergence analysis for a class of generic shuffling-type gradient methods for solving finite-sum optimization problems. Our analysis works with any sampling without replacement strategy and covers many known variants such as randomized reshuffling, deterministic or randomized single permutation, and cyclic and incremental gradient schemes. We focus on two different settings: strongly convex and nonconvex problems, but also discuss the non-strongly convex case. Our main contribution consists of new non-asymptotic and asymptotic convergence rates for a wide class of shuffling-type gradient methods in both nonconvex and convex settings. We also study uniformly randomized shuffling variants with different learning rates and model assumptions. While our rate in the nonconvex case is new and significantly improved over existing works under standard assumptions, the rate on the strongly convex one matches the existing best-known rates prior to this paper up to a constant factor without imposing a bounded gradient condition. Finally, we empirically illustrate our theoretical results via two numerical examples: nonconvex logistic regression and neural network training examples. As byproducts, our results suggest some appropriate choices for diminishing learning rates in certain shuffling variants.

Keywords: Stochastic gradient algorithm, shuffling-type gradient scheme, sampling without replacement, non-convex finite-sum minimization, strongly convex minimization

. Corresponding author: Lam M. Nguyen

1. Introduction

This paper aims at analyzing convergence rates of a general class of shuffling-type gradient methods for solving the following well-known finite sum minimization problem:

$$\min_{w \in \mathbb{R}^d} \left\{ F(w) := \frac{1}{n} \sum_{i=1}^n f(w; i) \right\}, \quad (\text{P})$$

where $f(\cdot; i) : \mathbb{R}^d \rightarrow \mathbb{R}$ is smooth and possibly nonconvex for $i \in [n] := \{1, \dots, n\}$.

Problem (P) covers a wide range of convex and nonconvex models in machine learning and statistical learning, including logistic regression, multi-kernel learning, conditional random fields, and neural networks. Especially, it covers *empirical risk minimization* as a special case. Very often, (P) lives in a high dimensional space, and/or it has a large number of components n . Therefore, deterministic optimization methods relying on full gradients are usually inefficient to solve (P), see, e.g., (Bottou et al., 2018; Sra et al., 2012).

The stochastic gradient descent (SGD) method has been widely used to solve (P) due to its efficiency in dealing with large-scale problems in big data regimes. The first variant of SGD (called *stochastic approximation method*) was introduced by (Robbins and Monro, 1951). In the last fifteen years, there has been a tremendous progress of research in SGD, where various stochastic and randomized-based algorithms have been proposed, making it be one of the most active research areas in optimization as well as in machine learning. In addition, due to the deep learning revolution, research on SGD for nonconvex optimization for deep learning also becomes extremely active nowadays.

SGD is also a method of choice to solve the following stochastic optimization problem:

$$\min_{w \in \mathbb{R}^d} \left\{ F(w) := \mathbb{E}_{(x,y) \sim \mathcal{D}} [f(w; x, y)] \right\}, \quad (1)$$

where \mathcal{D} is some probability distribution. Clearly, (P) can be cast into a special case of (1).

To solve (P), at each iteration k (for $k := 0, 1, \dots, K$), SGD chooses an index $i_k \in [n]$ (or a minibatch) at random and updates an iterate sequence $\{w_k\}$ as $w_{k+1} := w_k - \eta_k \nabla f(w_k; i_k)$ from a given starting point w_0 , which is up to n times ‘‘component gradient’’ cheaper than one iteration of a full gradient method with the update $w_{k+1} := w_k - \frac{\eta_k}{n} \sum_{i=1}^n \nabla f(w_k; i)$, where $\eta_k > 0$ is called a learning rate at the k -th iteration and $i_k = (x_k, y_k)$, a single sample or a mini-batch of the input data (x, y) . Although the first variant of SGD was introduced in the 1950s, its convergence rate was investigated much later, see, e.g., (Polyak and Juditsky, 1992). The convergence rate of SGD for solving (1) under strong convexity is $\mathcal{O}(K^{-1})$ (Nemirovski et al., 2009; Polyak and Juditsky, 1992; Nguyen et al., 2018), and for finding a stationary point of (1) in the nonconvex case is $\mathcal{O}(K^{-1/2})$ (Ghadimi and Lan, 2013), where K is the total iteration number. In particular, these rates also apply to (P).

Motivation. This paper is motivated by a number of observations as follows.

- Firstly, shuffling gradient-type methods are widely used in practice. The classical SGD scheme (we refer to it as the *standard SGD* method in this paper) for solving (P) relies on an i.i.d. sampling scheme to select components $\nabla f(\cdot; i)$ for updating the iterates w_k . In practice, however, other mechanisms for selecting components $\nabla f(\cdot; i)$ such as *randomized [re]shuffling techniques* are more desirable to use for implementing stochastic gradient algorithms. Shuffling strategies are easier and faster to implement in practice. They have

been implemented in several well-known packages such as TensorFlow and PyTorch. In addition, it has been recognized that shuffling-type methods often decrease the training loss faster than standard SGD, see, e.g., (Bottou, 2009, 2012; Hiroyuki, 2018). However, convergence guarantee of shuffling-type algorithms has just recently emerged. This motivates us to conduct a unified analysis for a general class of shuffling-type algorithms.

- Secondly, convergence analysis of shuffling-type schemes and cyclic strategies is much more challenging than that of the standard SGD or its variants due to the lack of independence. Hitherto, there is a limited number of theoretical works that analyze the convergence rates of shuffling techniques, where a majority has focussed on the strongly convex case (Gürbüzbalaban et al., 2019; HaoChen and Sra, 2019; Safran and Shamir, 2020; Nagaraj et al., 2019). To the best of our knowledge, hitherto and prior to our work, only (Li et al., 2020; Meng et al., 2019) have studied convergence rates of shuffling-type gradient methods for solving nonconvex instances of (P). There exists no unified analysis that can cover a wide class of shuffling-type gradient algorithms under different assumptions ranging from strongly convex to nonconvex cases. In this paper, we will provide some key elements to form a unified analysis framework for shuffling-type gradient methods, which can be applied to different variants. While we only focus on the strongly convex and nonconvex cases without specifying shuffling strategy, we believe that our framework can be customized to take into account additional or alternative assumptions to achieve a possibly better convergence rate (see, e.g., (Mishchenko et al., 2020) as an example).

- Thirdly, existing shuffling-type gradient algorithms for the strongly convex case of (P) such as (Ahn et al., 2020; Gürbüzbalaban et al., 2019; HaoChen and Sra, 2019; Safran and Shamir, 2020; Nagaraj et al., 2019) require a **bounded gradient assumption**. Although the bounded gradient condition is widely used and accepted for strongly convex problems, it leads to implicitly imposing a ball constraint on (P). More precisely, if F is μ -strongly convex and G -bounded gradient as $\sup_{i \in [n]} \|\nabla f(w; i)\| \leq G$ for $\forall w \in \text{dom}(F)$ (the domain of F), then it is easy to show that $\frac{\mu}{2} \|w - w^*\|^2 \leq F(w) - F(w^*) \leq \frac{1}{2\mu} \|\nabla F(w)\|^2 \leq \frac{G^2}{2\mu}$ for all $w \in \text{dom}(F)$, where w^* is the unique minimizer of F (Nguyen et al., 2018). This expression implies that $\|w - w^*\| \leq \frac{G}{\mu}$ as an implicit ball constraint on (P). However, we usually do not know w^* to quantify this condition in practice. Note that since F is strongly convex, it is also coercive. Thus any sublevel set $\{w \in \text{dom}(F) : F(w) \leq F(w_0)\}$ (for a fixed w_0) is bounded, and consequently, G exists. Nevertheless, it is still difficult to quantify G over this sublevel set since we do not know the size of this sublevel set explicitly. If G is quantified inappropriately, it will change problem (P) from the unconstrained to the constrained setting which may not be equivalent to (P). In addition, a projection is required to guarantee the feasibility of the iterates in this case, adding another computational cost, see, e.g., (Nagaraj et al., 2019). A recent work in (Mishchenko et al., 2020) can also avoid the bounded gradient condition in their analysis, but only focuses on the case where each component $f(\cdot; i)$ is convex.

- Fourthly, for the strongly convex case, prior to this work, existing analysis relies on a set of strong assumptions. Such assumptions often include: strong convexity, L -smoothness, Lipschitz Hessian, and bounded gradient. HaoChen and Sra prove an $\mathcal{O}((nT)^{-2} + T^{-3})$ convergence rate for randomized reshuffling scheme in (HaoChen and Sra, 2019), where T is the number of epochs (i.e. the number of passes over n component functions $f(\cdot; i)$), but under stronger assumptions than ours. Further improvement can be found in (Ahn et al.,

2020; Mishchenko et al., 2020). The rates in these papers nearly match the lower bound proven in (Safran and Shamir, 2020) or a sharper one in (Rajput et al., 2020). Our rate is $\mathcal{O}(T^{-2})$ in the general case, and $\mathcal{O}(1/(nT^2))$ in the randomized reshuffling case, but only requires $f(\cdot; i)$ to be L -smooth and F to be strongly convex with bounded variance. Hence, it is unclear if one can fairly compare ours and (HaoChen and Sra, 2019) since they consider two different classes of problems due to the use of different sets of assumptions.

- Finally, prior to our work, convergence analysis of shuffling-type gradient schemes has not been rigorously investigated for the nonconvex setting of (P). Existing works only focus on special variants such as incremental gradient or using nonstandard criterion (Li et al., 2020; Meng et al., 2019). Our work is the first analyzing convergence rates of the general shuffling-type gradient scheme in the nonconvex case under standard assumptions, which achieves the best known $\mathcal{O}(T^{-2/3})$ or $\mathcal{O}(n^{-1/3}T^{-2/3})$ rate in epoch.

Contribution. In this paper, we develop a new and unified convergence analysis framework for general shuffling-type gradient methods to solve (P) and apply it to different shuffling variants in both nonconvex and strongly convex settings under standard assumptions. More specifically, our contribution can be summarized as follows:

- (a) We prove $\mathcal{O}(1/T^2)$ -convergence rate in epoch of a generic shuffling-type gradient scheme for the strongly convex case without imposing “gradient boundedness” and/or Lipschitz Hessian assumptions, e.g., in (Gürbüzbalaban et al., 2019; HaoChen and Sra, 2019). In addition, our analysis does not require convexity of each component function as in some existing works. Similar to (Gürbüzbalaban et al., 2019) our rate also can be viewed as $\mathcal{O}(1/t^2)$ for any $1 \leq t \leq T$ without fixing T a priori as in other works.
- (b) If either a general bounded variance condition is imposed on F (see Assumption 2) or each component $f(\cdot; i)$ is convex for all $i \in [n]$, then by using a uniformly randomized reshuffling strategy, our convergence rate in the strongly convex case is improved to $\mathcal{O}(1/(nT^2))$, which matches the best-known results in recent works, including (Ahn et al., 2020; Mishchenko et al., 2020). Our latter case (i.e. when $f(\cdot; i)$ is convex) holds for both constant and diminishing learning rates.
- (c) We prove $\mathcal{O}(T^{-2/3})$ -convergence rate in epoch for the constant stepsizes and $\tilde{\mathcal{O}}(T^{-2/3})$ -convergence rate¹ for the diminishing stepsizes of a general shuffling-type gradient method (Algorithm 1) to approximate a stationary point of the nonconvex problem (P), where $T := K/n$ is number of epochs. Our rate is significantly improved over $\mathcal{O}(T^{-1/2})$ rate of the special incremental gradient method in (Li et al., 2020). To the best of our knowledge, these are the first improved non-asymptotic rates for SGD with shuffling for both constant and diminishing learning rates under standard assumptions. When a uniformly randomized reshuffling strategy is used, our rate is improved by a factor of $n^{1/3}$ to $\mathcal{O}(n^{-1/3}T^{-2/3})$.
- (d) We establish asymptotic convergence to a stationary point under diminishing learning rate scheme. We theoretically and empirically show that the shuffling-type gradient algorithm achieves the best performance with the learning rate $\eta_t = \mathcal{O}(t^{-1/3})$, where t is the epoch counter. Our learning rate closely relates to a “scheduled” one, i.e. it is constant within each epoch t and decreases w.r.t. t . When a uniformly randomized reshuffling strategy is used, our rate is also improved by a factor of $n^{1/3}$.

1. The notation $\tilde{\mathcal{O}}(\cdot)$ hides all logarithmic terms of the input (\cdot) compared to the standard $\mathcal{O}(\cdot)$ notation.

Comparison. Our theoretical results are new and different from existing works in several aspects. For the nonconvex case, (Meng et al., 2019) only proves $\mathcal{O}(1/\sqrt{nT} + \log(n)/n)$ -convergence rate to a neighborhood of a stationary point of a randomized reshuffling variant, (Li et al., 2020) shows $\mathcal{O}(T^{-1/2})$ convergence rate under the bounded subgradients and weak convexity conditions for an incremental subgradient variant. Our convergence rate is $\mathcal{O}(T^{-2/3})$ in epoch, and hence is $T^{-1/6}$ factor better than (Li et al., 2020) but using the smoothness of $f(\cdot; i)$ instead of a weak convexity condition. Compared to standard SGD, our rate is $\mathcal{O}(n^{2/3}K^{-2/3})$ in the total of iterations, while the rate of SGD is $\mathcal{O}(K^{-1/2})$. Our bound is only better than SGD if $n < \mathcal{O}(K^{1/4})$ under any shuffling strategy. If a randomized reshuffling strategy is used, our rate is improved to $\mathcal{O}(n^{1/3}K^{-2/3})$. This rate is better than SGD if $n < \mathcal{O}(K^{1/2})$. Note that our results and the ones of standard SGD are using different assumptions which may lead to the dependency on some constants for both bounds. In theory, shuffling gradient methods seem not better than SGD when n is large. However, these methods are often implemented in practice, especially in machine learning, as we have mentioned earlier. Our detailed discussion on the comparison with SGD is given in Remark 2. Further comparison between randomized reshuffling methods and SGD, GD, and other deterministic shuffling schemes for strongly convex problems can be found, e.g., in (HaoChen and Sra, 2019; Safran and Shamir, 2020).

Related work. Let us briefly review the most related works to our methods in this paper. The random shuffling-type method has been empirically studied in early works such as (Bottou, 2009) and also discussed in (Bottou, 2012). Its cyclic variant, known as an incremental gradient method was proposed even much earlier, see (Nedic and Bertsekas, 2001), where the convergence analysis was given in (Nedić and Bertsekas, 2001) for a subgradient variant, and in (Gürbüzbalaban et al., 2015) for gradient variants. These results are only for convex problems. Other incremental gradient variants can be found, e.g., in (Defazio et al., 2014a,c) known as SAGA-based methods. Our method is more general since it covers different shuffling variants in both deterministic and randomized worlds. When a randomized reshuffling strategy is used, we can improve our results to match the best known rates in both convex and nonconvex settings.

In (Gürbüzbalaban et al., 2019), the authors showed that if T is large, the randomized shuffling gradient method asymptotically converges with $\mathcal{O}(T^{-2})$ rate under a proper stepsize. However, this rate was only shown for strongly convex problems with bounded gradient/sequence, smoothness, and Lipschitz Hessian. These assumptions all together are very restrictive to hold in practice. Under the same conditions, (HaoChen and Sra, 2019) improved the convergence rate to $\mathcal{O}((nT)^{-2} + T^{-3})$ non-asymptotically, but in the regime of $T/\log(T) \geq \mathcal{O}(n)$. Another related work is (Nagaraj et al., 2019), which achieves $\tilde{\mathcal{O}}(1/(nT^2))$ convergence rates without Lipschitz Hessian when T is above some order of the condition number. Such a paper still requires F to have uniformly bounded gradient on its domain. Recently, an $\Omega(T^{-2} + n^2T^{-3})$ lower bound is proved in (Safran and Shamir, 2020) under the same assumptions as (HaoChen and Sra, 2019). Another $\Omega(nT^{-2})$ shaper lower bound is recently established in (Rajput et al., 2020) when $n > \mathcal{O}(T)$.

In (Ying et al., 2020), the authors replaced the i.i.d. sampling scheme by a randomized reshuffling strategy and established that variance reduced methods such as SAGA and SVRG still have a linear convergence rate for strongly convex problems but using an unusual

energy function. Unfortunately, it is unclear how to transform such a criterion to standard convergence criteria such as loss residuals or solution distances. The stochastic gradient method and its variance-reduced variants (e.g., SAG (Roux et al., 2012), SAGA (Defazio et al., 2014a), SVRG (Johnson and Zhang, 2013), and SARAH (Nguyen et al., 2017)) for solving (P) under strong convexity and smoothness structures usually have theoretically linear convergence rates compared to a sublinear rate in shuffling gradient variants. These variance-reduced techniques have also been widely exploited in nonconvex settings, but often require full gradient evaluations due to the use of double loops (Pham et al., 2020; Reddi et al., 2016; Tran-Dinh et al., 2020). However, these algorithms are different from the standard shuffling-type gradient method we study in this paper.

In (Shamir, 2016), a convergence rate to a neighborhood of the optimal value of an SGD variant using a without-replacement sampling strategy is studied for general convex problems. Clearly, this type of convergence is different from ours, and requires n to be large to get a suitable bound. If (P) is generalized linear and strongly convex, then a faster $\mathcal{O}(\log(K)/K)$ rate is achieved. Another recent work is (Meng et al., 2019) which considers different distributed SGD variants with shuffling for strongly convex, general convex, and nonconvex problems. The authors can only show convergence to a neighborhood of an optimal solution or a stationary point as in (Shamir, 2016). In addition, their rates are much slower than existing results for the strongly convex case, and also slower than ours, while requiring stronger assumptions. After the first draft of this paper was online, (Mishchenko et al., 2020) have made a step further by improving the convergence rates of the randomized reshuffling variant for the strongly convex case, but require convexity of each component function $f(\cdot; i)$. Moreover, they only focus on constant stepsize, while our results cover both constant and diminishing stepsizes.

Paper outline. The rest of this paper is organized as follows. Section 2 describes our general shuffling gradient algorithm to solve (P). We state our main assumptions and provides necessary mathematical tools in Section 2.3. Sections 3 and 4 analyze convergence for the nonconvex and convex cases, respectively. Several numerical experiments are presented in Section 5. For the sake of presentation, all details and proofs are deferred to Appendix.

2. The Shuffling-Type Gradient Algorithm and Technical Lemmas

Let us first describe our generic shuffling-type gradient algorithm for (P). Next, we state two standard assumptions imposed on (P), which will be used in this paper. Finally, we prove two technical lemmas that serve as key steps for our convergence analysis.

2.1 The generic shuffling-type gradient algorithm

Shuffling-type gradient methods are widely used in practice due to their efficiency (Bottou, 2009). Moreover, these methods have been investigated in many recent papers, including (Gürbüzbalaban et al., 2019; HaoChen and Sra, 2019; Nagaraj et al., 2019). In this paper, we analyze convergence rates for a wide class of shuffling-type schemes to solve (P) in both convex and nonconvex settings as described in Algorithm 1.

Note that $\pi^{(t)}(i)$ is the i -th element of $\pi^{(t)}$ for $i \in [n]$. Each outer iteration t of Algorithm 1 can be counted for one epoch. The inner loop updates the iterate sequence $\{w_i^{(t)}\}_{i=1}^n$

Algorithm 1 (Generic Shuffling-Type Gradient Algorithm for Solving (P))

- 1: **Initialization:** Choose an initial point $\tilde{w}_0 \in \text{dom}(F)$.
 - 2: **for** $t = 1, 2, \dots, T$ **do**
 - 3: Set $w_0^{(t)} := \tilde{w}_{t-1}$;
 - 4: Generate any permutation $\pi^{(t)}$ of $[n]$ (either deterministic or random);
 - 5: **for** $i = 1, \dots, n$ **do**
 - 6: Update $w_i^{(t)} := w_{i-1}^{(t)} - \eta_i^{(t)} \nabla f(w_{i-1}^{(t)}; \pi^{(t)}(i))$;
 - 7: **end for**
 - 8: Set $\tilde{w}_t := w_n^{(t)}$;
 - 9: **end for**
-

using only one component per iteration as in SGD by shuffling the objective components. Our analysis will be done in epoch-wise (i.e. the convergence guarantee is on $F(\tilde{w}_t) - F_*$ evaluated at the outer iterate \tilde{w}_t instead of $w_i^{(t)}$). Depending on the choice of $\pi^{(t)}$ we obtain different variants, especially the following methods:

- If $\pi^{(t)} = \{1, 2, \dots, n\}$ or some fixed permutation of $\{1, 2, \dots, n\}$ for all epochs t , then Algorithm 1 is equivalent to a cyclic gradient method. This method can also be viewed as the incremental gradient scheme studied in (Nedic and Bertsekas, 2001), and recently in (Li et al., 2020).
- If $\pi^{(t)}$ is randomly generated one time and repeatedly used at each iteration t , then Algorithm 1 becomes a single shuffling variant (Safran and Shamir, 2020).
- If $\pi^{(t)}$ is randomly generated at each epoch t , then Algorithm 1 reduces to a randomized reshuffling scheme, broadly used in practice, see, e.g., (Hiroyuki, 2018).

The randomized reshuffling schemes have been studied, e.g., in (Gürbüzbalaban et al., 2019; HaoChen and Sra, 2019; Nagaraj et al., 2019), but their convergence analysis has mainly been investigated for the strongly convex case and often under a strong set of assumptions.

2.2 Model assumptions

Our analysis throughout the paper relies on the following standard assumptions of (P).

Assumption 1 *Assume that problem (P) satisfies the following conditions:*

- (i) $\text{dom}(F) := \{x \in \mathbb{R}^d : F(x) < +\infty\} \neq \emptyset$ and $F_* := \inf_{w \in \mathbb{R}^d} F(w) > -\infty$.
- (ii) $f(\cdot; i)$ is L -smooth for all $i \in [n]$, i.e. there exists a constant $L \in (0, +\infty)$ such that:

$$\|\nabla f(w; i) - \nabla f(\hat{w}; i)\| \leq L\|w - \hat{w}\|, \quad \forall w, \hat{w} \in \text{dom}(F). \quad (2)$$

Note that Assumption 1(i) is required in any algorithm to guarantee the well-definedness of (P). Assumption 1(ii) is standard in gradient-type methods. Hence, we refer to Assumption 1 as the standard assumption required throughout the paper. Assumption 1(ii) implies that the objective function F is also L -smooth. Moreover, as proven in (Nesterov, 2004), we have

$$F(w) \leq F(\hat{w}) + \langle \nabla F(\hat{w}), w - \hat{w} \rangle + \frac{L}{2} \|w - \hat{w}\|^2, \quad \forall w, \hat{w} \in \text{dom}(F). \quad (3)$$

Apart from Assumption 1, we also require the following one in most of our results.

Assumption 2 *There exists two constants $\Theta \in [0, +\infty)$ and $\sigma \in (0, +\infty)$ such that*

$$\frac{1}{n} \sum_{i=1}^n \|\nabla f(w; i) - \nabla F(w)\|^2 \leq \Theta \|\nabla F(w)\|^2 + \sigma^2, \quad \forall w \in \text{dom}(F). \quad (4)$$

Note that if $\Theta = 0$, then Assumption 2 reduces to the standard bounded variance assumption $\mathbb{E}_i[\|f(w; i) - \nabla F(w)\|^2] \leq \sigma^2$, which is often used in nonconvex problems, see, e.g., (Ghadimi and Lan, 2013; Pham et al., 2020).

2.3 Technical lemmas for convergence analysis

The following two lemmas provide key estimates for our convergence analysis in this paper. We first state them here and provide their proof in Appendix A.2.

Lemma 1 *Let $\{Y_t\}_{t \geq 1}$ be a nonnegative sequence in \mathbb{R} and q be a positive integer number. Let $\rho > 0$ and $D > 0$ be two given constants and $0 < \eta_t \leq \frac{1}{\rho}$ be given for all $t \geq 1$. Assume that, for all $t \geq 1$, we have*

$$Y_{t+1} \leq (1 - \rho \cdot \eta_t) Y_t + D \cdot \eta_t^{q+1}. \quad (5)$$

If we choose $\eta_t := \frac{q}{\rho(t+\beta)}$ for all $t \geq 1$, where $\beta \geq q - 1$, then we have

$$Y_{t+1} \leq \frac{\beta \cdots (\beta - q + 1)}{(t + \beta - q + 1) \cdots (t + \beta)} Y_1 + \frac{q^{q+1} D \log(t + \beta)}{\rho^{q+1} (t + \beta - q + 1) \cdots (t + \beta)}. \quad (6)$$

If we choose $\eta_t := \eta \in (0, \rho^{-1})$ for all $t \geq 1$, then we have

$$Y_{t+1} \leq (1 - \rho\eta)^t Y_1 + \frac{D\eta^q [1 - (1 - \rho\eta)^t]}{\rho} \leq Y_1 \exp(-\rho\eta t) + \frac{D\eta^q}{\rho}. \quad (7)$$

Lemma 2 *Let $\{Y_t\}_{t \geq 1}$ and $\{Z_t\}_{t \geq 1}$ be two nonnegative sequences in \mathbb{R} and m and q be two positive numbers such that $q > m$. For positive constants $\rho, \alpha, \beta, \gamma$, and D , assume that*

$$Y_{t+1} \leq Y_t - \rho\eta_t^m \cdot Z_t + \eta_t^q \cdot D, \quad \text{where } \eta_t := \frac{\gamma}{(t + \beta)^\alpha} \text{ and } \alpha m \leq \frac{1}{2}. \quad (8)$$

Suppose that $Y_t \leq C + H \log(t + \theta)$ for some $C > 0$, $H \geq 0$, $\theta > 0$, and $1 + \theta - \beta > (1 - \alpha m)e^{\frac{\alpha m}{1 - \alpha m}}$ for all $t \geq 1$, (where e is the natural number). Then, we have

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T Z_t &\leq \frac{1}{T} \left[\frac{(1 + \beta)^{\alpha m} Y_1}{\rho \gamma^m} + \frac{C(T - 1 + \beta)^{\alpha m}}{2\rho \alpha m \gamma^m} + \frac{H(T - 1 + \beta)^{\alpha m} \log(T + \theta)}{2\rho \alpha m \gamma^m} \right] \\ &\quad + \frac{D\gamma^{q-m}}{\rho} \cdot \frac{A(T)}{T}, \end{aligned} \quad (9)$$

where

$$A(T) := \begin{cases} \log(T + \beta) - \log(\beta) & \text{if } \alpha(q - m) = 1, \\ \frac{(T + \beta)^{1 - \alpha(q - m)}}{1 - \alpha(q - m)} & \text{otherwise.} \end{cases}$$

Lemmas 1 and 2 are independent of interest, and play an important role not only for deriving the results for shuffling-type gradient methods but also for applying our new convergence analysis to the standard SGD algorithm (see Appendix D for more details).

Remark 1 (Types of guarantee) *Since we can choose the permutation $\pi^{(t)}$ of Algorithm 1 either deterministically or randomly, our bounds in the sequel will hold either deterministically or in expectation, respectively. Without loss of generality, we write these results in the context of expectation taken overall the randomness generated by the algorithm.*

3. Convergence Analysis for Convex Case

In this section, we mainly consider two cases. In the first case, we only assume that the sum function F is strongly convex while some components $f(\cdot; i)$ for $i \in [n]$ are not necessarily convex. In the second case, we assume that F is strongly convex and each $f(\cdot; i)$ for $i \in [n]$ is also convex. Let us state the strong convexity assumption of F as follows.

Assumption 3 (μ -strong convexity) *The objective function F of (P) is μ -strongly convex on $\text{dom}(F)$, i.e. there exists a constant $\mu \in (0, +\infty)$ such that*

$$F(w) \geq F(\hat{w}) + \langle \nabla F(\hat{w}), w - \hat{w} \rangle + \frac{\mu}{2} \|w - \hat{w}\|^2, \quad \forall w, \hat{w} \in \text{dom}(F). \quad (10)$$

It is well-known from the literature (Nesterov, 2004; Bottou et al., 2018) that Assumption 3 implies the existence and uniqueness of the optimal solution w_* of (P), and

$$\frac{\mu}{2} \|w - w_*\|^2 \leq F(w) - F(w_*) \leq \frac{1}{2\mu} \|\nabla F(w)\|^2, \quad \forall w \in \text{dom}(F). \quad (11)$$

As we have mentioned above, Assumption 3 only requires the sum function F to be strongly convex, but some components $f(\cdot; i)$ can even be nonconvex.

Under Assumption 3, since problem (P) has a unique optimal solution w_* , we introduce the following variance of F at w_* :

$$\sigma_*^2 := \frac{1}{n} \sum_{i=1}^n \|\nabla f(w_*; i)\|^2 \in [0, +\infty). \quad (12)$$

Now, we state the convergence of Algorithm 1 with a constant learning rate in Theorem 1, whose proof is given in Appendix B.1.

Theorem 1 *Assume that Assumptions 1 and 3 hold (but some $f(\cdot; i)$ for $i \in [n]$ are not necessarily convex), σ_*^2 is defined by (12), and $\kappa := \frac{L}{\mu}$ is the condition number of F . Let $\{w_i^{(t)}\}_{t=1}^T$ be generated by Algorithm 1 after T epochs with a constant stepsize $\eta_i^{(t)} := \frac{6 \log(T)}{\mu n T}$ and any shuffling strategy $\pi^{(t)}$. Then, for $T \geq 1$ such that $T \geq 12\kappa^2 \log(T)$, we have*

$$\begin{aligned} \mathbb{E}[F(\tilde{w}_T) - F(w_*)] &\leq \frac{1}{T^2} \left[(F(\tilde{w}_0) - F(w_*)) + \frac{54(\mu^2 + L^2)\sigma_*^2 \log(T)^2}{\mu^3} \right] \\ &= \mathcal{O} \left(\frac{\log(T)^2}{T^2} \right). \end{aligned} \quad (13)$$

Assume additionally that $\pi^{(t)}$ is sampled uniformly at random without replacement from $[n]$ and Assumption 2 holds. Then, by choosing a learning rate $\eta_i^{(t)} := \frac{4 \log(\sqrt{n}T)}{\mu n T}$ for $T \geq 1$ such that $T \geq \frac{8L\sqrt{\Theta/n+1}}{\mu^2} \log(\sqrt{n}T)$, we have

$$\mathbb{E}[F(\tilde{w}_t) - F(w_*)] \leq \frac{1}{nT^2} \left[[F(\tilde{w}_0) - F(w_*)] + \frac{2L^2\sigma^2 \log(\sqrt{n}T)^2}{\mu^3} \right]. \quad (14)$$

Consequently, the convergence rate of $\{\mathbb{E}[F(\tilde{w}_T) - F(w_*)]\}$ in this case is $\mathcal{O}\left(\frac{\log(T)^2}{nT^2}\right)$.

Alternatively, assume additionally that $\pi^{(t)}$ is sampled uniformly at random without replacement from $[n]$ and each $f(\cdot; i)$ is convex for $i \in [n]$. Then, by choosing a constant learning rate $\eta_i^{(t)} := \frac{2 \log(\sqrt{n}T)}{\mu n T}$ for $T \geq 1$ such that $\log(T\sqrt{n}) \leq \frac{T}{2} \min\{1, \frac{\sqrt{5}-1}{\kappa}\}$, we have

$$\mathbb{E}[F(\tilde{w}_t) - F(w_*)] \leq \frac{L}{2} \mathbb{E}[\|\tilde{w}_t - w_*\|^2] \leq \frac{L}{2nT^2} \left[\mathbb{E}[\|\tilde{w}_0 - w_*\|^2] + \frac{8L\sigma_*^2 \log(\sqrt{n}T)^2}{3\mu^3} \right]. \quad (15)$$

Consequently, the convergence rate of both $\{\mathbb{E}[F(\tilde{w}_T) - F(w_*)]\}$ and $\{\mathbb{E}[\|\tilde{w}_T - w_*\|^2]\}$ in this case is $\mathcal{O}\left(\frac{\log(T)^2}{nT^2}\right)$.

The condition $\log(T\sqrt{n}) \leq \frac{T}{2} \min\{1, \frac{\sqrt{5}-1}{\kappa}\}$ on T of (15) shows that $\frac{T}{\log(T\sqrt{n})} \geq \mathcal{O}(\kappa)$, where $\kappa := \frac{L}{\mu}$ is the condition number of F . This condition has been shown in previous works, and aligns with recent results in (Ahn et al., 2020; Mishchenko et al., 2020). However, unlike (Mishchenko et al., 2020), we have new results stated in (13) and (14).

Next, we prove the following result for the strongly convex case using diminishing learning rates. The detailed proof of this theorem is given in Appendix B.1.

Theorem 2 *Assume that Assumptions 1 and 3 hold (but some $f(\cdot; i)$ for $i \in [n]$ are not necessarily convex), σ_*^2 is defined by (12), and $\kappa := \frac{L}{\mu}$ is the condition number of F . Let $\{w_i^{(t)}\}$ be generated by Algorithm 1 with $\eta_i^{(t)} := \frac{\eta t}{n}$ and any shuffling strategy $\pi^{(t)}$ for solving (P). Let η_t be updated by $\eta_t := \frac{6}{\mu(t+\beta)}$ for all $t \geq 1$, where $\beta \geq 12\kappa^2 - 1$. Then, we have*

$$\begin{aligned} \mathbb{E}[F(\tilde{w}_t) - F(w_*)] &\leq \frac{\beta(\beta-1)}{(t+\beta)(t+\beta-1)} \left[(F(\tilde{w}_0) - F(w_*)) + \frac{216(L^2 + \mu^2)\sigma_*^2 \log(t+\beta)}{\mu^3\beta(\beta-1)} \right] \\ &= \mathcal{O}\left(\frac{\log(t)}{t^2}\right). \end{aligned} \quad (16)$$

If, additionally, $\pi^{(t)}$ is sampled uniformly at random without replacement from $[n]$, $f(\cdot; i)$ (for all $i \in [n]$) are convex, and $L \leq \frac{\sqrt{5}-1}{2}$, then by choosing $\eta_t := \frac{2}{\mu(t+1+1/n)}$ for $t \geq 1$, we have

$$\begin{aligned} \mathbb{E}[F(\tilde{w}_t) - F(w_*)] &\leq \frac{L}{2} \mathbb{E}[\|\tilde{w}_t - w_*\|^2] \\ &\leq \frac{2L}{n(t+1/n)(t+1/n+1)} \left[\mathbb{E}[\|\tilde{w}_0 - w_*\|^2] + \frac{L\sigma_*^2 \log(t+\beta)}{3\mu^3} \right] \\ &= \mathcal{O}\left(\frac{\log(t)}{nt^2}\right). \end{aligned} \quad (17)$$

The condition $L \leq \frac{\sqrt{5}-1}{2}$ in Theorem 2 holds without loss of generality. Indeed, if it does not hold, then by rescaling $f(\cdot; i)$ by $\frac{\sqrt{5}-1}{2L} \cdot f(\cdot; i)$, we obtain this condition. Clearly, with a randomized reshuffling strategy, the convergence rate of Algorithm 1 is better than a general shuffling one. This rate nearly matches the lower bound in previous works, e.g., in (Nagaraj et al., 2019; Rajput et al., 2020). While our proof of (16) is new, the proof of (17) is inspired by (Mishchenko et al., 2020), but it is rather different (see Lemma 7 in the appendix). Moreover, our bound (17) is established for a diminishing stepsize η_t .

Since the total number of iterations is $K := nT$, if we write the convergence rates in terms of K , then for any shuffling strategy, we have $\mathbb{E}[F(\tilde{w}_T) - F(w_*)] \leq \mathcal{O}(n^2 K^{-2})$. However, for a randomized reshuffling one, we have $\mathbb{E}[F(\tilde{w}_T) - F(w_*)] \leq \mathcal{O}(nK^{-2})$, matching the result in (HaoChen and Sra, 2019; Mishchenko et al., 2020). As mentioned earlier, our assumptions for (17) are as in (Mishchenko et al., 2020) and weaker than those in (HaoChen and Sra, 2019; Nagaraj et al., 2019). Here, we use diminishing learning rates in Theorem 2 instead of constant ones as in (HaoChen and Sra, 2019; Mishchenko et al., 2020; Nagaraj et al., 2019). Note that Algorithm 1 covers much broader class of algorithms compared to existing methods in the literature.

Remark 1 (Non-strongly convex case) *Similar to (Mishchenko et al., 2020), we can use (39) of Lemma 7 to prove the following convergence rate for Algorithm 1 under only convexity of $f(\cdot; i)$ for all $i \in [n]$. We state this result as follows without proof.*

$$\mathbb{E}[F(\hat{w}_T) - F(w_*)] \leq \frac{1}{\Sigma_T} \sum_{t=1}^T \eta_t \mathbb{E}[F(\tilde{w}_{t-1}) - F(w_*)] \leq \frac{1}{2\Sigma_T} \|\tilde{w}_0 - w_*\|^2 + \frac{L\sigma_*^2}{3n\Sigma_T} \cdot \sum_{t=1}^T \eta_t^3,$$

where $\Sigma_T := \sum_{t=1}^T \eta_t$ and $\hat{w}_T := \frac{1}{\Sigma_T} \sum_{t=1}^T \eta_t \tilde{w}_{t-1}$ is a weighted averaging sequence.

- If we choose $\eta := \frac{\gamma n^{1/3}}{T^{1/3}} \leq \frac{1}{2L}$ as a constant learning rate, then we have

$$\mathbb{E}[F(\hat{w}_T) - F(w_*)] \leq \frac{1}{n^{1/3} T^{2/3}} \cdot \left[\frac{1}{2\gamma} \|\tilde{w}_0 - w_*\|^2 + \frac{\gamma^2 L \sigma_*^2}{3} \right].$$

- If we choose $\eta_t := \frac{\gamma n^{1/3}}{(t+\beta)^{1/3}}$ for $1 \leq t \leq T$ as a diminishing learning rate for some $\gamma > 0$ and $\beta \geq 1$ such that $\beta \geq 8L^3 \gamma^3 n - 1$, then we have

$$\mathbb{E}[F(\hat{w}_T) - F(w_*)] \leq \frac{1}{n^{1/3} T^{2/3}} \cdot \left[\frac{1}{2\gamma} \|\tilde{w}_0 - w_*\|^2 + \frac{\gamma^2 L \sigma_*^2 \log(T + \beta)}{3} \right].$$

The constant learning rate case stated here is discussed in (Mishchenko et al., 2020). However, we add a new statement on diminishing learning rate that may be practically favorable.

4. Convergence Analysis for Nonconvex Case

We now provide convergence analysis for Algorithm 1 to solve nonconvex smooth instances of (P). We consider two cases in Subsection 4.1 and 4.2, respectively.

4.1 The general case

We state our first result on the nonconvex case, whose proof is given in Appendix C.1.

Theorem 3 *Suppose that Assumptions 1 and 2 hold for (P). Let $\{\tilde{w}_t\}_{t=1}^T$ be generated by Algorithm 1 with any shuffling strategy for $\pi^{(t)}$ and any learning rate $\eta_i^{(t)} = \frac{\eta t}{n} = \frac{\eta}{n}$ such that $0 < \eta \leq \frac{1}{L\sqrt{2(3\Theta+2)}}$. Then, we have*

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] \leq \frac{4}{T\eta} [F(\tilde{w}_0) - F_*] + 6L^2\sigma^2\eta^2. \quad (18)$$

If, additionally, $\pi^{(t)}$ is sampled uniformly at random without replacement from $[n]$ and η_t is chosen such that $0 < \eta \leq \frac{1}{L\sqrt{2(\Theta/n+1)}}$, then we have

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] \leq \frac{4}{T\eta} [F(\tilde{w}_0) - F_*] + \frac{4L^2\sigma^2\eta^2}{n}. \quad (19)$$

If L , σ , and Θ are known, then we can choose the following learning rate to get a concrete bound as stated in Corollary 1, whose proof can be found in Appendix C.1.

Corollary 1 *Let $\{\tilde{w}_t\}_{t=1}^T$ be generated by Algorithm 1 for solving (P). For a given ϵ such that $0 < \epsilon \leq 2\sigma^2$, under the same conditions as of (18) in Theorem 3, if we choose a constant learning rate $\eta := \frac{\sqrt{\epsilon}}{2L\sigma\sqrt{3\Theta+2}}$, then to guarantee $\frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] \leq \epsilon$ for (P), it requires at most $T := \left\lceil \frac{16L\sigma(3\Theta+2)^{3/2}[F(\tilde{w}_0)-F_*]}{(6\Theta+1)} \cdot \frac{1}{\epsilon^{3/2}} \right\rceil$ outer iterations. As a result, the total number of gradient evaluations is at most $\mathcal{T}_{\nabla f} := \left\lceil \frac{16L\sigma(3\Theta+2)^{3/2}[F(\tilde{w}_0)-F_*]}{(6\Theta+1)} \cdot \frac{n}{\epsilon^{3/2}} \right\rceil$.*

If, in addition, $\pi^{(t)}$ is sampled uniformly at random without replacement from $[n]$, then by choosing $\eta := \frac{\sqrt{n\epsilon}}{2L\sigma\sqrt{2(\Theta/n+1)}}$ for $0 < \epsilon \leq \frac{4\sigma^2}{n}$, to guarantee $\frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] \leq \epsilon$, it requires at most $T := \left\lceil \frac{16(\Theta/n+1)^{3/2}[F(\tilde{w}_0)-F_]}{(2\Theta/n+1)} \cdot \frac{L\sigma}{\sqrt{n\epsilon^{3/2}}} \right\rceil$ outer iterations. Consequently, the total number of gradient evaluations is at most $\mathcal{T}_{\nabla f} := \left\lceil \frac{16(\Theta/n+1)^{3/2}[F(\tilde{w}_0)-F_*]}{(2\Theta/n+1)} \cdot \frac{L\sigma\sqrt{n}}{\epsilon^{3/2}} \right\rceil$.*

Remark 2 To obtain the same bound, the total complexity of the standard SGD is $\mathcal{O}(L_F\sigma_S^2\epsilon^{-2})$ for solving (1) under the bounded variance $\mathbb{E}_i[\|\nabla f(w; i) - \nabla F(w)\|^2] \leq \sigma_S^2$ for some $\sigma_S > 0$ and the L_F -smoothness of F . Note that the standard SGD only requires F to be L_F -smooth while we impose the smoothness on individual realizations. Therefore, L_F and L may be different (Ghadimi and Lan, 2013). For a rough comparison, if $n < \mathcal{O}\left(\frac{L_F\sigma_S^2}{L\sigma} \cdot \frac{1}{\epsilon^{1/2}}\right)$, then Algorithm 1 with any shuffling strategy seems to have advantages over the standard SGD method in the nonconvex setting. In addition, if a randomized reshuffling strategy is used, then for $n \leq \frac{4\sigma^2}{\epsilon}$, the complexity of Algorithm 1 is $\mathcal{O}\left(\frac{L\sigma\sqrt{n}}{\epsilon^{3/2}}\right)$, which is better than SGD by a factor $\frac{\sigma}{\epsilon^{1/3}}$. From this point of view, it seems that Algorithm 1 with a general shuffling strategy is theoretically less efficient than SGD when a low accuracy solution is desirable (i.e. ϵ is not too small) or when $n \gg 1$. However, we believe that Algorithm 1 allows more

flexible strategy to choose $f(\cdot; i)$ rather than that i.i.d. sampling. For instance, choosing a randomized reshuffling strategy significantly improves the complexity of Algorithm 1. We also note that our convergence guarantee is completely different from (Meng et al., 2019) as mentioned earlier. Nevertheless, Assumptions 1 and 2 are very standard and hold for various applications in machine learning.

Let us propose another choice of the learning rate η_t in the following result, whose proof can also be founded in Appendix C.1.

Corollary 2 *Let $\{\tilde{w}_t\}_{t=1}^T$ be generated by Algorithm 1. Under the same conditions as of (18) in Theorem 3, if we choose a constant learning rate $\eta := \frac{\gamma}{T^{1/3}}$ for some $\gamma > 0$ and $T \geq 1$ such that $T^{1/3} \geq \gamma L \sqrt{2(3\Theta + 2)}$, then we have*

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] \leq \frac{1}{T^{2/3}} \left[\frac{4(F(\tilde{w}_0) - F_*)}{\gamma} + 6L^2\sigma^2\gamma^2 \right] = \mathcal{O}\left(\frac{1}{T^{2/3}}\right). \quad (20)$$

If, in addition, $\pi^{(t)}$ is sampled uniformly at random without replacement from $[n]$, then by choosing $\eta := \frac{\gamma n^{1/3}}{T^{1/3}}$ for some $\gamma > 0$ such that $T^{1/3} \geq \gamma L n^{1/3} \sqrt{2(\Theta/n + 1)}$, we have

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] \leq \frac{4}{n^{1/3}T^{2/3}} \left[\frac{(F(\tilde{w}_0) - F_*)}{\gamma} + L^2\sigma^2\gamma^2 \right] = \mathcal{O}\left(\frac{1}{n^{1/3}T^{2/3}}\right). \quad (21)$$

Since the total number of iterations is $K := nT$, if we express (20) in terms of K , then we have $\frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] \leq \frac{n^{2/3}\Delta_0}{K^{2/3}}$, where $\Delta_0 := F(\tilde{w}_0) - F_*$. Alternatively, we can express (21) in terms of K as $\frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] \leq \frac{n^{1/3}\Delta_0}{K^{2/3}}$.

The following theorem characterizes an asymptotic convergence for general diminishing stepsizes, whose proof can be found in Appendix C.2.

Theorem 4 *Suppose that Assumptions 1 and 2 hold for (P). Let $\{\tilde{w}_t\}_{t \geq 1}$ be generated by Algorithm 1 with diminishing learning rate $\eta_i^{(t)} = \frac{\eta_t}{n}$ such that $\sum_{t=1}^{\infty} \eta_t = \infty$ and $\sum_{t=1}^{\infty} \eta_t^3 < \infty$. Then, w.p.1. (i.e. almost surely), we have $\liminf_{t \rightarrow \infty} \|\nabla F(\tilde{w}_{t-1})\|^2 = 0$.*

Now, if we vary η_t , then Theorem 5 shows how η_t affects our rates (see Appendix C.3).

Theorem 5 *Suppose that Assumptions 1 and 2 hold for (P). Let $\{\tilde{w}_t\}_{t=1}^T$ be generated by Algorithm 1 with $\eta_i^{(t)} = \frac{\eta_t}{n}$, where $\eta_t := \frac{\gamma}{(t+\beta)^\alpha}$, for some $\gamma > 0$, $\beta > 0$, and $\frac{1}{3} < \alpha < 1$. If a generic shuffling strategy is used, then we let $D := \frac{3}{2}L^2\sigma^2$ and assume that $\gamma L \sqrt{2(3\Theta + 2)} \leq (\beta+1)^\alpha$. Otherwise, if a uniformly randomized reshuffling strategy is used, then we set $D := \frac{L^2\sigma^2}{n}$ and assume that $\gamma L \sqrt{2(\Theta/n + 1)} \leq (\beta+1)^\alpha$. Let $C := [F(\tilde{w}_0) - F_*] + \frac{D\gamma^3}{(3\alpha-1)\beta^{3\alpha-1}} > 0$ be a given constant. Then, the following statements hold:*

- If $\alpha = \frac{1}{2}$, then the following bound holds:

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] &\leq \frac{4(1+\beta)^{1/2} [F(\tilde{w}_0) - F_*]}{\gamma} \cdot \frac{1}{T} + \frac{4C}{\gamma} \cdot \frac{(T-1+\beta)^{1/2}}{T} \\ &\quad + 4D\gamma^2 \cdot \frac{\log(T+\beta) - \log(\beta)}{T}. \end{aligned}$$

- If $\alpha \neq \frac{1}{2}$, then we have

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] &\leq \frac{4(1+\beta)^\alpha [F(\tilde{w}_0) - F_*]}{\gamma} \cdot \frac{1}{T} + \frac{2C}{\alpha\gamma} \cdot \frac{(T-1+\beta)^\alpha}{T} \\ &\quad + \frac{4D\gamma^2}{(1-2\alpha)} \cdot \frac{(T+\beta)^{1-2\alpha}}{T}. \end{aligned}$$

If a uniformly randomized reshuffling strategy is used, then by replacing γ by $n^{1/3}\gamma$, we have

- If $\alpha = \frac{1}{2}$, then the convergence rate of Algorithm 1 is $\mathcal{O}\left(\frac{1+T^{1/2}+\log(T)}{n^{1/3}T}\right)$.
- If $\alpha \neq \frac{1}{2}$, then the convergence rate of Algorithm 1 is $\mathcal{O}\left(\frac{1+T^{1-2\alpha}+\log(T)}{n^{1/3}T}\right)$.

Remark 3 In Theorem 5, if we choose $\alpha := \frac{1}{3} + \delta$ for some $0 < \delta < \frac{1}{6}$, then we have

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] &\leq \frac{4(1+\beta)^{\frac{1}{3}+\delta} [F(\tilde{w}_0) - F_*]}{\gamma} \cdot \frac{1}{T} + \frac{2C}{\gamma(\frac{1}{3}+\delta)} \cdot \frac{(T-1+\beta)^{\frac{1}{3}+\delta}}{T} \\ &\quad + \frac{12D\gamma^2}{1-6\delta} \cdot \frac{(T+\beta)^{\frac{1}{3}-2\delta}}{T}, \end{aligned}$$

where $C := [F(\tilde{w}_0) - F_*] + \frac{D\gamma^3}{3\delta\beta^{3\delta}}$. Hence, the convergence rate of Algorithm 1 is $\mathcal{O}(T^{-(\frac{2}{3}-\delta)})$ in general. If a randomized reshuffling strategy is used, then this rate is $\mathcal{O}(n^{-1/3}T^{-(\frac{2}{3}-\delta)})$.

For the extreme case $\alpha := \frac{1}{3}$, we have the following result (see Appendix C.3).

Theorem 6 Suppose that Assumptions 1 and 2 hold for (P). Let $\{\tilde{w}_t\}_{t=1}^T$ be generated by Algorithm 1 with $\eta_i^{(t)} := \frac{\eta_t}{n}$, where $\eta_t := \frac{\gamma}{(t+\beta)^{1/3}}$ for some $\gamma > 0$ and $\beta > 0$. If any shuffling strategy is used, then let $D := \frac{3}{2}L^2\sigma^2$ and assume that $\gamma L\sqrt{2(3\Theta+2)} \leq (\beta+1)^{1/3}$. Otherwise, if a uniformly randomized reshuffling strategy is used, then let $D := \frac{L^2\sigma^2}{n}$ and assume that $\gamma L\sqrt{2(\Theta/n+1)} \leq (\beta+1)^{1/3}$. Let $C := [F(\tilde{w}_0) - F_*] + \frac{D\gamma^3}{(1+\beta)} > 0$ be a given constant. Then the following bound holds:

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] &\leq \frac{4(1+\beta)^{1/3} [F(\tilde{w}_0) - F_*]}{\gamma} \cdot \frac{1}{T} + 12D\gamma^2 \cdot \frac{(T+\beta)^{1/3}}{T} \\ &\quad + 6D\gamma^2 \cdot \frac{(T-1+\beta)^{1/3} \log(T+1+\beta)}{T} + \frac{6C}{\gamma} \cdot \frac{(T-1+\beta)^{1/3}}{T}. \end{aligned}$$

Consequently, the convergence rate of Algorithm 1 is $\mathcal{O}\left(\frac{1}{T} + \frac{1}{T^{2/3}} + \frac{\log(T)}{T^{2/3}}\right)$. In addition, if a uniformly randomized reshuffling strategy is used, then by replacing γ by $n^{1/3}\gamma$ into the above estimate, the convergence rate of Algorithm 1 is $\mathcal{O}\left(\frac{1}{n^{1/3}T} + \frac{1}{n^{1/3}T^{2/3}} + \frac{\log(T)}{n^{1/3}T^{2/3}}\right)$, provided that $\gamma Ln^{1/3}\sqrt{2(\Theta/n+1)} \leq (\beta+1)^{1/3}$.

Remark 4 The choice of the learning rates in Theorems 5 and 6 is not necessarily dependent on the smoothness constant L and Θ of Assumption 2. Since η_t is diminishing and both L and Θ are finite, by choosing β large, the condition $\eta_t \leq \frac{1}{L\sqrt{2(3\Theta+2)}}$ or $\eta_t \leq \frac{1}{L\sqrt{2(\Theta/n+1)}}$ automatically holds. Therefore, the choices of γ , β , and α makes our results more flexible to adjust in particular practical implementation. We believe that a diminishing or scheduled diminishing learning rate is more favorable in practice than a constant one.

4.2 Convergence under gradient dominance

We can further improve convergence rates of Algorithm 1 in the nonconvex case by imposing the following gradient dominance condition.

Assumption 4 A function F is said to be τ -gradient dominant if there exists a constant $\tau \in (0, +\infty)$ such that

$$F(w) - F_* \leq \tau \|\nabla F(w)\|^2, \quad \forall w \in \text{dom}(F), \quad (22)$$

where $F_* := \inf_{w \in \mathbb{R}^d} F(w)$.

This assumption is well-known and widely used in the literature, see, e.g., (Karimi et al., 2016; Nesterov and Polyak, 2006; Polyak, 1964). It is also weaker than a strong convexity assumption. When F_* is achievable (i.e. $F_* = F(w_*)$ for some w_*), then we can observe that every stationary point w_* of the τ -gradient dominant function F is a global minimizer. However, such a function F is not necessarily convex.

The following theorem states the convergence rate of Algorithm 1 under gradient dominance, whose proof is deferred to Appendix C.4.

Theorem 7 Suppose that Assumptions 1, 2, and 4 hold for (P). Let $\{w_i^{(t)}\}$ be generated by Algorithm 1 for solving (P) using $\eta_i^{(t)} := \frac{\eta_t}{n}$ and any shuffling strategy. Let η_t be updated as $\eta_t := \frac{2}{t+\beta}$ for some $\beta \geq \max\{2L\sqrt{2(3\Theta+2)} - 1, 1\}$. Then, for all $t \geq 1$, we have

$$\mathbb{E}[F(\tilde{w}_t) - F_*] \leq \frac{1}{(t+\beta-1)(t+\beta)} \left[\beta(\beta-1)(F(\tilde{w}_0) - F_*) + 768 \cdot \tau^3 L^2 \sigma^2 \log(t+\beta) \right]. \quad (23)$$

Consequently, the convergence rate of $\{\mathbb{E}[F(\tilde{w}_t) - F_*]\}$ is $\mathcal{O}\left(\frac{\log(t)}{t^2}\right)$.

If, in addition, $\pi^{(t)}$ is uniformly sampled at random without replacement from $[n]$ and $L\sqrt{2(\Theta/n+1)} \leq 1$, then by choosing $\eta_t := \frac{2}{t+1+1/n}$, for all $t \geq 1$, we have

$$\mathbb{E}[F(\tilde{w}_t) - F_*] \leq \frac{2}{n(t+1/n)(t+1+1/n)} \left[(F(\tilde{w}_0) - F_*) + 265 \cdot \tau^3 L^2 \sigma^2 \log(t+\beta) \right]. \quad (24)$$

Consequently, the convergence rate of $\{\mathbb{E}[F(\tilde{w}_t) - F_*]\}$ is $\mathcal{O}\left(\frac{\log(t)}{nt^2}\right)$.

As we mentioned earlier, the condition $L\sqrt{2(\Theta/n+1)} \leq 1$ for (24) is not restrictive. One can always scale $f(\cdot; i)$ to guarantee this condition. The rate stated in (23) for the general shuffling strategy is $\tilde{\mathcal{O}}(1/t^2)$ for any $1 \leq t \leq T$ without fixing T a priori. If,

in addition, a randomized reshuffling strategy is used, then our rate in (24) is improved to $\tilde{\mathcal{O}}(1/(nt^2))$. Our rates in both cases are for diminishing learning rates. Note that (HaoChen and Sra, 2019) provide $\mathcal{O}(1/(nT^2))$ convergence rate but under a constant learning rate and stronger assumptions, including Lipschitz Hessian continuity. A recent work in (Ahn et al., 2020) also considers this gradient dominance case and achieves the $\mathcal{O}(1/(nT^2))$ rate but still requires a bounded gradient condition and using constant learning rate.

5. Numerical Experiments

In this section, we provide various numerical experiments to illustrate the theoretical convergence results of Algorithm 1 for solving nonconvex problem instances of (P). We only focus on the nonconvex setting since the convex case has been intensively studied in previous works, e.g., in (Ahn et al., 2020; Gürbüzbalaban et al., 2019; HaoChen and Sra, 2019; Mishchenko et al., 2020). We implement Algorithm 1 in Python and compare between different variants. Our code is available online at <https://github.com/lamnguyen-mltd/shuffling>. For each experiment, we conduct 10 runs and reported the average results.

5.1 Nonconvex logistic regression example

We consider the following well-studied binary classification problem with nonconvex F :

$$\min_{w \in \mathbb{R}^d} \left\{ F(w) := \frac{1}{n} \sum_{i=1}^n \left[\log(1 + \exp(-y_i x_i^\top w)) + \frac{\lambda}{2} \sum_{j=1}^d \frac{w_j^2}{1 + w_j^2} \right] \right\}, \quad (25)$$

where $w \in \mathbb{R}^d$ is the vector of model parameters and w_j is the element of w , $\{(x_i, y_i)\}_{i=1}^n$ is a set of training examples, and $\lambda > 0$ is a given regularization parameter.

We first conduct experiments to demonstrate the performance of Algorithm 1 on two classification datasets *w8a* ($n = 49,749$ samples) and *ijcnn1* ($n = 91,701$) from LIBSVM (Chang and Lin, 2011). Since we only aim at testing the nonconvexity of each f_i instead of statistical properties, we simply choose $\lambda := 0.01$, but other values of λ also work. The input data x_i ($i \in [n]$) have been scaled in the range of $[0, 1]$.

We apply Algorithm 1 with mini-batch size of 1 and $\eta_i^{(t)} := \frac{\eta_t}{n}$ to solve (25), where $\eta_t := \frac{\gamma}{(t+\beta)^\alpha}$ and $\pi^{(t)}$ is generated randomly to obtain an SGD variant with randomized reshuffling strategy. We experiment using different configurations: $\alpha = \{1/3, 1/2, 1\}$ and $\gamma/n = \{0.001, 0.005, 0.01\}$, respectively on the two datasets: *w8a* and *ijcnn1*.

Figures 1 and 2 show our comparison on the loss value $F(\tilde{w}_t)$ and the test accuracy using different configurations: $\alpha = \{1/3, 1/2, 1\}$ and $\gamma/n = \{0.001, 0.005, 0.01\}$, respectively on the *w8a* and *ijcnn1* datasets. The choices $\alpha = 1$ and $\gamma/n = 0.001$ do not give good training performances, hence we omit them in our plots.

Discussion. We observe from Figures 1 and 2 that the value $\alpha := 1/3$ used in the learning rate of Algorithm 1 usually gives the best performance. If we fix $\alpha := 1/3$ and use different ratios γ/n , then as showed in these plots, $\gamma/n = 0.01$ seems to work best. Note that we plotted the confidence intervals in every figure, however these intervals for train loss can barely be seen because the loss values for different random seeds do not deviate much from the mean value in this experiment. The configuration $\alpha = 1/3$ and $\gamma/n = 0.01$




Figure 1: The behavior of the train loss $F(\tilde{w}_t)$ and the test accuracy (starting from the 2nd epoch) of (25) produced by different values of α and γ/n in Algorithm 1 using the *w8a* dataset.

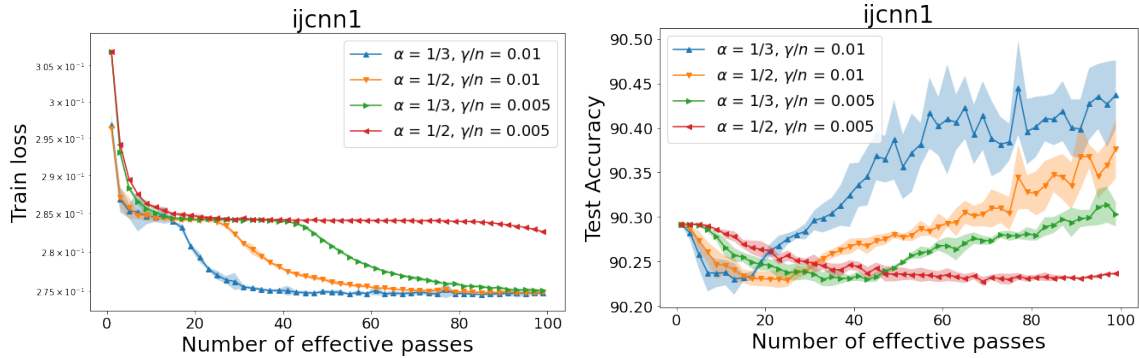


Figure 2: The behavior of the train loss $F(\tilde{w}_t)$ and the test accuracy (starting from the 2nd epoch) of (25) produced by different values of α and γ/n in Algorithm 1 using the *ijcnn1* dataset.

for the datasets *w8a* and *ijcnn1* has larger confidence intervals for the test accuracy and can be seen easily in Figures 1 and 2.

5.2 Fully connected neural network training example

Our second example is to test Algorithm 1 on a neural network training problem. We perform this test on a fully connected neural network with two hidden layers of 300 and 100 nodes, followed by a fully connected output layer which fits into a soft-max cross-entropy loss. We use PyTorch to train this model on the well-known MNIST dataset with $n = 60,000$ (LeCun et al., 1998). This data set has 10 classes corresponding to 10 soft-max output nodes. We also conduct another test on the CIFAR-10 dataset ($n = 50,000$ samples and 10 classes) (Krizhevsky and Hinton, 2009). We scale the datasets by standardization. To accelerate the performance, we run Algorithm 1 with a mini-batch size of 256 instead of single sample. We use the same setting as in the previous experiment and do not use any weight decay or any data augmentation techniques.

We apply Algorithm 1 with a learning rate $\eta_i^{(t)} := \frac{\eta_t}{n}$, where $\eta_t := \frac{\gamma}{(t+\beta)^\alpha}$ to solve this training problem. We repeatedly run the algorithm 10 times and report the average results in our figures. These plots compare the algorithms on different values of $\alpha = \{1/3, 1/2, 1\}$ and $\gamma/n = \{0.01, 0.05, 0.1, 0.5\}$, respectively, on the two datasets.

Figures 3 and 4 show our comparison on the loss value $F(\tilde{w}_t)$ and test accuracy using MNIST and CIFAR-10 datasets, respectively. For each dataset, we only plot the results of experiments that yield the best training performance.

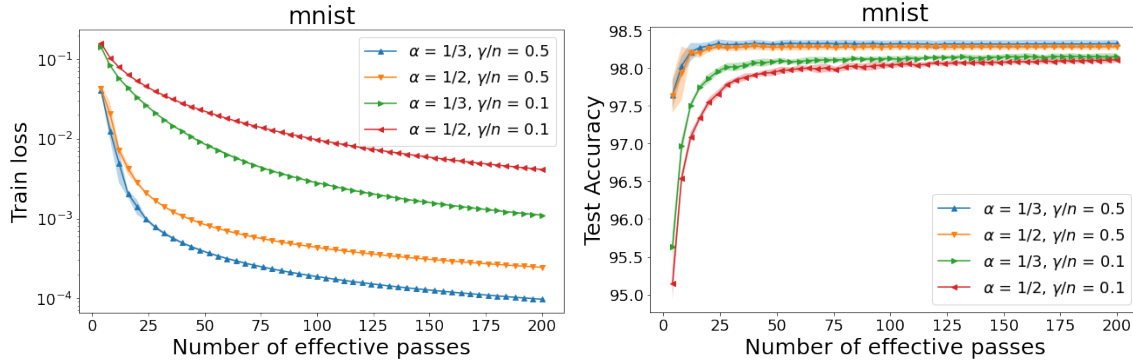


Figure 3: The behavior of the train loss $F(\tilde{w}_t)$ and the test accuracy (from the 4th epoch) produced by Algorithm 1 for solving a neural network training problem on different values of α and γ/n using the *MNIST* dataset.




Figure 4: The behavior of the train loss $F(\tilde{w}_t)$ and the test accuracy (from the 4th epoch) produced by Algorithm 1 for solving a neural network training problem on different values of α and γ/n using the *CIFAR-10* dataset.

Discussion. We observe again from Figures 3 and 4 that $\alpha = 1/3$ works best when fixing γ/n . Once we fix $\alpha := 1/3$ and test on γ/n , the ratios $\gamma/n = 0.5$ and $\gamma/n = 0.1$ give the best performance for the *MNIST* and *CIFAR-10* datasets, respectively. Similarly to the nonconvex logistic regression datasets, the train loss values and the test accuracy for *MNIST* do not deviate much in most cases. On the other hand, the *CIFAR-10* dataset’s performance is known to be noisy and we observe this phenomenon again in our experiments. The test

accuracy for the choice $\alpha := 1/3$ is not ideal in the early stage. However, it gives the best performance and becomes stable toward the end of the training process.

5.3 Comparing different shuffling schemes

In this last subsection, we compare different shuffling strategies for Algorithm 1: Random Reshuffling (RR), Shuffle Once (SO), and Incremental Gradient (IG). The behaviors of Algorithm 1 for the CIFAR-10 dataset are particularly interesting since its optimization problem is challenging to train. We experimented with the same neural network model as in the previous experiment, but with a constant learning rate $\eta_i^{(t)} = \eta_t/n := 0.05$. The train loss $F(\tilde{w}_t)$ and the test accuracy are shown in Figure 5.

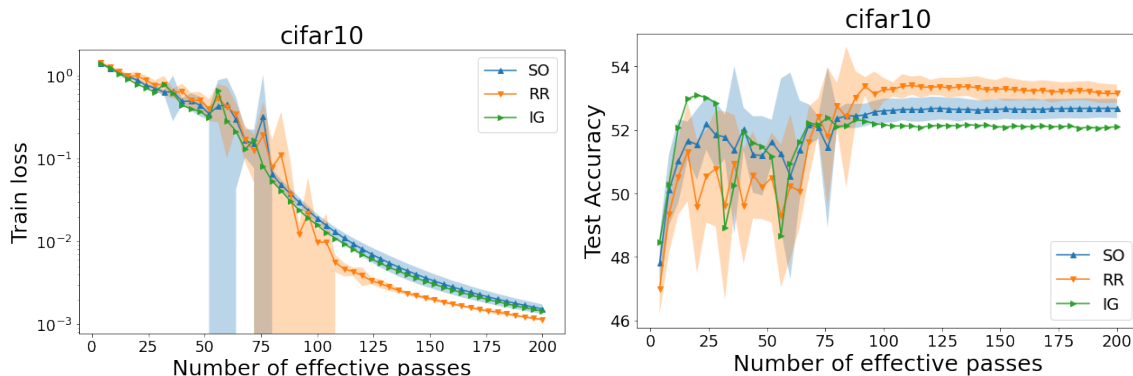


Figure 5: The train loss $F(w)$ and test accuracy produced by Algorithm 1 for the neural network training problem using different shuffling schemes on the CIFAR-10 dataset.

Note that we start all experiments with the same initialization and run these algorithms for 10 different random seeds. The only exception is the Incremental Gradient (IG) scheme where a deterministic permutation $\pi^{(t)} := \{1, 2, \dots, n\}$ is used for all $t \geq 1$. For this reason, the IG scheme has only one run and we do not see its confidence intervals in Figure 5.

Discussion. We observe that the Random Reshuffling scheme works efficiently toward the end of the training process. All shuffling schemes we test here are comparable at the early stage, but the deviation seems to decrease along the training epochs.

6. Conclusions

We have conducted an intensive convergence analysis for a wide class of shuffling-type gradient methods for solving a finite-sum minimization problem. In the strongly convex case, we have established $\mathcal{O}(T^{-2})$ convergence rate under just strong convexity of the sum function and the smoothness for any shuffling strategy. When a randomized reshuffling strategy is used, our rate has been improved to $\mathcal{O}(n^{-1}T^{-2})$, matching the results in the literature but under different assumptions. For the nonconvex case, we have proved a non-asymptotic $\mathcal{O}(T^{-2/3})$ convergence rate of our algorithm with any shuffling strategy under standard assumptions, which is significantly better than some previous works such as (Li et al., 2020). When a randomized reshuffling strategy is used, our rate has been improved to $\mathcal{O}(n^{-1/3}T^{-2/3})$, matching the recent result in (Mishchenko et al., 2020). We have also

considered these rates in both constant and diminishing learning rates, and investigated an asymptotic convergence. We believe that our results provide a unified analysis for shuffling-type algorithms using both randomized and deterministic sampling strategies, where it covers the well-known incremental gradient scheme as a special case. We have conducted different numerical experiments to highlight some theoretical aspects of our results. We believe that our analysis framework could be extended to study non-asymptotic convergence rates of SGDs and minimax algorithms, including adaptive SGD variants such as Adam (Kingma and Ba, 2014) and AdaGrad (Duchi et al., 2011) under shuffling strategies.

Acknowledgements

The authors would like to thank Trang H. Tran for her valuable comments on some technical proofs and her great help during the revision of this paper. The work of Q. Tran-Dinh has partly been supported by the National Science Foundation (NSF), grant no. DMS-1619884, the Office of Naval Research (ONR), grant no. N00014-20-1-2088 (2020-2023), and The Statistical and Applied Mathematical Sciences Institute (SAMSI).

Appendix A. Key Technical Lemmas for Convergence Analysis

This appendix provides the full proof of different technical lemmas using for our convergence analysis in the entire paper. However, let us first outline the key idea of our analysis.

A.1 The outline of our convergence analysis

Let us briefly outline the key steps of our convergence analysis to help the readers easily follow our main proofs.

- The key step of our analysis is to form a “quasi-descent” inequality between $\mathbb{E}[F(\tilde{w}_t) - F_*]$, $\eta_t \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2]$, and η_t^3 as, e.g., in (45) or (48), relying on (43) of Lemma 8.
- For the strongly convex case, we can form a “quasi-descent” inequality (see (39) of Lemma 7) between $\mathbb{E}[F(\tilde{w}_t) - F_*]$, $\mathbb{E}[\|\tilde{w}_t - w_*\|^2]$, and η_t^3 .
- To obtain such a desired bound, we need to upper bound the average deviation $\frac{1}{n} \sum_{j=0}^{n-1} \|w_j^{(t)} - w_0^{(t)}\|^2$ between the inner iterates $w_j^{(t)}$ and its epoch iterate $w_0^{(t)}$ via $\|\nabla F(\tilde{w}_{t-1})\|^2$ as in Lemma 6.
- The final step is to apply either Lemma 1 or Lemma 2 to obtain our results.

To improve our convergence rate for uniformly randomized reshuffling variants, we exploit Lemma 3 below from (Mishchenko et al., 2020, Lemma 1).

A.2 General lemmas

In order to improve our theoretical results for randomized reshuffling variants, we will use (Mishchenko et al., 2020, Lemma 1), which is stated as follows.

Lemma 3 *Let X_1, \dots, X_n be n given vectors in \mathbb{R}^d , $\bar{X} := \frac{1}{n} \sum_{i=1}^n X_i$ be their average, and $\sigma^2 := \frac{1}{n} \sum_{i=1}^n \|X_i - \bar{X}\|^2$ be their population variance. Fix any $k \in \{1, \dots, n\}$, let $X_{\pi_1}, \dots, X_{\pi_k}$ be sampled uniformly without replacement from $\{X_1, \dots, X_n\}$ and $\bar{X}_\pi :=$*

$\frac{1}{k} \sum_{i=1}^k X_{\pi_i}$ be their average. Then, we have

$$\mathbb{E}[\bar{X}_\pi] = \bar{X} \quad \text{and} \quad \mathbb{E}[\|\bar{X}_\pi - \bar{X}\|^2] = \frac{n-k}{k(n-1)} \sigma^2.$$

Now, let us prove Lemma 1 and Lemma 2 in Subsection 2.3 of the main text.

Proof [The proof of Lemma 1] Let us choose $\eta_t := \frac{q}{\rho(t+\beta)}$ for all $t \geq 1$, where $\beta \geq q-1$. Then, $0 < \rho\eta_t \leq 1$ for all $t \geq 1$. From the condition $Y_{t+1} \leq (1 - \rho \cdot \eta_t)Y_t + D \cdot \eta_t^{q+1}$ in (5), by induction, we can show that

$$Y_{t+1} \leq \prod_{i=1}^t (1 - \rho \cdot \eta_i) Y_1 + D \sum_{i=1}^t \eta_i^{q+1} \prod_{j=i+1}^t (1 - \rho \cdot \eta_j). \quad (26)$$

Using $\eta_t := \frac{q}{\rho(t+\beta)}$, we can directly compute the first coefficient as

$$\begin{aligned} C_t &:= \prod_{i=1}^t (1 - \rho \cdot \eta_i) = \prod_{i=1}^t \left(1 - \frac{q}{i+\beta}\right) = \prod_{i=1}^t \frac{i+\beta-q}{i+\beta} \\ &= \frac{\beta+1-q}{\beta+1} \cdot \frac{\beta+2-q}{\beta+2} \cdots \frac{\beta+1}{\beta+q+1} \cdots \frac{t+\beta-q}{t+\beta} \\ &= \frac{(\beta+1-q) \cdots \beta}{(t+\beta-q+1) \cdots (t+\beta)}. \end{aligned}$$

Similarly, we can show that, for any $1 \leq i \leq t$, we have

$$\begin{aligned} E_{i,t} &:= \eta_i^{q+1} \prod_{j=i+1}^t (1 - \rho \cdot \eta_j) = \frac{q^{q+1}}{\rho^{q+1}(i+\beta)^{q+1}} \cdot \frac{(i+\beta+1-q) \cdots (i+\beta)}{(t+\beta+1-q) \cdots (t+\beta)} \\ &\leq \frac{q^{q+1}}{\rho^{q+1}(t+\beta+1-q) \cdots (t+\beta)} \cdot \frac{1}{(i+\beta)}. \end{aligned}$$

Therefore, we obtain

$$\begin{aligned} \sum_{i=1}^t \eta_i^{q+1} \prod_{j=i+1}^t (1 - \rho \cdot \eta_j) &= \sum_{i=1}^t E_{i,t} \leq \frac{q^{q+1}}{\rho^{q+1}(t+\beta+1-q) \cdots (t+\beta)} \sum_{i=1}^t \frac{1}{i+\beta} \\ &\leq \frac{q^{q+1} \log(t+\beta)}{\rho^{q+1}(t+\beta-q+1) \cdots (t+\beta)} \end{aligned}$$

Substituting this sum and C_t above into (26), we finally obtain (6), i.e.:

$$Y_{t+1} \leq \frac{(\beta+1-q) \cdots \beta}{(t+\beta-q+1) \cdots (t+\beta)} \cdot Y_1 + \frac{Dq^{q+1} \log(t+\beta)}{\rho^{q+1}(t+\beta-q+1) \cdots (t+\beta)},$$

If we choose $\eta_t := \eta \in (0, \rho^{-1})$ for all $t \geq 1$, then $0 < \rho\eta_t \leq 1$, $C_t = (1 - \rho\eta)^t$, and $E_{i,t} = (1 - \rho\eta)^{t-i} \eta^{q+1}$. Hence, we get $\sum_{i=1}^t E_{i,t} = \eta^{q+1} \sum_{i=1}^t (1 - \rho\eta)^{t-i} = \frac{\eta^q [1 - (1 - \rho\eta)^t]}{\rho}$. Substituting these C_t and $E_{i,t}$ into (26), we obtain $Y_{t+1} \leq (1 - \rho\eta)^t Y_1 + \frac{D\eta^q [1 - (1 - \rho\eta)^t]}{\rho}$, which proves the first inequality of (7). Now, since $1 - \rho\eta \leq \exp(-\rho\eta)$ for $0 \leq \rho\eta \leq 1$ and $1 - (1 - \rho\eta)^t \leq 1$, we can easily prove the second inequality of (7). \blacksquare

Next, we prove the following elementary results, which will be used to prove Lemma 2.

Lemma 4 *The following statements hold:*

- (a) For any $0 \leq \nu \leq \frac{1}{2}$ and $s > 0$, we have $(s+1)^\nu - s^\nu \leq \frac{1}{2s^{1-\nu}}$.

- (b) For any $c > 0$, $\theta > 0$, $\beta > 0$, and $1 + \theta - \beta > ce^{\frac{1-c}{c}}$, the function $f(t) := \frac{\log(t+1+\theta)}{(t+\beta)^c}$ is monotonically decreasing on $[0, +\infty)$.
- (c) Suppose that f is a real-valued and monotonically decreasing function on $[a, +\infty)$ such that $f(x) \geq 0$ for all $x \in [a, +\infty)$. Then, for any integers t and t_0 such that $t \geq t_0 \geq a$, we have

$$\sum_{i=t_0+1}^t f(i) \leq \int_{t_0}^t f(x)dx \leq \sum_{i=t_0}^{t-1} f(i). \quad (27)$$

Proof (a) If $2\nu \leq 1$, then $\left(\frac{s+1}{s}\right)^{1-2\nu} \geq 1$, which is equivalent to $\frac{s+1}{s} \geq \left(\frac{s+1}{s}\right)^{2\nu}$. This leads to $(s+1)^\nu s^{1-\nu} - s^\nu (s+1)^{1-\nu} \leq 0$. Hence, we have

$$(s+1)^\nu - s^\nu = \frac{1 + (s+1)^\nu s^{1-\nu} - s^\nu (s+1)^{1-\nu}}{(s+1)^{1-\nu} + s^{1-\nu}} \leq \frac{1}{(s+1)^{1-\nu} + s^{1-\nu}} \leq \frac{1}{2s^{1-\nu}},$$

which proves assertion (a).

(b) Our goal is to show that $f'(t) < 0$ for all $t \geq 0$. We can directly compute $f'(t)$ as

$$f'(t) = (t+\beta)^{-c-1} \left[1 - \frac{1+\theta-\beta}{t+1+\theta} - c \cdot \log(t+1+\theta) \right] = (t+\beta)^{-c-1} g(t+1+\theta),$$

where $g(\tau) := 1 - \frac{1+\theta-\beta}{\tau} - c \log(\tau)$. We consider $g(\tau)$ for $\tau > 0$. It is obvious to show that $g'(\tau) = \frac{1+\theta-\beta}{\tau^2} - \frac{c}{\tau} = \frac{(1+\theta-\beta)-c\tau}{\tau^2}$ and $g''(\tau) = \frac{c\tau-2(1+\theta-\beta)}{\tau^3}$. Hence, $g'(\tau) = 0$ has a unique solution $\tau^* := \frac{1+\theta-\beta}{c} > 0$ and $g''(\tau^*) = -\frac{c^3}{(1+\theta-\beta)^2} < 0$. Consequently, g attains its unique local maximum at τ^* . Moreover, for $\tau \geq \tau^*$, we have $g'(\tau) \leq 0$. Hence, g is nonincreasing on $[\tau^*, +\infty)$, which leads to

$$g(\tau) \leq g(\tau^*) = 1 - c - c \log\left(\frac{1+\theta-\beta}{c}\right) < 0.$$

Here, the last inequality holds since $1 + \theta - \beta > ce^{\frac{1-c}{c}}$. Since $f'(t) = (t+\beta)^{-c-1}g(t+1+\theta)$, where $(t+\beta)^{-c-1} > 0$ for any $t \geq 0$ and c , we have $f'(t) < 0$ for all $t \geq 0$. Hence, f is monotonically decreasing on $[0, +\infty)$.

(c) If f is monotonically decreasing and nonnegative on $[a, +\infty)$, then $f(i+1) \leq \int_i^{i+1} f(x)dx \leq f(i)$ for any integer $i \geq a$. Hence, summing this inequality from $i := t_0$ to $t-1$, we have

$$\sum_{i=t_0+1}^t f(i) = \sum_{i=t_0}^{t-1} f(i+1) \leq \sum_{i=t_0}^{t-1} \int_i^{i+1} f(x)dx = \int_{t_0}^t f(x)dx \leq \sum_{i=t_0}^{t-1} f(i),$$

which proves (27). ■

Proof [The proof of Lemma 2] From the inequality (8) and $\eta_t := \frac{\gamma}{(t+\beta)^\alpha}$, we have

$$Z_t \leq \frac{1}{\rho \eta_t^m} (Y_t - Y_{t+1}) + \frac{D \eta_t^{q-m}}{\rho} = \frac{(t+\beta)^{\alpha m}}{\rho \gamma^m} (Y_t - Y_{t+1}) + \frac{D \gamma^{q-m}}{\rho} \cdot \frac{1}{(t+\beta)^{\alpha(q-m)}}.$$

Next, using Lemma 4(a) with $s := t + \beta$ and $\nu := m\alpha$ we have

$$(t + \beta + 1)^{\alpha m} - (t + \beta)^{\alpha m} \leq \frac{1}{2(t + \beta)^{1-\alpha m}}, \quad (28)$$

because we assume that $m\alpha \leq \frac{1}{2}$. Summing up the first inequality from $t = 1, \dots, T$ and taking average, we have

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T Z_t &\leq \frac{1}{\rho\gamma^m} \cdot \frac{1}{T} \sum_{t=1}^T (t + \beta)^{\alpha m} (Y_t - Y_{t+1}) + \frac{D\gamma^{q-m}}{\rho} \cdot \frac{1}{T} \sum_{t=1}^T \frac{1}{(t+\beta)^{\alpha(q-m)}} \\ &= \frac{1}{\rho\gamma^m} \cdot \frac{1}{T} [(1 + \beta)^{\alpha m} Y_1 - (T + \beta)^{\alpha m} Y_{T+1}] + \frac{D\gamma^{q-m}}{\rho} \cdot \frac{1}{T} \sum_{t=1}^T \frac{1}{(t+\beta)^{\alpha(q-m)}} \\ &\quad + \frac{1}{\rho\gamma^m} \cdot \frac{1}{T} \sum_{t=1}^{T-1} ((t + 1 + \beta)^{\alpha m} - (t + \beta)^{\alpha m}) Y_{t+1} \\ &\stackrel{(28)}{\leq} \frac{(1+\beta)^{\alpha m} Y_1}{\rho\gamma^m} \cdot \frac{1}{T} + \frac{1}{2\rho\gamma^m} \cdot \frac{1}{T} \sum_{t=1}^{T-1} \frac{C+H \log(t+1+\theta)}{(t+\beta)^{1-\alpha m}} + \frac{D\gamma^{q-m}}{\rho} \cdot \frac{1}{T} \sum_{t=1}^T \frac{1}{(t+\beta)^{\alpha(q-m)}} \\ &\stackrel{(27)}{\leq} \frac{(1+\beta)^{\alpha m} Y_1}{\rho\gamma^m} \cdot \frac{1}{T} + \frac{C}{2\rho\gamma^m} \cdot \frac{1}{T} \int_{t=0}^{T-1} \frac{dt}{(t+\beta)^{1-\alpha m}} + \frac{H}{2\rho\gamma^m} \cdot \frac{1}{T} \int_{t=0}^{T-1} \frac{\log(t+1+\theta)}{(t+\beta)^{1-\alpha m}} dt \\ &\quad + \frac{D\gamma^{q-m}}{\rho} \cdot \frac{1}{T} \int_{t=0}^T \frac{dt}{(t+\beta)^{\alpha(q-m)}}, \end{aligned}$$

where the second inequality follows since $0 \leq Y_t \leq C + H \log(t + \theta)$ for some $C > 0$, $H \geq 0$, and $\theta > 0$, for all $t \geq 1$, and $\alpha m \leq \frac{1}{2}$. The last inequality follows since $\frac{\log(t+1+\theta)}{(t+\beta)^{1-\alpha m}}$ is nonnegative and monotonically decreasing on $[0, \infty)$ according to Lemma 4(b) with $1 - \alpha m \geq \frac{1}{2} > 0$ and $1 + \theta - \beta > (1 - \alpha m)e^{\frac{\alpha m}{1 - \alpha m}}$, and both $\frac{1}{(t+\beta)^{1-\alpha m}}$ and $\frac{1}{(t+\beta)^{\alpha(q-m)}}$ are also nonnegative and monotonically decreasing on $[0, \infty)$. Note that

$$\begin{aligned} \int_{t=0}^{T-1} \frac{\log(t+1+\theta)}{(t+\beta)^{1-\alpha m}} dt &= \frac{1}{\alpha m} (t + \beta)^{\alpha m} \log(t + 1 + \theta) \Big|_{t=0}^{T-1} - \frac{1}{\alpha m} \int_{t=0}^{T-1} \frac{(t+\beta)^{\alpha m}}{(t+1+\theta)} dt \\ &\leq \frac{1}{\alpha m} (T - 1 + \beta)^{\alpha m} \log(T + \theta). \end{aligned}$$

Therefore, we consider two cases:

- If $\alpha(q - m) = 1$, we have

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T Z_t &\leq \frac{(1+\beta)^{\alpha m} Y_1}{\rho\gamma^m} \cdot \frac{1}{T} + \frac{C}{2\rho\alpha m\gamma^m} \cdot \frac{(T-1+\beta)^{\alpha m} - \beta^{\alpha m}}{T} \\ &\quad + \frac{H}{2\rho\alpha m\gamma^m} \cdot \frac{(T-1+\beta)^{\alpha m} \log(T+\theta)}{T} + \frac{D\gamma^{q-m}}{\rho} \cdot \frac{\log(T+\beta) - \log(\beta)}{T} \\ &\leq \frac{(1+\beta)^{\alpha m} Y_1}{\rho\gamma^m} \cdot \frac{1}{T} + \frac{C}{2\rho\alpha m\gamma^m} \cdot \frac{(T-1+\beta)^{\alpha m}}{T} \\ &\quad + \frac{H}{2\rho\alpha m\gamma^m} \cdot \frac{(T-1+\beta)^{\alpha m} \log(T+\theta)}{T} + \frac{D\gamma^{q-m}}{\rho} \cdot \frac{\log(T+\beta) - \log(\beta)}{T}. \end{aligned}$$

- If $\alpha(q - m) \neq 1$, we have

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T Z_t &\leq \frac{(1+\beta)^{\alpha m} Y_1}{\rho\gamma^m} \cdot \frac{1}{T} + \frac{C}{2\rho\alpha m\gamma^m} \cdot \frac{(T-1+\beta)^{\alpha m} - \beta^{\alpha m}}{T} + \frac{H}{2\rho\alpha m\gamma^m} \cdot \frac{(T-1+\beta)^{\alpha m} \log(T+\theta)}{T} \\ &\quad + \frac{D\gamma^{q-m}}{\rho(1-\alpha(q-m))} \cdot \frac{(T+\beta)^{1-\alpha(q-m)} - \beta^{1-\alpha(q-m)}}{T} \\ &\leq \frac{(1+\beta)^{\alpha m} Y_1}{\rho\gamma^m} \cdot \frac{1}{T} + \frac{C}{2\rho\alpha m\gamma^m} \cdot \frac{(T-1+\beta)^{\alpha m}}{T} + \frac{H}{2\rho\alpha m\gamma^m} \cdot \frac{(T-1+\beta)^{\alpha m} \log(T+\theta)}{T} \\ &\quad + \frac{D\gamma^{q-m}}{\rho(1-\alpha(q-m))} \cdot \frac{(T+\beta)^{1-\alpha(q-m)}}{T}. \end{aligned}$$

Here, the result is obtained by directly computing the integrals. Hence, (9) is proved. \blacksquare

A.3 Key estimates

This appendix provides four technical lemmas for the next steps of our analysis. However, let us first state the following facts.

- Let w_* be a stationary point of F , i.e. $F(w_*) = 0$. Then, if F is convex, then $w_* \in \arg \min_{w \in \mathbb{R}^d} F(w)$. Consequently, for any permutation $\pi^{(t)}$ of $[n]$, we have

$$\sum_{j=0}^{n-1} \nabla f(w_*; \pi^{(t)}(j+1)) = 0. \quad (29)$$

- For any $1 \leq i \leq n$, from the update of $w_i^{(t)}$ in Algorithm 1, we have

$$\begin{aligned} w_i^{(t)} &= \tilde{w}_{t-1} - \frac{\eta_t}{n} \sum_{j=0}^{i-1} \nabla f(w_j^{(t)}; \pi^{(t)}(j+1)), \\ \tilde{w}_t &= \tilde{w}_{t-1} - \frac{\eta_t}{n} \sum_{j=0}^{n-1} \nabla f(w_j^{(t)}; \pi^{(t)}(j+1)). \end{aligned} \quad (30)$$

- Let us recalled $\sigma_*^2 := \frac{1}{n} \sum_{i=0}^n \|\nabla f(w_*; i)\|^2$, the variance of F , defined by (12).

Now, we first upper bound $\frac{1}{n} \sum_{i=0}^{n-1} \|w_i^{(t)} - w_0^{(t)}\|^2$ in the following lemma. This lemma only requires Assumption 1(ii) to hold without convexity.

Lemma 5 *Suppose that Assumption 1(ii) holds for (P). Let $\{w_i^{(t)}\}$ be generated by Algorithm 1 with the learning rate $\eta_i^{(t)} := \frac{\eta_t}{n} > 0$ for a given positive sequence $\{\eta_t\}$. Then*

$$\begin{aligned} \|w_i^{(t)} - w_0^{(t)}\|^2 &\leq \frac{2L^2 \eta_t^2 \cdot i}{n^2} \sum_{j=0}^{i-1} \|w_j^{(t)} - w_*\|^2 + \frac{2\eta_t^2(n-i)}{n} \cdot \sigma_*^2, \\ \|w_i^{(t)} - w_*\|^2 &\leq 2\|w_0^{(t)} - w_*\|^2 + \frac{4L^2 \eta_t^2 \cdot i}{n^2} \cdot \sum_{j=0}^{i-1} \|w_j^{(t)} - w_*\|^2 + \frac{4\eta_t^2(n-i)\sigma_*^2}{n}. \end{aligned} \quad (31)$$

If, in addition, $0 < \eta_t \leq \frac{1}{2L}$, then, for any $1 \leq i \leq n$, we have

$$\sum_{j=0}^{i-1} \|w_j^{(t)} - w_*\|^2 \leq 4i \cdot [\|w_0^{(t)} - w_*\|^2 + 2\eta_t^2 \sigma_*^2]. \quad (32)$$

Consequently, if $0 < \eta_t \leq \frac{1}{2L}$ for all $t \geq 1$, then we have

$$\frac{1}{n} \sum_{i=0}^{n-1} \|w_i^{(t)} - w_0^{(t)}\|^2 \leq \eta_t^2 \cdot \frac{8L^2}{3} \|w_0^{(t)} - w_*\|^2 + \frac{16L^2 \sigma_*^2}{3} \cdot \eta_t^4 + 2\sigma_*^2 \cdot \eta_t^2. \quad (33)$$

Proof Using the first line of (30), the optimality condition $\nabla F(w_*) = 0$ in (a), and $(u+v)^2 \leq 2u^2 + 2v^2$ and the Cauchy-Schwarz inequality in (b), for $i \in [n]$, we can derive

$$\begin{aligned} \|w_i^{(t)} - w_0^{(t)}\|^2 &= \frac{\eta_t^2}{n^2} \left\| \sum_{j=0}^{i-1} \nabla f(w_j^{(t)}; \pi^{(t)}(j+1)) \right\|^2 \\ &\stackrel{(a)}{=} \frac{\eta_t^2}{n^2} \left\| \sum_{j=0}^{i-1} (\nabla f(w_j^{(t)}; \pi^{(t)}(j+1)) - \nabla f(w_*; \pi^{(t)}(j+1))) \right. \\ &\quad \left. - \sum_{j=i}^{n-1} \nabla f(w_*; \pi^{(t)}(j+1)) \right\|^2 \\ &\stackrel{(b)}{\leq} \frac{2\eta_t^2 \cdot i}{n^2} \sum_{j=0}^{i-1} \|\nabla f(w_j^{(t)}; \pi^{(t)}(j+1)) - \nabla f(w_*; \pi^{(t)}(j+1))\|^2 \\ &\quad + \frac{2\eta_t^2 \cdot (n-i)}{n^2} \sum_{j=i}^{n-1} \|\nabla f(w_*; \pi^{(t)}(j+1))\|^2. \end{aligned} \quad (34)$$

Using (2) and (12), we can further estimate (34) as

$$\begin{aligned} \|w_i^{(t)} - w_0^{(t)}\|^2 &\stackrel{(2)}{\leq} \frac{2L^2\eta_t^2 \cdot i}{n^2} \sum_{j=0}^{i-1} \|w_j^{(t)} - w_*\|^2 + \frac{2(n-i)\cdot\eta_t^2}{n} \cdot \frac{1}{n} \sum_{j=0}^{n-1} \|\nabla f(w_*; \pi^{(t)}(j+1))\|^2 \\ &\stackrel{(12)}{\leq} \frac{2L^2\eta_t^2 \cdot i}{n^2} \sum_{j=0}^{i-1} \|w_j^{(t)} - w_*\|^2 + \frac{2(n-i)\cdot\eta_t^2}{n} \cdot \sigma_*^2. \end{aligned}$$

This is exactly the first inequality of (31).

Next, by $\|u+v\|^2 \leq 2\|u\|^2 + 2\|v\|^2$ for any u and v , for $i \in [n]$, using the last inequality we can easily show that

$$\begin{aligned} \|w_i^{(t)} - w_*\|^2 &\leq 2\|w_0^{(t)} - w_*\|^2 + 2\|w_i^{(t)} - w_0^{(t)}\|^2 \\ &\leq 2\|w_0^{(t)} - w_*\|^2 + \eta_t^2 \cdot \frac{4iL^2}{n^2} \sum_{j=0}^{i-1} \|w_j^{(t)} - w_*\|^2 + \frac{4\eta_t^2(n-i)}{n} \cdot \sigma_*^2, \end{aligned}$$

which proves the second estimate of (31).

Now, summing up the second estimate of (31) from $j = 0$ to $j = i - 1$, we obtain

$$\begin{aligned} \sum_{j=0}^{i-1} \|w_j^{(t)} - w_*\|^2 &\leq 2i \cdot \|w_0^{(t)} - w_*\|^2 + \frac{4\eta_t^2\sigma_*^2}{n} \sum_{j=0}^{i-1} (n-j) + \frac{4L^2\eta_t^2}{n^2} \sum_{j=0}^{i-1} j \sum_{k=0}^{j-1} \|w_k^{(t)} - w_*\|^2 \\ &\leq 2i \cdot \|w_0^{(t)} - w_*\|^2 + 4\eta_t^2\sigma_*^2 \cdot i + \frac{2L^2\eta_t^2 \cdot i(i-1)}{n^2} \sum_{j=0}^{i-1} \|w_j^{(t)} - w_*\|^2 \\ &\leq 2i [\|w_0^{(t)} - w_*\|^2 + 2\eta_t^2\sigma_*^2] + 2L^2\eta_t^2 \sum_{j=0}^{i-1} \|w_j^{(t)} - w_*\|^2. \end{aligned}$$

Here, we obtain the second inequality by first rearranging the double sum and then upper bound each term. Since $0 < \eta_t \leq \frac{1}{2L}$, we have $1 - 2L^2\eta_t^2 \geq \frac{1}{2}$. Rearranging the last inequality and using the last fact, we obtain (32).

Finally, combining the first inequality of (31) and (32), we can derive that

$$\begin{aligned} \sum_{i=0}^{n-1} \|w_i^{(t)} - w_0^{(t)}\|^2 &\stackrel{(31)}{\leq} \frac{2L^2\eta_t^2}{n^2} \sum_{i=0}^{n-1} i \cdot \sum_{j=0}^{i-1} \|w_j^{(t)} - w_*\|^2 + \frac{2\sigma_*^2\eta_t^2}{n} \sum_{i=0}^{n-1} (n-i) \\ &\stackrel{(32)}{\leq} \frac{2L^2\eta_t^2}{n^2} \sum_{i=0}^{n-1} 4i^2 [\|w_0^{(t)} - w_*\|^2 + 2\eta_t^2\sigma_*^2] + \sigma_*^2\eta_t^2(n+1) \\ &\leq \frac{8L^2\eta_t^2 \cdot n}{3} \|w_0^{(t)} - w_*\|^2 + \frac{16L^2\sigma_*^2 \cdot n}{3} \cdot \eta_t^4 + 2\sigma_*^2\eta_t^2 \cdot n, \end{aligned}$$

which implies (33) after multiplying both sides by $\frac{1}{n}$. ■

Lemma 6 *Suppose that Assumption 1(ii) and Assumption 2 hold for (P). Let $\{w_i^{(t)}\}$ be generated by Algorithm 1 with any shuffling strategy $\pi^{(t)}$ and a learning rate $\eta_i^{(t)} := \frac{\eta_t}{n} > 0$ for a given positive sequence $\{\eta_t\}$ such that $0 < \eta_t \leq \frac{1}{L\sqrt{3}}$. Then, we have*

$$\sum_{j=0}^{n-1} \|w_j^{(t)} - w_0^{(t)}\|^2 \leq n\eta_t^2 \cdot \left[(3\Theta + 2) \|\nabla F(w_0^{(t)})\|^2 + 3\sigma^2 \right]. \quad (35)$$

If $\pi^{(t)}$ is uniformly sampled at random without replacement from $[n]$ and the learning rate η_t satisfies $0 < \eta_t \leq \frac{1}{L\sqrt{3}}$ for all $t \geq 1$, then we have

$$\mathbb{E} \left[\sum_{j=0}^{n-1} \|w_j^{(t)} - w_0^{(t)}\|^2 \right] \leq 2\eta_t^2 \cdot \left[(\Theta + n) \mathbb{E} [\|\nabla F(w_0^{(t)})\|^2] + \sigma^2 \right]. \quad (36)$$

Proof First, from the first line of (30), by using $\|\sum_{i=1}^3 a_i\|^2 \leq 3 \sum_{i=1}^3 \|a_i\|^2$ in (a) and the Cauchy-Schwarz inequality in (b), we can derive that

$$\begin{aligned}
 \|w_i^{(t)} - w_0^{(t)}\|^2 &\stackrel{(30)}{=} \frac{i^2 \cdot \eta_t^2}{n^2} \left\| \frac{1}{i} \sum_{j=0}^{i-1} \nabla f(w_j^{(t)}; \pi^{(t)}(j+1)) \right\|^2 \\
 &\stackrel{(a)}{\leq} \frac{3i^2 \cdot \eta_t^2}{n^2} \left[\left\| \frac{1}{i} \sum_{j=0}^{i-1} \left(\nabla f(w_0^{(t)}; \pi^{(t)}(j+1)) - \nabla F(w_0^{(t)}) \right) \right\|^2 + \|\nabla F(w_0^{(t)})\|^2 \right] \\
 &\quad + \frac{3i^2 \cdot \eta_t^2}{n^2} \left\| \frac{1}{i} \sum_{j=0}^{i-1} \left(\nabla f(w_j^{(t)}; \pi^{(t)}(j+1)) - \nabla f(w_0^{(t)}; \pi^{(t)}(j+1)) \right) \right\|^2 \quad (37) \\
 &\stackrel{(b)}{\leq} \frac{3i^2 \cdot \eta_t^2}{n^2} \frac{1}{i} \sum_{j=0}^{i-1} \left\| \nabla f(w_j^{(t)}; \pi^{(t)}(j+1)) - \nabla f(w_0^{(t)}; \pi^{(t)}(j+1)) \right\|^2 \\
 &\quad + \frac{3i^2 \cdot \eta_t^2}{n^2} \left[\frac{1}{i} \sum_{j=0}^{i-1} \left\| \nabla f(w_0^{(t)}; \pi^{(t)}(j+1)) - \nabla F(w_0^{(t)}) \right\|^2 + \|\nabla F(w_0^{(t)})\|^2 \right].
 \end{aligned}$$

Let us introduce $\Delta := \sum_{j=0}^{n-1} \|w_j^{(t)} - w_0^{(t)}\|^2$. Then, using (2) from Assumption 1(ii) and (4) from Assumption 2, we can further derive from (37) that

$$\begin{aligned}
 \|w_i^{(t)} - w_0^{(t)}\|^2 &\stackrel{(2)}{\leq} \frac{3i^2 \cdot \eta_t^2}{n^2} \left[\frac{1}{i} \sum_{j=0}^{i-1} \left\| \nabla f(w_0^{(t)}; \pi^{(t)}(j+1)) - \nabla F(w_0^{(t)}) \right\|^2 + \|\nabla F(w_0^{(t)})\|^2 \right] \\
 &\quad + \frac{3i^2 \cdot \eta_t^2}{n^2} \frac{L^2}{i} \sum_{j=0}^{i-1} \|w_j^{(t)} - w_0^{(t)}\|^2 \\
 &\leq \frac{3i^2 \cdot \eta_t^2}{n^2} \left[\frac{n}{i} \cdot \frac{1}{n} \sum_{j=0}^{n-1} \left\| \nabla f(w_0^{(t)}; \pi^{(t)}(j+1)) - \nabla F(w_0^{(t)}) \right\|^2 + \|\nabla F(w_0^{(t)})\|^2 \right] \\
 &\quad + \frac{3iL^2 \eta_t^2}{n^2} \sum_{j=0}^{i-1} \|w_j^{(t)} - w_0^{(t)}\|^2 \\
 &\stackrel{(4)}{\leq} \frac{3iL^2 \eta_t^2}{n^2} \sum_{j=0}^{i-1} \|w_j^{(t)} - w_0^{(t)}\|^2 + \frac{3i^2 \cdot \eta_t^2}{n^2} \left[\frac{n}{i} (\Theta \|\nabla F(w_0^{(t)})\|^2 + \sigma^2) + \|\nabla F(w_0^{(t)})\|^2 \right] \\
 &\leq \frac{3iL^2 \eta_t^2}{n^2} \Delta + \frac{3\eta_t^2}{n^2} \left[n \cdot i (\Theta \|\nabla F(w_0^{(t)})\|^2 + \sigma^2) + i^2 \|\nabla F(w_0^{(t)})\|^2 \right].
 \end{aligned}$$

Using this estimate and the definition of Δ , we have

$$\begin{aligned}
 \Delta &= \sum_{i=0}^{n-1} \|w_i^{(t)} - w_0^{(t)}\|^2 \\
 &\leq \frac{3L^2 \eta_t^2}{n^2} (\sum_{i=0}^{n-1} i) \Delta + \frac{3\eta_t^2}{n^2} \left[n (\Theta \|\nabla F(w_0^{(t)})\|^2 + \sigma^2) \sum_{i=0}^{n-1} i + \|\nabla F(w_0^{(t)})\|^2 \sum_{i=0}^{n-1} i^2 \right] \\
 &\leq \frac{3L^2 \eta_t^2}{2} \Delta + \frac{3n\eta_t^2}{2} (\Theta \|\nabla F(w_0^{(t)})\|^2 + \sigma^2) + n\eta_t^2 \|\nabla F(w_0^{(t)})\|^2 \\
 &\leq \frac{3L^2 \eta_t^2}{2} \Delta + \frac{n\eta_t^2}{2} \left[(3\Theta + 2) \|\nabla F(w_0^{(t)})\|^2 + 3\sigma^2 \right],
 \end{aligned}$$

where the second inequality follows since $\sum_{i=0}^{n-1} i = \frac{n(n-1)}{2} \leq \frac{n^2}{2}$ and $\sum_{i=0}^{n-1} i^2 = \frac{n(n-1)(2n-1)}{6} \leq \frac{n^3}{3}$. Rearranging the last inequality and noticing that $1 - \frac{3L^2 \eta_t^2}{2} \geq \frac{1}{2}$ due to the condition $0 < \eta_t \leq \frac{1}{\sqrt{3L}}$ on η_t , we obtain (35).

Now, let $\pi^{(t)} := (\pi^{(t)}(1), \dots, \pi^{(t)}(n))$ be sampled uniformly at random without replacement from $[n]$. For each epoch $t \geq 1$, we denote $\mathcal{F}_t := \sigma(w_0^{(1)}, \dots, w_0^{(t)})$, the σ -algebra generated by the iterates of Algorithm 1. Similar to the proof of (37), we can show that

$$\begin{aligned}
 \|w_i^{(t)} - w_0^{(t)}\|^2 &\leq \frac{3i^2 \cdot \eta_t^2}{n^2} \left\| \frac{1}{i} \sum_{j=0}^{i-1} \left(\nabla f(w_0^{(t)}; \pi^{(t)}(j+1)) - \nabla F(w_0^{(t)}) \right) \right\|^2 + \frac{3i^2 \cdot \eta_t^2}{n^2} \|\nabla F(w_0^{(t)})\|^2 \\
 &\quad + \frac{3iL^2 \eta_t^2}{n^2} \sum_{j=0}^{i-1} \|w_j^{(t)} - w_0^{(t)}\|^2.
 \end{aligned}$$

Taking expectation conditioned on \mathcal{F}_t both sides of this estimate and using Δ , we get

$$\begin{aligned} \mathbb{E} \left[\|w_i^{(t)} - w_0^{(t)}\|^2 \mid \mathcal{F}_t \right] &\leq \frac{3i^2 \cdot \eta_t^2}{n^2} \mathbb{E} \left[\left\| \frac{1}{i} \sum_{j=0}^{i-1} (\nabla f(w_0^{(t)}; \pi^{(t)}(j+1)) - \nabla F(w_0^{(t)})) \right\|^2 \mid \mathcal{F}_t \right] \\ &\quad + \frac{3i^2 \cdot \eta_t^2}{n^2} \|\nabla F(w_0^{(t)})\|^2 + \frac{3iL^2 \eta_t^2}{n^2} \mathbb{E}[\Delta \mid \mathcal{F}_t]. \end{aligned} \quad (38)$$

Applying Lemma 3 and (4), we can upper bound the first term of (38) as

$$\begin{aligned} \mathcal{T}_{[2]} &:= \mathbb{E} \left[\left\| \frac{1}{i} \sum_{j=0}^{i-1} \nabla f(w_0^{(t)}; \pi^{(t)}(j+1)) - \nabla F(w_0^{(t)}) \right\|^2 \mid \mathcal{F}_t \right] \\ &= \frac{n-i}{i(n-1)} \frac{1}{n} \sum_{j=0}^{n-1} \|\nabla f(w_0^{(t)}; j+1) - \nabla F(w_0^{(t)})\|^2 \\ &\stackrel{(4)}{\leq} \frac{i(n-i)}{i^2(n-1)} \left[\Theta \|\nabla F(w_0^{(t)})\|^2 + \sigma^2 \right]. \end{aligned}$$

Substituting this inequality into (38), we get

$$\mathbb{E} [\|w_i^{(t)} - w_0^{(t)}\|^2 \mid \mathcal{F}_t] \leq \frac{3iL^2 \eta_t^2}{n^2} \mathbb{E}[\Delta \mid \mathcal{F}_t] + \frac{3\eta_t^2}{n^2} \frac{i(n-i)}{(n-1)} \left[\Theta \|\nabla F(w_0^{(t)})\|^2 + \sigma^2 \right] + \frac{3i^2 \cdot \eta_t^2}{n^2} \|\nabla F(w_0^{(t)})\|^2.$$

Taking full expectation over \mathcal{F}_t of both sides of the last estimate, we have

$$\mathbb{E} [\|w_i^{(t)} - w_0^{(t)}\|^2] \leq \frac{3iL^2 \eta_t^2}{n^2} \mathbb{E}[\Delta] + \frac{3\eta_t^2}{n^2} \frac{i(n-i)}{(n-1)} \left[\Theta \mathbb{E} [\|\nabla F(w_0^{(t)})\|^2] + \sigma^2 \right] + \frac{3i^2 \cdot \eta_t^2}{n^2} \mathbb{E} [\|\nabla F(w_0^{(t)})\|^2].$$

Using the last estimate and the definition of Δ , we can derive that

$$\begin{aligned} \mathbb{E}[\Delta] &= \sum_{i=0}^{n-1} \mathbb{E} [\|w_i^{(t)} - w_0^{(t)}\|^2] \\ &\leq \frac{3L^2 \eta_t^2}{n^2} \cdot \mathbb{E}[\Delta] \left(\sum_{i=0}^{n-1} i \right) + \frac{3\eta_t^2}{n^2(n-1)} \cdot \left[\Theta \mathbb{E} [\|\nabla F(w_0^{(t)})\|^2] + \sigma^2 \right] \cdot \left[\sum_{i=0}^{n-1} i(n-i) \right] \\ &\quad + \frac{3\eta_t^2}{n^2} \cdot \mathbb{E} [\|\nabla F(w_0^{(t)})\|^2] \left(\sum_{i=0}^{n-1} i^2 \right) \\ &\leq \frac{3L^2 \eta_t^2}{2} \cdot \mathbb{E}[\Delta] + \eta_t^2 \cdot \left[(\Theta + n) \mathbb{E} [\|\nabla F(w_0^{(t)})\|^2] + \sigma^2 \right], \end{aligned}$$

where the second inequality follows since $\sum_{i=0}^{n-1} i = \frac{n(n-1)}{2} \leq \frac{n^2}{2}$, $\sum_{i=0}^{n-1} i^2 = \frac{n(n-1)(2n-1)}{6} \leq \frac{n^3}{3}$, and $\sum_{i=0}^{n-1} i(n-i) = \frac{(n-1)n(n+1)}{6} \leq \frac{n^2(n-1)}{3}$. Now, since $0 < \eta_t \leq \frac{1}{\sqrt{3}L}$, we have $1 - \frac{3L^2 \eta_t^2}{2} \geq \frac{1}{2}$. Rearranging the last inequality and using this fact, we finally get (36). \blacksquare

Let us improve Lemma 5 above by using a randomized reshuffling strategy, the convexity of each $f(\cdot; i)$ for $i \in [n]$, and the strong convexity of F .

Lemma 7 *Suppose that Assumption 1(ii) holds and each $f(\cdot; i)$ is convex for $i \in [n]$. Let $\{w_i^{(t)}\}$ be generated by Algorithm 1 and σ_*^2 be defined by (12). Let $\pi^{(t)} := (\pi^{(t)}(1), \dots, \pi^{(t)}(n))$ be sampled uniformly at random without replacement from $[n]$. Then, if we choose η_t such that $0 < \eta_t \leq \frac{\sqrt{5}-1}{2L}$, then, for any $t \geq 1$, we have*

$$\mathbb{E} [\|\tilde{w}_t - w_*\|^2] \leq \mathbb{E} [\|\tilde{w}_{t-1} - w_*\|^2] - 2\eta_t \cdot [F(\tilde{w}_{t-1}) - F(w_*)] + \frac{2L\eta_t^3 \sigma_*^2}{3n}. \quad (39)$$

Proof Using the first line of (30), with the same proof as of (34), we have

$$\|w_i^{(t)} - \tilde{w}_{t-1}\|^2 \leq \frac{2\eta_t^2 i}{n^2} \sum_{j=0}^{i-1} \|\nabla f(w_j^{(t)}; \pi^{(t)}(j+1)) - \nabla f(w_*; \pi^{(t)}(j+1))\|^2 + \frac{2\eta_t^2 B_i^*}{n^2},$$

where $B_i^* := \|\sum_{j=i}^{n-1} \nabla f(w_*; \pi^{(t)}(j+1))\|^2$. Summing up this inequality from $i := 0$ to $i := n-1$, we can derive that

$$\begin{aligned} \sum_{i=0}^{n-1} \|w_i^{(t)} - \tilde{w}_{t-1}\|^2 &\leq \frac{2\eta_t^2}{n^2} \sum_{i=0}^{n-1} i \cdot \sum_{j=0}^{i-1} \|\nabla f(w_j^{(t)}; \pi^{(t)}(j+1)) - \nabla f(w_*; \pi^{(t)}(j+1))\|^2 + \frac{2\eta_t^2}{n^2} \sum_{i=0}^{n-1} B_i^* \\ &\leq \eta_t^2 \sum_{j=0}^{n-1} \|\nabla f(w_j^{(t)}; \pi^{(t)}(j+1)) - \nabla f(w_*; \pi^{(t)}(j+1))\|^2 + \frac{2\eta_t^2 B^*}{n^2}. \end{aligned} \quad (40)$$

where $B^* := \sum_{i=0}^{n-1} B_i^* = \sum_{i=0}^{n-1} \|\sum_{j=i}^{n-1} \nabla f(w_*; \pi^{(t)}(j+1))\|^2$.

Next, using the second line of (30) and the Cauchy-Schwarz inequality, we have

$$\begin{aligned} \|\tilde{w}_t - w_*\|^2 &= \|\tilde{w}_{t-1} - w_*\|^2 + \frac{2\eta_t}{n} \sum_{j=0}^{n-1} \langle \nabla f(w_j^{(t)}; \pi^{(t)}(j+1)), w_* - \tilde{w}_{t-1} \rangle \\ &\quad + \frac{\eta_t^2}{n^2} \|\sum_{j=0}^{n-1} (\nabla f(w_j^{(t)}; \pi^{(t)}(j+1)) - \nabla f(w_*; \pi^{(t)}(j+1)))\|^2 \\ &\leq \|\tilde{w}_{t-1} - w_*\|^2 + \frac{2\eta_t}{n} \sum_{j=0}^{n-1} \langle \nabla f(w_j^{(t)}; \pi^{(t)}(j+1)), w_* - \tilde{w}_{t-1} \rangle \\ &\quad + \frac{\eta_t^2}{n} \sum_{j=0}^{n-1} \|\nabla f(w_j^{(t)}; \pi^{(t)}(j+1)) - \nabla f(w_*; \pi^{(t)}(j+1))\|^2. \end{aligned} \quad (41)$$

We can upper bound the second term on the right-hand side of (41) as follows:

$$\begin{aligned} \mathcal{T}_{[1]} &:= \sum_{j=0}^{n-1} \langle \nabla f(w_j^{(t)}; \pi^{(t)}(j+1)), w_* - \tilde{w}_{t-1} \rangle \\ &= \sum_{j=0}^{n-1} \langle \nabla f(w_j^{(t)}; \pi^{(t)}(j+1)), w_* - w_j^{(t)} \rangle + \sum_{j=0}^{n-1} \langle \nabla f(w_j^{(t)}; \pi^{(t)}(j+1)), w_j^{(t)} - \tilde{w}_{t-1} \rangle \\ &\stackrel{(a)}{\leq} \sum_{j=0}^{n-1} \langle \nabla f(w_j^{(t)}; \pi^{(t)}(j+1)), w_* - w_j^{(t)} \rangle + \frac{L}{2} \sum_{j=0}^{n-1} \|w_j^{(t)} - \tilde{w}_{t-1}\|^2 \\ &\quad + \sum_{j=0}^{n-1} [f(w_j^{(t)}; \pi^{(t)}(j+1)) - f(\tilde{w}_{t-1}; \pi^{(t)}(j+1))] \\ &= -\sum_{j=0}^{n-1} [f(w_*; \pi^{(t)}(j+1)) - f(w_j^{(t)}; \pi^{(t)}(j+1))] - \langle \nabla f(w_j^{(t)}; \pi^{(t)}(j+1)), w_* - w_j^{(t)} \rangle \\ &\quad + \frac{L}{2} \sum_{j=0}^{n-1} \|w_j^{(t)} - \tilde{w}_{t-1}\|^2 + \sum_{j=0}^{n-1} [f(w_*; \pi^{(t)}(j+1)) - f(\tilde{w}_{t-1}; \pi^{(t)}(j+1))] \\ &\stackrel{(b)}{\leq} \frac{L}{2} \sum_{j=0}^{n-1} \|w_j^{(t)} - \tilde{w}_{t-1}\|^2 - \frac{1}{2L} \sum_{j=0}^{n-1} \|\nabla f(w_j^{(t)}; \pi^{(t)}(j+1)) - \nabla f(w_*; \pi^{(t)}(j+1))\|^2 \\ &\quad - n[F(\tilde{w}_{t-1}) - F(w_*)]. \end{aligned}$$

Here, we have used in (a) the following estimate

$$\langle \nabla f(w_j^{(t)}; \cdot), w_j^{(t)} - \tilde{w}_{t-1} \rangle \leq f(w_j^{(t)}; \cdot) - f(\tilde{w}_{t-1}; \cdot) + \frac{L}{2} \|w_j^{(t)} - \tilde{w}_{t-1}\|^2$$

and in (b) the following two estimates:

$$f(w_*; \cdot) - f(w_j^{(t)}; \cdot) - \langle \nabla f(w_j^{(t)}; \cdot), w_* - w_j^{(t)} \rangle \geq \frac{1}{2L} \|\nabla f(w_j^{(t)}; \cdot) - \nabla f(w_*; \cdot)\|^2,$$

$$\text{and } \sum_{j=0}^{n-1} [f(w_*; \pi^{(t)}(j+1)) - f(\tilde{w}_{t-1}; \pi^{(t)}(j+1))] = n[F(w_*) - F(\tilde{w}_{t-1})].$$

Substituting (40) into $\mathcal{T}_{[1]}$, we can further upper bound it as

$$\begin{aligned} \mathcal{T}_{[1]} &\leq \left(\frac{L\eta_t^2}{2} - \frac{1}{2L} \right) \sum_{j=0}^{n-1} \|\nabla f(w_j^{(t)}; \pi^{(t)}(j+1)) - \nabla f(w_*; \pi^{(t)}(j+1))\|^2 \\ &\quad - n [F(\tilde{w}_{t-1}) - F(w_*)] + \frac{L\eta_t^2 B^*}{n^2}. \end{aligned}$$

Using this upper bound of $\mathcal{T}_{[1]}$ into (41), we get

$$\begin{aligned} \|\tilde{w}_t - w_*\|^2 &\leq \|\tilde{w}_{t-1} - w_*\|^2 - 2\eta_t [F(\tilde{w}_{t-1}) - F(w_*)] + \frac{2L\eta_t^3 B^*}{n^3} \\ &\quad + \left(\frac{L\eta_t^3}{n} - \frac{\eta_t}{Ln} + \frac{\eta_t^2}{n} \right) \sum_{j=0}^{n-1} \|\nabla f(w_j^{(t)}; \pi^{(t)}(j+1)) - \nabla f(w_*; \pi^{(t)}(j+1))\|^2. \end{aligned}$$

Let us choose $\eta_t > 0$ such that $L^2\eta_t^3 + L\eta_t^2 - \eta_t \leq 0$ (or equivalently, $0 < \eta_t \leq \frac{\sqrt{5}-1}{2L}$). Then, this estimate reduces to

$$\|\tilde{w}_t - w_*\|^2 \leq \|\tilde{w}_{t-1} - w_*\|^2 - 2\eta_t [F(\tilde{w}_{t-1}) - F(w_*)] + \frac{2L\eta_t^3 B^*}{n^3}. \quad (42)$$

If $\pi^{(t)} := (\pi^{(t)}(1), \dots, \pi^{(t)}(n))$ is uniformly sampled at random without replacement from $[n]$, then by Lemma 3, we have $\mathbb{E} \left[\frac{1}{n-i} \sum_{j=i}^{n-1} \nabla f(w_*; \pi^{(t)}(j+1)) \right] = \nabla F(w_*)$ and

$$\begin{aligned} \mathbb{E}[B^*] &= \mathbb{E} \left[\sum_{i=0}^{n-1} \left\| \sum_{j=i}^{n-1} \nabla f(w_*; \pi^{(t)}(j+1)) \right\|^2 \right] \\ &= \sum_{i=0}^{n-1} (n-i)^2 \mathbb{E} \left[\left\| \frac{1}{n-i} \sum_{j=i}^{n-1} \nabla f(w_*; \pi^{(t)}(j+1)) \right\|^2 \right] \\ &= \sum_{i=0}^{n-1} (n-i)^2 \mathbb{E} \left[\left\| \frac{1}{n-i} \sum_{j=i}^{n-1} \nabla f(w_*; \pi^{(t)}(j+1)) - \nabla F(w_*) \right\|^2 \right] \\ &= \sum_{i=0}^{n-1} \frac{(n-i)^2 i}{(n-i)(n-1)} \frac{1}{n} \sum_{j=0}^{n-1} \|\nabla f(w_*; \pi^{(t)}(j+1))\|^2 \\ &\stackrel{(12)}{=} \frac{\sigma_*^2}{n-1} \sum_{i=0}^{n-1} i(n-i) \\ &= \frac{n(n+1)\sigma_*^2}{6}. \end{aligned}$$

Taking expectation both sides of (42) and using this upper bound of B^* , we obtain (39). ■

Finally, we will need the following bound on F in the sequel.

Lemma 8 *Suppose that Assumption 1(ii) holds for (P). Let $\{w_i^{(t)}\}$ be generated by Algorithm 1 with any shuffling strategy $\pi^{(t)}$ and a learning rate $\eta_i^{(t)} := \frac{\eta_t}{n} > 0$ for a given positive sequence $\{\eta_t\}$ such that $0 < \eta_t \leq \frac{1}{L}$. Then, for any $t \geq 1$, we have*

$$F(w_0^{(t+1)}) \leq F(w_0^{(t)}) - \frac{\eta_t}{2} \|\nabla F(w_0^{(t)})\|^2 + \frac{L^2\eta_t}{2n} \sum_{i=0}^{n-1} \|w_i^{(t)} - w_0^{(t)}\|^2. \quad (43)$$

Proof Since F is L -smooth by Assumption 1(ii), we can derive

$$\begin{aligned}
 F(w_0^{(t+1)}) &\stackrel{(3)}{\leq} F(w_0^{(t)}) + \nabla F(w_0^{(t)})^\top (w_0^{(t+1)} - w_0^{(t)}) + \frac{L}{2} \|w_0^{(t+1)} - w_0^{(t)}\|^2 \\
 &\stackrel{(30)}{=} F(w_0^{(t)}) - \eta_t \nabla F(w_0^{(t)})^\top \left(\frac{1}{n} \sum_{i=0}^{n-1} \nabla f(w_i^{(t)}; \pi^{(t)}(i+1)) \right) \\
 &\quad + \frac{L\eta_t^2}{2} \left\| \frac{1}{n} \sum_{i=0}^{n-1} \nabla f(w_i^{(t)}; \pi^{(t)}(i+1)) \right\|^2 \\
 &\stackrel{(a)}{=} F(w_0^{(t)}) - \frac{\eta_t}{2} \|\nabla F(w_0^{(t)})\|^2 + \frac{\eta_t}{2} \left\| \nabla F(w_0^{(t)}) - \frac{1}{n} \sum_{i=0}^{n-1} \nabla f(w_i^{(t)}; \pi^{(t)}(i+1)) \right\|^2 \\
 &\quad - \frac{\eta_t}{2} (1 - L\eta_t) \left\| \frac{1}{n} \sum_{i=0}^{n-1} \nabla f(w_i^{(t)}; \pi^{(t)}(i+1)) \right\|^2 \\
 &\stackrel{(b)}{\leq} F(w_0^{(t)}) + \frac{\eta_t}{2} \left\| \frac{1}{n} \sum_{i=0}^{n-1} \nabla f(w_0^{(t)}; \pi^{(t)}(i+1)) - \frac{1}{n} \sum_{i=0}^{n-1} \nabla f(w_i^{(t)}; \pi^{(t)}(i+1)) \right\|^2 \\
 &\quad - \frac{\eta_t}{2} \|\nabla F(w_0^{(t)})\|^2 \\
 &\stackrel{(c)}{\leq} F(w_0^{(t)}) + \frac{\eta_t}{2n} \sum_{i=0}^{n-1} \left\| \nabla f(w_0^{(t)}; \pi^{(t)}(i+1)) - \nabla f(w_i^{(t)}; \pi^{(t)}(i+1)) \right\|^2 \\
 &\quad - \frac{\eta_t}{2} \|\nabla F(w_0^{(t)})\|^2 \\
 &\stackrel{(2)}{\leq} F(w_0^{(t)}) - \frac{\eta_t}{2} \|\nabla F(w_0^{(t)})\|^2 + \frac{L^2\eta_t}{2n} \sum_{i=0}^{n-1} \|w_i^{(t)} - w_0^{(t)}\|^2,
 \end{aligned}$$

where (a) follows from $u^\top v = \frac{1}{2}(\|u\|^2 + \|v\|^2 - \|u - v\|^2)$, (b) follows from the fact that $\eta_t \leq \frac{1}{L}$, and (c) is from the Cauchy-Schwarz inequality. \blacksquare

Appendix B. Convergence Analysis for Strongly Convex Case

In this section, we present the full proof of the results in the main text of Section 3.

B.1 Proofs of Theorem 1 and Theorem 2: The strongly convex case

Proof [The proof of Theorem 1] Using (33), we can further estimate (43) as follows:

$$\begin{aligned}
 F(w_0^{(t+1)}) &\leq F(w_0^{(t)}) - \frac{\eta_t}{2} \|\nabla F(w_0^{(t)})\|^2 + \frac{L^2\eta_t}{2} \frac{1}{n} \sum_{i=0}^{n-1} \|w_i^{(t)} - w_0^{(t)}\|^2 \\
 &\stackrel{(33)}{\leq} F(w_0^{(t)}) - \frac{\eta_t}{2} \|\nabla F(w_0^{(t)})\|^2 \\
 &\quad + \frac{L^2\eta_t}{2} \left(\eta_t^2 \cdot \frac{8L^2}{3} \|w_0^{(t)} - w_*\|^2 + \eta_t^4 \cdot \frac{16L^2\sigma_*^2}{3} + \eta_t^2 \cdot 2\sigma_*^2 \right) \\
 &\stackrel{(11)}{\leq} F(w_0^{(t)}) - \mu\eta_t [F(w_0^{(t)}) - F(w_*)] + \frac{8L^4\eta_t^3}{3\mu} [F(w_0^{(t)}) - F(w_*)] \\
 &\quad + \frac{8L^4\sigma_*^2}{3} \cdot \eta_t^5 + L^2\sigma_*^2 \cdot \eta_t^3,
 \end{aligned}$$

Subtracting $F(w_*)$ from both sides of the last inequality, we can further derive

$$F(w_0^{(t+1)}) - F(w_*) \leq \left[1 - \eta_t \left(\mu - \frac{8L^4}{3\mu} \eta_t^2 \right) \right] [F(w_0^{(t)}) - F(w_*)] + L^2\sigma_*^2\eta_t^3 \left(1 + \frac{8L^2\eta_t^2}{3} \right). \quad (44)$$

Now, assume that $0 < \eta_t < \sqrt{\frac{3}{8} \frac{\mu}{L^2}}$. Then, one can show that

$$\mu - \frac{4L^4}{3\mu} \eta_t^2 \geq \mu - \frac{2}{3}\mu = \frac{\mu}{3} > 0 \quad \text{and} \quad \frac{8L^4 \sigma_*^2}{3} \eta_t^2 < \frac{8L^4 \sigma_*^2}{3} \cdot \frac{3}{8} \frac{\mu^2}{L^4} = \mu^2 \sigma_*^2.$$

Using these bounds into (44), we can further upper bound it as

$$F(w_0^{(t+1)}) - F(w_*) \leq \left(1 - \frac{\mu}{3} \eta_t\right) [F(w_0^{(t)}) - F(w_*)] + \eta_t^3 (\mu^2 + L^2) \sigma_*^2. \quad (45)$$

Note that we have imposed $\eta_t \leq \min\left\{\frac{1}{2L}, \sqrt{\frac{3}{8} \frac{\mu}{L^2}}\right\}$ due to $\eta_t \leq \frac{1}{2L}$ in Lemma 5.

Now, let us define $Y_t := F(w_0^{(t)}) - F(w_*) = F(\tilde{w}_{t-1}) - F(w_*) \geq 0$, $\rho := \frac{\mu}{3}$, and $D := (\mu^2 + L^2) \sigma_*^2$. The estimate (45) becomes

$$Y_{t+1} \leq (1 - \rho \cdot \eta_t) Y_t + D \eta_t^3.$$

Applying (7) of Lemma 1 with $q = 2$ and $\eta := \frac{2 \log(T)}{\rho T} = \frac{6 \log(T)}{\mu T}$, we obtain

$$\begin{aligned} Y_{T+1} &\leq (1 - \rho \eta)^T Y_1 + \frac{D \eta^2 [1 - (1 - \rho \eta)^T]}{\rho} \leq Y_1 \exp(-\rho \eta T) + \frac{D \eta^2}{\rho} \\ &= [F(\tilde{w}_0) - F_*] \exp(-2 \log(T)) + \frac{54(\mu^2 + L^2) \sigma_*^2 \log(T)^2}{\mu^3 T^2}. \end{aligned}$$

This estimate leads to $F(\tilde{w}_T) - F(w_*) \leq \frac{[F(\tilde{w}_0) - F(w_*)]}{T^2} + \frac{54(\mu^2 + L^2) \sigma_*^2 \log(T)^2}{\mu^3 T^2}$, which is exactly (13) after taking expectation. To guarantee $\eta_t = \frac{6 \log(T)}{\mu T} \leq \min\left\{\frac{1}{2L}, \sqrt{\frac{3}{8} \frac{\mu}{L^2}}\right\}$, we need to choose T such that $\frac{\log(T)}{T} \leq \min\left\{\frac{\mu}{12L}, \frac{\mu^2}{6L^2} \sqrt{\frac{3}{8}}\right\}$. This condition holds if $T \geq 12\kappa^2 \log(T)$.

Next, under Assumption 2 and the randomized reshuffling strategy of $\pi^{(t)}$, using (36), we can further estimate (43) as follows:

$$\begin{aligned} \mathbb{E}[F(w_0^{(t+1)})] &\leq \mathbb{E}[F(w_0^{(t)})] - \frac{\eta_t}{2} \mathbb{E}[\|\nabla F(w_0^{(t)})\|^2] + \frac{L^2 \eta_t}{2n} \sum_{i=0}^{n-1} \mathbb{E}[\|w_i^{(t)} - w_0^{(t)}\|^2] \\ &\stackrel{(36)}{\leq} \mathbb{E}[F(w_0^{(t)})] - \frac{\eta_t}{2} \mathbb{E}[\|\nabla F(w_0^{(t)})\|^2] + \frac{L^2 \eta_t^3}{n} \left[(\Theta + n) \mathbb{E}[\|\nabla F(w_0^{(t)})\|^2] + \sigma^2 \right] \\ &\stackrel{(11)}{\leq} \mathbb{E}[F(w_0^{(t)})] - \mu \eta_t \mathbb{E}[F(w_0^{(t)}) - F(w_*)] + \frac{2L^2(\Theta+n)\eta_t^3}{n\mu} \mathbb{E}[F(w_0^{(t)}) - F(w_*)] \\ &\quad + \frac{L^2 \sigma_*^2 \eta_t^3}{n}. \end{aligned}$$

Subtracting $F(w_*)$ from both sides of the last inequality, we can further derive

$$\mathbb{E}[F(w_0^{(t+1)}) - F(w_*)] \leq \left[1 - \eta_t \left(\mu - \frac{2L^2(\Theta+n)\eta_t^2}{n\mu}\right)\right] \mathbb{E}[F(w_0^{(t)}) - F(w_*)] + \frac{L^2 \sigma^2 \eta_t^3}{n}. \quad (46)$$

Now, assume that $0 < \eta_t \leq \frac{\mu}{2L\sqrt{\Theta/n+1}} = \frac{1}{2\kappa\sqrt{1+\Theta/n}}$. Then, one can show that $\mu - \frac{2L^2(\Theta+n)\eta_t^2}{n\mu} \geq \mu - \frac{1}{2}\mu = \frac{\mu}{2} > 0$. Using this bound into (46) and noticing that $\tilde{w}_t = w_0^{(t+1)}$, we can further upper bound it as

$$\mathbb{E}[F(\tilde{w}_t) - F(w_*)] \leq \left(1 - \frac{\mu}{2} \eta_t\right) \mathbb{E}[F(\tilde{w}_{t-1}) - F(w_*)] + \frac{L^2 \sigma^2}{n} \cdot \eta_t^3. \quad (47)$$

Note that we have imposed $\eta_t \leq \min \left\{ \frac{1}{\sqrt{3}L}, \frac{1}{2\kappa\sqrt{1+\Theta/n}} \right\}$ due to $\eta_t \leq \frac{1}{\sqrt{3}L}$ in Lemma 6.

Let $Y_t := \mathbb{E}[F(\tilde{w}_{t-1}) - F(w_*)]$, $\eta_t := \eta > 0$ be fixed for all $t \geq 1$, and $\rho := \frac{\mu}{2}$. Then, from (47), we have $Y_{t+1} \leq (1 - \rho\eta)Y_t + \frac{L^2\sigma^2\eta^3}{n}$. Applying (7) of Lemma 1 with $q = 2$ and $\eta := \frac{4\log(\sqrt{n}T)}{\mu T}$, we obtain

$$\begin{aligned} Y_{T+1} &= \mathbb{E}[F(\tilde{w}_T) - F(w_*)] \leq (1 - \rho\eta)^T Y_1 + \frac{2L^2\sigma^2\eta^2[1-(1-\rho\eta)^T]}{n\mu} \leq Y_1 \exp(-\mu\eta T) + \frac{2L^2\sigma^2\eta^2}{n\mu} \\ &= \mathbb{E}[F(\tilde{w}_0) - F(w_*)] \exp(-2\log(\sqrt{n}T)) + \frac{2L^2\sigma^2\log(\sqrt{n}T)^2}{\mu^3 n T^2}, \end{aligned}$$

which is exactly (14). Note that to guarantee $\rho\eta = \frac{2\log(\sqrt{n}T)}{T} \leq 1$ and $\eta = \frac{4\log(\sqrt{n}T)}{\mu T} \leq \frac{1}{2\kappa\sqrt{1+\Theta/n}}$, we need to choose $T \geq \frac{8L\sqrt{\Theta/n+1}}{\mu^2} \log(\sqrt{n}T)$.

Finally, if $\pi^{(t)}$ is sampled at random without replacement and each $f(\cdot; i)$ is convex for $i \in [n]$, from (39) and $F(\tilde{w}_{t-1}) - F(w_*) \geq \frac{\mu}{2}\|\tilde{w}_{t-1} - w_*\|^2$, we have

$$\mathbb{E}[\|\tilde{w}_t - w_*\|^2] \leq (1 - \mu\eta_t)\mathbb{E}[\|\tilde{w}_{t-1} - w_*\|^2] + \frac{2L\eta_t^3\sigma_*^2}{3n}.$$

Let us denote $Y_t := \mathbb{E}[\|\tilde{w}_{t-1} - w_*\|^2]$. Then, the last inequality can be written as $Y_{t+1} \leq (1 - \mu\eta)Y_t + \frac{2L\eta^3\sigma_*^2}{3n}$. Applying (7) of Lemma 1 with $q = 2$ and $\eta := \frac{2\log(\sqrt{n}T)}{\mu T}$, we obtain

$$\begin{aligned} Y_{T+1} &= \mathbb{E}[\|\tilde{w}_T - w_*\|^2] \leq (1 - \mu\eta)^T Y_1 + \frac{2L\sigma_*^2\eta^2[1-(1-\mu\eta)^T]}{3n\mu} \leq Y_1 \exp(-\mu\eta T) + \frac{2L\sigma_*^2\eta^2}{3n\mu} \\ &= \mathbb{E}[\|\tilde{w}_0 - w_*\|^2] \exp(-2\log(\sqrt{n}T)) + \frac{8L\sigma_*^2\log(\sqrt{n}T)^2}{3\mu^3 n T^2}, \end{aligned}$$

which is exactly (15). Note that to guarantee $\mu\eta = \frac{2\log(\sqrt{n}T)}{T} \leq 1$ and $\eta = \frac{2\log(\sqrt{n}T)}{\mu T} \leq \frac{\sqrt{5}-1}{2L}$, we need to choose T such that $\frac{\log(T\sqrt{n})}{T} \leq \min \left\{ \frac{1}{2}, \frac{(\sqrt{5}-1)\mu}{2L} \right\}$. \blacksquare

Proof [The proof of Theorem 2] Similar to the proof of Theorem 1, we define $Y_t := \mathbb{E}[F(\tilde{w}_{t-1}) - F(w_*)] \geq 0$, $\rho := \frac{\mu}{3}$, and $D := (\mu^2 + L^2)\sigma_*^2$. The estimate (45) implies $Y_{t+1} \leq (1 - \rho\eta)Y_t + D\eta^3$. Moreover, since $\eta_t = \frac{6}{\mu(t+\beta)} = \frac{2}{\rho(t+\beta)}$ for $\beta \geq 1$, apply Lemma 1 with $q = 2$, we obtain

$$Y_{t+1} \leq \frac{\beta(\beta-1)}{(t+\beta-1)(t+\beta)} Y_1 + \frac{8D\log(t+\beta+1)}{\rho^3(t+\beta-1)(t+\beta)},$$

which leads to (16) after substituting $Y_{t+1} := \mathbb{E}[F(\tilde{w}_t) - F(w_*)]$, $Y_1 := F(\tilde{w}_0) - F(w_*)$, $D := (\mu^2 + L^2)\sigma_*^2$, and $\rho := \frac{\mu}{3}$ into the last estimate. However, to guarantee $\eta_t = \frac{6}{\mu(t+\beta)} \leq \min \left\{ \frac{1}{2L}, \sqrt{\frac{3}{4}} \frac{\mu}{L^2} \right\}$, we need to impose $\frac{6}{\beta+1} \leq \min \left\{ \frac{\mu}{2L}, \sqrt{\frac{3}{4}} \frac{\mu^2}{L^2} \right\}$, which holds if $\beta \geq 12\kappa^2 - 1$.

To prove (16), we use (39) from Lemma 7 and $F(\tilde{w}_{t-1}) - F(w_*) \geq \frac{\mu}{2}\|\tilde{w}_{t-1} - w_*\|^2$ to get

$$\mathbb{E}[\|\tilde{w}_t - w_*\|^2] \leq (1 - \mu\eta_t)\mathbb{E}[\|\tilde{w}_{t-1} - w_*\|^2] + \frac{2L\eta_t^3\sigma_*^2}{3n}.$$

By letting $\eta_t := \frac{2}{t+\beta}$, $Y_{t+1} := \mathbb{E}[\|\tilde{w}_t - w_*\|^2]$, $\rho := \mu$, and $D := \frac{2L\sigma_*^2}{3n}$ for all $t \geq 1$. The last estimate becomes

$$Y_{t+1} \leq (1 - \rho\eta_t)Y_t + D\eta_t^3, \quad \forall t \geq 1.$$

By applying Lemma 1 with $q = 2$ we have $Y_{t+1} \leq \frac{\beta(\beta-1)}{(t+\beta-1)(t+\beta)}Y_1 + \frac{8D}{\rho^3(t+\beta-1)(t+\beta)}$. In order to guarantee that $\eta_t = \frac{2}{t+\beta} \leq \frac{\sqrt{5}-1}{2L}$ for all $t \geq 1$, we need to choose $\beta \geq \frac{4L}{\sqrt{5}-1} - 1$. However, since $L \leq \frac{\sqrt{5}-1}{2}$ and $\beta \geq 1$, this condition automatically holds. \blacksquare

Appendix C. Convergence Analysis for Nonconvex Case

In this appendix, we provide the full proofs of the results in Section 4.

C.1 Proofs of Theorem 3, Corollary 1, and Corollary 2

Proof [The proof of Theorem 3] First, using (35) into (43), we can derive that

$$\begin{aligned} F(w_0^{(t+1)}) &\stackrel{(35)}{\leq} F(w_0^{(t)}) - \frac{\eta_t}{2} \|\nabla F(w_0^{(t)})\|^2 + \frac{L^2\eta_t^3}{2} \left[(3\Theta + 2) \|\nabla F(w_0^{(t)})\|^2 + 3\sigma^2 \right] \\ &= F(w_0^{(t)}) - \frac{\eta_t}{2} (1 - L^2\eta_t^2(3\Theta + 2)) \|\nabla F(w_0^{(t)})\|^2 + \frac{3L^2\sigma^2\eta_t^3}{2} \\ &\leq F(w_0^{(t)}) - \frac{\eta_t}{4} \|\nabla F(w_0^{(t)})\|^2 + \frac{3L^2\sigma^2\eta_t^3}{2}. \end{aligned}$$

where the last inequality follows since $\eta_t^2 \leq \frac{1}{2(3\Theta+2)L^2}$. Note that $\tilde{w}_t = w_0^{(t+1)}$ and $\tilde{w}_{t-1} = w_0^{(t)}$ in Algorithm 1, the last estimate becomes

$$F(\tilde{w}_t) \leq F(\tilde{w}_{t-1}) - \frac{\eta_t}{4} \|\nabla F(\tilde{w}_{t-1})\|^2 + \frac{3L^2\sigma^2\eta_t^3}{2}. \quad (48)$$

Using $\eta_t := \eta$ into (48) and rearranging its result, then taking expectation we end up with

$$\mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] \leq \frac{4}{\eta} \mathbb{E}[F(\tilde{w}_{t-1}) - F(\tilde{w}_t)] + 6L^2\sigma^2\eta^2.$$

Taking average the last inequality from $t := 1$ to $t := T$ and using the fact that $\mathbb{E}[F(\tilde{w}_t)] \geq F_*$ from Assumption 1(i), we finally obtain

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] \leq \frac{4}{T\eta} [F(\tilde{w}_0) - F_*] + 6L^2\sigma^2\eta^2,$$

which is exactly (18). Note that, since \tilde{w}_0 is deterministic, we drop the expectation on the right-hand side of this estimate (18).

If $\pi^{(t)}$ is sampled uniformly at random without replacement from $[n]$, then taking expectation both sides of (43), and then using (36), we obtain

$$\begin{aligned} \mathbb{E}[F(w_0^{(t+1)})] &\leq \mathbb{E}[F(w_0^{(t)})] - \frac{\eta_t}{2} \mathbb{E}[\|\nabla F(w_0^{(t)})\|^2] + \frac{L^2\eta_t}{2n} \mathbb{E}[\sum_{i=0}^{n-1} \|w_i^{(t)} - w_0^{(t)}\|^2] \\ &\stackrel{(36)}{\leq} \mathbb{E}[F(w_0^{(t)})] - \frac{\eta_t}{2} \mathbb{E}[\|\nabla F(w_0^{(t)})\|^2] + \frac{L^2\eta_t^3}{n} \left[(\Theta + n) \mathbb{E}[\|\nabla F(w_0^{(t)})\|^2] + \sigma^2 \right] \\ &\leq \mathbb{E}[F(w_0^{(t)})] - \frac{\eta_t}{4} \mathbb{E}[\|\nabla F(w_0^{(t)})\|^2] + \frac{L^2\sigma^2\eta_t^3}{n}, \end{aligned}$$

where the last inequality follows since $\eta_t^2 \leq \frac{n}{2(\Theta+n)L^2}$. Note that $\tilde{w}_t = w_0^{(t+1)}$ and $\tilde{w}_{t-1} = w_0^{(t)}$ in Algorithm 1, the last estimate becomes

$$\mathbb{E}[F(\tilde{w}_t) - F_*] \leq \mathbb{E}[F(\tilde{w}_{t-1}) - F_*] - \frac{\eta_t}{4} \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] + \frac{L^2 \sigma^2 \eta_t^3}{n}. \quad (49)$$

Using this estimate and with a similar proof as of (18), we obtain (19). \blacksquare

Proof [The proof of Corollary 1] Given $2\sigma^2 \geq \epsilon$ for a given accuracy $\epsilon > 0$, to guarantee $\frac{1}{T} \sum_{t=1}^T \|\nabla F(\tilde{w}_{t-1})\|^2 \leq \epsilon$, by using (18) in Theorem 3, we impose

$$\frac{4}{T\eta} [F(\tilde{w}_0) - F_*] + 6L^2 \sigma^2 \eta^2 \leq \epsilon.$$

Using $\eta = \frac{\sqrt{\epsilon}}{2L\sigma\sqrt{3\Theta+2}} \leq \frac{1}{L\sqrt{2(3\Theta+2)}}$ into this inequation, we can easily get

$$\frac{8L\sigma\sqrt{3\Theta+2}}{T\sqrt{\epsilon}} [F(\tilde{w}_0) - F_*] \leq \epsilon \left[\frac{6\Theta+1}{6\Theta+4} \right] \Rightarrow T \geq \frac{16L\sigma(3\Theta+2)^{3/2} [F(\tilde{w}_0) - F_*]}{(6\Theta+1)} \cdot \frac{1}{\epsilon^{3/2}}$$

Rounding this expression we get $T := \left\lceil \frac{16L\sigma(3\Theta+2)^{3/2} [F(\tilde{w}_0) - F_*]}{(6\Theta+1)} \cdot \frac{1}{\epsilon^{3/2}} \right\rceil$. As a result, the total number of gradient evaluations is $\mathcal{T}_{\nabla f} := nT = \left\lceil \frac{16L\sigma(3\Theta+2)^{3/2} [F(\tilde{w}_0) - F_*]}{(6\Theta+1)} \cdot \frac{n}{\epsilon^{3/2}} \right\rceil$.

Alternatively, let us choose $\eta := \frac{\sqrt{n\epsilon}}{2L\sigma\sqrt{2(\Theta/n+1)}}$, where $0 < \epsilon \leq \frac{4\sigma^2}{n}$. Then, we have $0 < \eta \leq \frac{1}{L\sqrt{2(\Theta/n+1)}}$. Similar to the above proof, but using (19), we have

$$\frac{4}{T\eta} [F(\tilde{w}_0) - F_*] + \frac{4L^2 \sigma^2 \eta^2}{n} = \frac{8L\sigma\sqrt{2(\Theta/n+1)}}{T\sqrt{n\epsilon}} [F(\tilde{w}_0) - F_*] + \frac{\epsilon}{2(\Theta/n+1)} \leq \epsilon.$$

This condition leads to $T \geq \frac{8L\sigma(2\Theta/n+2)^{3/2} [F(\tilde{w}_0) - F_*]}{(2\Theta/n+1)} \cdot \frac{1}{\sqrt{n\epsilon^{3/2}}}$. Hence, the total number of gradient evaluations is $\mathcal{T}_{\nabla f} = nT = \left\lceil \frac{8L\sigma(2\Theta/n+2)^{3/2} [F(\tilde{w}_0) - F_*]}{(2\Theta/n+1)\epsilon^{3/2}} \cdot \frac{\sqrt{n}}{\epsilon^{3/2}} \right\rceil$. \blacksquare

Proof [The proof of Corollary 2] Substituting $\eta = \frac{\gamma}{T^{1/3}} \leq \frac{1}{2L\sqrt{3\Theta+2}} \leq \frac{1}{L}$ into (18) of Theorem 3, we obtain

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] \leq \frac{4}{T\eta} (F(\tilde{w}_0) - F_*) + 6L^2 \sigma^2 \eta^2 = \frac{1}{T^{2/3}} \left[\frac{4(F(\tilde{w}_0) - F_*)}{\gamma} + 6L^2 \sigma^2 \gamma^2 \right],$$

which is exactly our desired estimate (20).

If, in addition, $\pi^{(t)}$ is sampled uniformly at random from $[n]$, then by choosing $\eta := \frac{\gamma n^{1/3}}{T^{1/3}}$ such that $\eta_t \leq \frac{1}{L\sqrt{2(\Theta/n+1)}} \leq \frac{1}{L}$, we obtain from (19) that

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] \leq \frac{4}{T\eta} [F(\tilde{w}_0) - F_*] + \frac{4L^2 \sigma^2 \eta^2}{n} = \frac{1}{n^{1/3} T^{2/3}} \left[\frac{4[F(\tilde{w}_0) - F_*]}{\gamma} + 4L^2 \sigma^2 \gamma^2 \right].$$

This proves (21). Here, we need to choose $T \geq 1$ such that $\frac{\gamma n^{1/3}}{T^{1/3}} \leq \frac{1}{L\sqrt{2(\Theta/n+1)}}$. \blacksquare

C.2 Proof of Theorem 4: Asymptotic convergence with diminishing stepsize

To establish Theorem 4, we will use the following lemma from (Bertsekas, 2015).

Lemma 9 ((Bertsekas, 2015)) *Let $\{Y_t\}_{t \geq 0}$, $\{Z_t\}_{t \geq 0}$, and $\{W_t\}_{t \geq 0}$ be three sequences of random variables. Let $\{\mathcal{F}_t\}_{t \geq 0}$ be a filtration, i.e. a σ -algebras such that $\mathcal{F}_t \subseteq \mathcal{F}_{t+1}$ for all $t \geq 0$. Suppose that the following conditions hold:*

- (a) Y_t , Z_t , and W_t are nonnegative and \mathcal{F}_t -measurable for all $t \geq 0$;
- (b) for each $t \geq 0$, we have $\mathbb{E}[Y_{t+1} | \mathcal{F}_t] \leq Y_t - Z_t + W_t$;
- (c) with probability 1 (w.p.1), it holds that $\sum_{t=0}^{\infty} W_t < +\infty$.

Then, w.p.1, we have

$$\sum_{t=0}^{\infty} Z_t < +\infty \quad \text{and} \quad Y_t \rightarrow Y \geq 0 \text{ as } t \rightarrow +\infty.$$

Using Lemma 9 we can now prove Theorem 4 in the main text as follows.

Proof [The proof of Theorem 4] First, following the same argument as in the proof of (48) of Theorem 3, we have

$$F(\tilde{w}_{t+1}) \leq F(\tilde{w}_t) - \frac{\eta_{t+1}}{4} \|\nabla F(\tilde{w}_t)\|^2 + \frac{3L^2\sigma^2\eta_{t+1}^3}{2}.$$

Let us define $\mathcal{F}_t = \sigma(\tilde{w}_0, \dots, \tilde{w}_t)$ the σ -algebra generated by $\{\tilde{w}_0, \dots, \tilde{w}_t\}$. Then, for $t \geq 0$, the last inequality implies

$$\mathbb{E}[F(\tilde{w}_{t+1}) - F_* | \mathcal{F}_t] \leq [F(\tilde{w}_t) - F_*] - \frac{\eta_{t+1}}{4} \|\nabla F(\tilde{w}_t)\|^2 + \frac{3L^2\sigma^2\eta_{t+1}^3}{2}.$$

Let us define $Y_t := [F(\tilde{w}_t) - F_*] \geq 0$, $Z_t := \frac{\eta_{t+1}}{4} \|\nabla F(\tilde{w}_t)\|^2 \geq 0$ and $W_t := \frac{3}{2}L^2\sigma^2\eta_{t+1}^3$. Then, the first condition (a) of Lemma 9 holds. Moreover, the last inequality shows that $\mathbb{E}[Y_{t+1} | \mathcal{F}_t] \leq Y_t - Z_t + W_t$, which means that the condition (b) of Lemma 9 holds. Since $\sum_{t=1}^{\infty} \eta_t^3 < +\infty$, we have $\sum_{t=0}^{\infty} W_t < +\infty$, which fulfills the condition (c) of Lemma 9. Then, by applying Lemma 9, we obtain w.p.1 that

$$F(\tilde{w}_t) - F_* \rightarrow Y \geq 0 \text{ as } t \rightarrow +\infty, \quad \text{and} \quad \sum_{t=0}^{\infty} \frac{\eta_{t+1}}{4} \|\nabla F(\tilde{w}_t)\|^2 < +\infty.$$

We prove $\liminf_{t \rightarrow \infty} \|\nabla F(\tilde{w}_{t-1})\| = 0$ w.p.1. by contradiction. Indeed, we assume that there exist $\epsilon > 0$ and $t_0 \geq 0$ such that $\|\nabla F(\tilde{w}_t)\|^2 \geq \epsilon$ for all $t \geq t_0$. In this case, since $\sum_{t=0}^{\infty} \eta_t = \infty$, we have

$$\infty > \sum_{t=t_0}^{\infty} \frac{\eta_{t+1}}{4} \|\nabla F(\tilde{w}_t)\|^2 \geq \frac{\epsilon}{4} \sum_{t=t_0}^{\infty} \eta_{t+1} = \infty.$$

This is a contradiction. As a result, w.p.1., we have $\liminf_{k \rightarrow \infty} \|\nabla F(\tilde{w}_k)\|^2 = 0$, or equivalently, it holds that $\liminf_{k \rightarrow \infty} \|\nabla F(\tilde{w}_k)\| = 0$. The proof still holds if we use (49) of Theorem 3. \blacksquare

C.3 Convergence analysis for different learning rates

This appendix provides convergence analysis for general choices of learning rate.

Proof [The proof of Theorem 5] By (48), we have

$$F(\tilde{w}_t) \leq F(\tilde{w}_{t-1}) - \frac{\eta_t}{4} \|\nabla F(\tilde{w}_{t-1})\|^2 + \frac{3L^2\sigma^2\eta_t^3}{2} \leq F(\tilde{w}_{t-1}) + \frac{3L^2\sigma^2\eta_t^3}{2}.$$

Alternatively, if a randomized reshuffling scheme is used, then by using (49), we have

$$\mathbb{E}[F(\tilde{w}_t)] \leq \mathbb{E}[F(\tilde{w}_{t-1})] - \frac{\eta_t}{4} \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] + \frac{L^2\sigma^2\eta_t^3}{n} \leq \mathbb{E}[F(\tilde{w}_{t-1})] + \frac{L^2\sigma^2\eta_t^3}{n}.$$

Combining both cases, we can write them in a single inequality as

$$\mathbb{E}[F(\tilde{w}_t)] \leq \mathbb{E}[F(\tilde{w}_{t-1})] - \frac{\eta_t}{4} \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] + D \cdot \eta_t^3 \leq \mathbb{E}[F(\tilde{w}_{t-1})] + D \cdot \eta_t^3, \quad (50)$$

where $D := \frac{3}{2}L^2\sigma^2$ for the general shuffling strategy and $D := \frac{L^2\sigma^2}{n}$ for the randomized reshuffling strategy.

Since $\eta_t = \frac{\gamma}{(t+\beta)^\alpha}$, summing up the last inequality (50) from $t = 1$ to $t = k \geq 1$, we have

$$\begin{aligned} F(\tilde{w}_k) &\leq F(\tilde{w}_0) + D \cdot \sum_{t=1}^k \eta_t^3 &&= F(\tilde{w}_0) + D \cdot \sum_{t=1}^k \frac{\gamma^3}{(t+\beta)^{3\alpha}} \\ &\stackrel{(27)}{\leq} F(\tilde{w}_0) + D \cdot \gamma^3 \int_{t=0}^k \frac{dt}{(t+\beta)^{3\alpha}} &&= F(\tilde{w}_0) + D \cdot \gamma^3 \left[-\frac{(t+\beta)^{-(3\alpha-1)}}{3\alpha-1} \Big|_{t=0}^k \right] \\ &\leq F(\tilde{w}_0) + \frac{D\gamma^3}{(3\alpha-1)\beta^{3\alpha-1}}. \end{aligned}$$

Here, we use the fact that $\frac{1}{(t+\beta)^{3\alpha}}$ is nonnegative and monotonically decreasing on $[0, +\infty)$ and $\frac{1}{3} < \alpha < 1$. Subtracting F_* from both sides of the last estimate, for $t \geq 1$, we have

$$\mathbb{E}[F(\tilde{w}_t) - F_*] \leq \mathbb{E}[F(\tilde{w}_0) - F_*] + \frac{D\gamma^3}{(3\alpha-1)\beta^{3\alpha-1}}. \quad (51)$$

On the other hand, subtracting F_* from both sides of (50), we have

$$\mathbb{E}[F(\tilde{w}_t) - F_*] \leq \mathbb{E}[F(\tilde{w}_{t-1}) - F_*] - \frac{\eta_t}{4} \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] + D \cdot \eta_t^3. \quad (52)$$

Now, let us define $Y_t := \mathbb{E}[F(\tilde{w}_{t-1}) - F_*] \geq 0$, $Z_t := \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] \geq 0$, for $t \geq 1$, and $\rho := \frac{1}{4}$. The estimate (52) becomes

$$Y_{t+1} \leq Y_t - \rho\eta_t Z_t + D\eta_t^3.$$

Let us define $C := [F(\tilde{w}_0) - F_*] + \frac{D\gamma^3}{(3\alpha-1)\beta^{3\alpha-1}} > 0$. By (51), we have $Y_t \leq C$ (note that $H = 0$ in Lemma 2), $t \geq 1$. Applying Lemma 2 with $q = 3$ and $m = 1$, we conclude that

- If $\alpha = \frac{1}{2}$, we have

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] &\leq \frac{4(1+\beta)^{1/2}[F(\tilde{w}_0) - F_*]}{\gamma} \cdot \frac{1}{T} + \frac{4C}{\gamma} \cdot \frac{(T-1+\beta)^{1/2}}{T} \\ &\quad + 4D\gamma^2 \cdot \frac{\log(T+\beta) - \log(\beta)}{T}. \end{aligned}$$

- If $\alpha \neq \frac{1}{2}$, we have

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] &\leq \frac{4(1+\beta)^\alpha [F(\tilde{w}_0) - F_*]}{\gamma} \cdot \frac{1}{T} + \frac{2C}{\alpha\gamma} \cdot \frac{(T-1+\beta)^\alpha}{T} \\ &\quad + \frac{4D\gamma^2}{(1-2\alpha)} \cdot \frac{(T+\beta)^{1-2\alpha}}{T}. \end{aligned}$$

For the general shuffling strategy, since $D := \frac{3}{2}L^2\sigma^2$, we need to impose $\eta_t = \frac{\gamma}{(t+\beta)^\alpha} \leq \frac{1}{L\sqrt{2(3\Theta+2)}}$ for all $t \geq 1$. This condition holds if $\gamma L\sqrt{2(3\Theta+2)} \leq (\beta+1)^\alpha$. For the randomized reshuffling strategy, since $D := \frac{L^2\sigma^2}{n}$, we need to impose $\eta_t = \frac{\gamma}{(t+\beta)^\alpha} \leq \frac{1}{L\sqrt{2(\Theta/n+1)}}$

for all $t \geq 1$. This condition holds if $\gamma L\sqrt{2(\Theta/n+1)} \leq (\beta+1)^\alpha$.

Finally, if we replace γ by $n^{1/3}\gamma$, then the last condition becomes $\gamma n^{1/3}L\sqrt{2(\Theta/n+1)} \leq (\beta+1)^\alpha$. Substituting this new quantity $n^{1/3}\gamma$ into γ of the right-hand side of our bounds, we obtain the remaining conclusions of this theorem. \blacksquare

Proof [The proof of Theorem 6] By (50), we have

$$\mathbb{E}[F(\tilde{w}_t)] \leq \mathbb{E}[F(\tilde{w}_{t-1})] - \frac{\eta_t}{4} \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] + D \cdot \eta_t^3 \leq \mathbb{E}[F(\tilde{w}_{t-1})] + D \cdot \eta_t^3,$$

where $D := \frac{3}{2}L^2\sigma^2$ for the general shuffling strategy, and $D := \frac{L^2\sigma^2}{n}$ for the randomized reshuffling strategy. Since $\eta_t = \frac{\gamma}{(t+\beta)^{1/3}}$, summing up this inequality from $t = 1$ to $t = k \geq 1$, we obtain

$$\begin{aligned} \mathbb{E}[F(\tilde{w}_k)] &\leq F(\tilde{w}_0) + D \cdot \sum_{t=1}^k \eta_t^3 &&= F(\tilde{w}_0) + D \cdot \sum_{t=1}^k \frac{\gamma^3}{(t+\beta)} \\ &\stackrel{(27)}{\leq} F(\tilde{w}_0) + \frac{D\gamma^3}{(1+\beta)} + D\gamma^3 \int_{t=1}^k \frac{dt}{(t+\beta)} &&\leq F(\tilde{w}_0) + \frac{D\gamma^3}{(1+\beta)} + D\gamma^3 \log(k+\beta). \end{aligned}$$

Here, we use the fact that $\frac{1}{t+\beta}$ is nonnegative and monotonically decreasing on $[0, +\infty)$. Subtracting F_* from both sides of the last estimate, for $t \geq 1$, we have

$$\mathbb{E}[F(\tilde{w}_t) - F_*] \leq [F(\tilde{w}_0) - F_*] + \frac{D\gamma^3}{(1+\beta)} + D\gamma^3 \log(t+\beta). \quad (53)$$

Define $Y_t := \mathbb{E}[F(\tilde{w}_{t-1}) - F_*] \geq 0$ and $Z_t := \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] \geq 0$ for $t \geq 1$, and $\rho := \frac{1}{4}$. Then, the estimate (53) becomes

$$Y_{t+1} \leq Y_t - \rho\eta_t Z_t + D\eta_t^3.$$

Let us define $C := [F(\tilde{w}_0) - F_*] + \frac{D\gamma^3}{(1+\beta)} > 0$, $H := D\gamma^3 > 0$, and $\theta := 1 + \beta > 0$. Clearly, we have $1 + \theta - \beta = 2 > \frac{2}{3}e^{1/2}$. By (53), we have $Y_t \leq C + H \log(t+\theta)$ for $t \geq 1$. Applying Lemma 2 with $q = 3$, $m = 1$, and $\alpha = \frac{1}{3}$, we conclude that

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] &\leq \frac{4(1+\beta)^{1/3} [F(\tilde{w}_0) - F_*]}{\gamma} \cdot \frac{1}{T} + \frac{6C}{\gamma} \cdot \frac{(T-1+\beta)^{1/3}}{T} \\ &\quad + 6D\gamma^2 \cdot \frac{(T-1+\beta)^{1/3} \log(T+\beta)}{T} + 12D\gamma^2 \cdot \frac{(T+\beta)^{1/3}}{T}, \end{aligned}$$

which is our main bound. The remaining conclusion is proved similarly as in Theorem 5. \blacksquare

C.4 Proofs of Theorem 7: The gradient dominance case

Proof [The proof Theorem 7] Using (50) from the proof of Theorem 5 and Assumption 4, we can derive that

$$\begin{aligned} \mathbb{E}[F(\tilde{w}_t) - F_*] &\leq \mathbb{E}[F(\tilde{w}_{t-1}) - F_*] - \frac{\eta_t}{4} \mathbb{E}[\|\nabla F(\tilde{w}_{t-1})\|^2] + D \cdot \eta_t^3 \\ &\stackrel{(22)}{\leq} \left(1 - \frac{\eta_t}{4\tau}\right) \mathbb{E}[F(\tilde{w}_{t-1}) - F_*] + D \cdot \eta_t^3. \end{aligned}$$

where $D := \frac{3}{2}L^2\sigma^2$ for the general shuffling strategy, and $D := \frac{L^2\sigma^2}{n}$ for the randomized reshuffling strategy. Let $Y_t := \mathbb{E}[F(\tilde{w}_{t-1}) - F_*] \geq 0$, and $\rho := \frac{1}{4\tau}$. We now verify that these quantities satisfy the conditions of Lemma 1 with $q = 2$, i.e. $Y_{t+1} \leq (1 - \rho\eta_t)Y_{t-1} + D \cdot \eta_t^3$. Applying Lemma 1 with $\eta_t := \frac{2}{t+\beta}$ for some $\beta \geq 1$, then we obtain

$$\mathbb{E}[F(\tilde{w}_t) - F_*] = Y_{t+1} \leq \frac{1}{(t+\beta-1)(t+\beta)} \left[\beta(\beta-1)[F(\tilde{w}_0) - F_*] + \frac{8D \log(t+\beta)}{\rho^3} \right].$$

For the general shuffling scheme, we need to choose $\beta \geq 1$ such that $\eta_t = \frac{2}{t+\beta} \leq \frac{1}{L\sqrt{2(3\Theta+2)}}$ for all $t \geq 1$. Hence, we can choose $\beta \geq 2L\sqrt{2(3\Theta+2)} - 1$, and in this case, we obtain (23).

For the randomized reshuffling strategy, if we choose $\eta_t := \frac{2}{t+1+1/n}$, then to guarantee $\eta_t \leq \frac{1}{L\sqrt{2(\Theta/n+1)}}$ for all $t \geq 1$, we require $L\sqrt{2(\Theta/n+1)} \leq 1$. In this case, by substituting $\beta := 1 + 1/n$ and $D := \frac{L^2\sigma^2}{n}$ into the last estimate, we obtain (24). \blacksquare

Appendix D. New Convergence Analysis for Standard SGD

As a side result of our stand-alone technical lemma, Lemma 2, we show in this appendix that we can apply the analysis framework of Lemma 2 to obtain convergence rate results for the standard stochastic gradient algorithm, abbreviated by SGD.

To keep it more general, we consider the stochastic optimization problem with respect to some distribution \mathcal{D} as in (1), i.e.:

$$\min_{w \in \mathbb{R}^d} \left\{ F(w) := \mathbb{E}_{\xi \sim \mathcal{D}} [f(w; \xi)] \right\}, \quad (54)$$

where ∇f is an unbiased gradient estimator of the gradient ∇F of F , i.e.:

$$\mathbb{E}_{\xi \sim \mathcal{D}} [\nabla f(w; \xi)] = \nabla F(w), \quad \forall w \in \text{dom}(F).$$

The standard SGD method without mini-batch for solving (54) can be described as in Algorithm 2.

To analyze convergence rate of Algorithm 2, we assume that problem (54) satisfies Assumptions 1(i) and Assumptions 5 and 6 below.

Assumption 5 (One-side L -smoothness) *The objective function F of (54) satisfies*

$$F(w) \leq F(w') + \langle \nabla F(w'), w - w' \rangle + \frac{L}{2} \|w - w'\|^2, \quad \forall w, w' \in \text{dom}(F). \quad (55)$$

Algorithm 2 Stochastic Gradient Descent (SGD) Method (without mini-batch)

- 1: **Initialize:** Choose an initial point $w_1 \in \text{dom}(F)$.
 - 2: **for** $t = 1, 2, \dots$ **do**
 - 3: Generate a realization of ξ_t and evaluate a stochastic gradient $\nabla f(w_t; \xi_t)$;
 - 4: Choose a step size (i.e. a learning rate) $\eta_t > 0$ (specified later);
 - 5: Update $w_{t+1} := w_t - \eta_t \nabla f(w_t; \xi_t)$;
 - 6: **end for**
-

Assumption 6 (Bounded variance) For (54), there exists $\sigma \in (0, +\infty)$ such that

$$\mathbb{E}[\|\nabla f(w; \xi) - \nabla F(w)\|^2] \leq \sigma^2, \quad \forall w \in \text{dom}(F). \quad (56)$$

We notice that we are able to derive the analysis based on the general bounded variance assumption $\mathbb{E}[\|\nabla f(w; \xi) - \nabla F(w)\|^2] \leq \Theta \|\nabla F(w)\|^2 + \sigma^2$ for some $\Theta \geq 0$ and $\sigma > 0$ as in Assumption 2. For simplicity, we only consider the special case where $\Theta = 0$ since Assumption 6 is commonly used in the literature for the standard SGD method. We prove our first result for Algorithm 2 to solve (54) in the following theorem.

Theorem 8 Assume that Assumptions 1(i), 5, and 6 hold for (54). Let $\{w_t\}$ be generated by Algorithm 2 with $0 < \eta_t := \frac{\gamma}{(t+\beta)^\alpha} \leq \frac{1}{L}$ for some $\gamma > 0$, $\beta > 0$, and $\frac{1}{2} < \alpha < 1$, and $C := [F(w_1) - F_*] + \frac{L\sigma^2\gamma^2}{2(2\alpha-1)\beta^{2\alpha-1}} > 0$. Then, the following bound holds:

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla F(w_t)\|^2] &\leq \frac{2(1+\beta)^\alpha [F(w_1) - F_*]}{\gamma} \cdot \frac{1}{T} + \frac{C}{\alpha\gamma} \cdot \frac{(T-1+\beta)^\alpha}{T} \\ &\quad + \frac{L\sigma^2\gamma}{(1-\alpha)} \cdot \frac{(T+\beta)^{1-\alpha}}{T}. \end{aligned} \quad (57)$$

Proof Let $\mathcal{F}_t = \sigma(w_1, \dots, w_t)$ be the σ -algebra generated by $\{w_1, \dots, w_t\}$. Then, from the one-side L -smoothness of F in (55), we have

$$\begin{aligned} \mathbb{E}[F(w_{t+1}) \mid \mathcal{F}_t] &\leq F(w_t) - \eta_t \|\nabla F(w_t)\|^2 + \frac{\eta_t^2 L}{2} \mathbb{E}[\|\nabla f(w_t; \xi_t)\|^2 \mid \mathcal{F}_t] \\ &= F(w_t) - \eta_t \left(1 - \frac{\eta_t L}{2}\right) \|\nabla F(w_t)\|^2 + \frac{\eta_t^2 L}{2} \mathbb{E}[\|\nabla f(w_t; \xi_t) - \nabla F(w_t)\|^2 \mid \mathcal{F}_t] \\ &\stackrel{(56)}{\leq} F(w_t) - \frac{\eta_t}{2} \|\nabla F(w_t)\|^2 + \frac{\eta_t^2 L}{2} \sigma^2, \end{aligned}$$

where the first equality follows since $\mathbb{E}[\|\nabla f(w_t; \xi_t) - \nabla F(w_t)\|^2 \mid \mathcal{F}_t] = \mathbb{E}[\|\nabla f(w_t; \xi_t)\|^2 \mid \mathcal{F}_t] - \|\nabla F(w_t)\|^2$; and the last inequality follows since F has bounded variance in Assumption 6. Note that $\eta_t \left(1 - \frac{\eta_t L}{2}\right) \geq \frac{\eta_t}{2}$ since $0 < \eta_t \leq \frac{1}{L}$. Subtracting F_* from, and then taking full expectation of both sides of the last estimate, we obtain

$$\begin{aligned} \mathbb{E}[F(w_{t+1}) - F_*] &\leq \mathbb{E}[F(w_t) - F_*] - \frac{\eta_t}{2} \mathbb{E}[\|\nabla F(w_t)\|^2] + \frac{\eta_t^2 L \sigma^2}{2} \\ &\leq \mathbb{E}[F(w_t) - F_*] + \frac{\eta_t^2 L \sigma^2}{2}. \end{aligned} \quad (58)$$

Summing up (58) from $t := 1$ to $t := k \geq 1$, we get

$$\begin{aligned} \mathbb{E}[F(w_{k+1}) - F_*] &\leq \mathbb{E}[F(w_1) - F_*] + \frac{L\sigma^2}{2} \sum_{t=1}^k \eta_t^2 = [F(w_1) - F_*] + \frac{L\sigma^2}{2} \sum_{t=1}^k \frac{\gamma^2}{(t+\beta)^{2\alpha}} \\ &\stackrel{(27)}{\leq} [F(w_1) - F_*] + \frac{L\sigma^2\gamma^2}{2} \int_{t=0}^k \frac{dt}{(t+\beta)^{2\alpha}}. \end{aligned} \quad (59)$$

If $\frac{1}{2} < \alpha < 1$, then for $k \geq 1$, we have

$$\mathbb{E}[F(w_{k+1}) - F_*] \leq [F(w_1) - F_*] + \frac{L\sigma^2\gamma^2}{2(2\alpha - 1)\beta^{2\alpha-1}}. \quad (60)$$

Now, let us define $Y_t := \mathbb{E}[F(w_t) - F_*] \geq 0$, $Z_t := \mathbb{E}[\|\nabla F(w_t)\|^2] \geq 0$ for $t \geq 1$, $\rho := \frac{1}{2}$, and $D := \frac{L\sigma^2}{2}$. Then, the estimate (58) becomes

$$Y_{t+1} \leq Y_t - \rho\eta_t Z_t + D\eta_t^2.$$

Let us define $C := [F(w_1) - F_*] + \frac{L\sigma^2\gamma^2}{2(2\alpha-1)\beta^{2\alpha-1}} > 0$. By (60), we have $Y_t \leq C$ (note that $H = 0$ in Lemma 2), $t \geq 1$. Applying Lemma 2 with $q = 2$, $m = 1$, and $\frac{1}{2} < \alpha < 1$, we conclude that

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla F(w_t)\|^2] \leq \frac{2(1+\beta)^\alpha [F(w_1) - F_*]}{\gamma} \cdot \frac{1}{T} + \frac{C}{\alpha\gamma} \cdot \frac{(T-1+\beta)^\alpha}{T} + \frac{L\sigma^2\gamma}{(1-\alpha)} \cdot \frac{(T+\beta)^{1-\alpha}}{T},$$

which proves (57). \blacksquare

Remark 5 In Theorem 8, if we choose $\alpha := \frac{1}{2} + \delta$ for some $0 < \delta < \frac{1}{2}$, then we have

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla F(w_t)\|^2] &\leq \frac{2(1+\beta)^{\frac{1}{2}+\delta} [F(w_1) - F_*]}{\gamma} \cdot \frac{1}{T} + \frac{C}{\gamma(\frac{1}{2}+\delta)} \cdot \frac{(T-1+\beta)^{\frac{1}{2}+\delta}}{T} + \frac{L\sigma^2\gamma}{(\frac{1}{2}-\delta)} \cdot \frac{(T+\beta)^{\frac{1}{2}-\delta}}{T} \\ &= \mathcal{O}\left(\frac{1}{T^{\frac{1}{2}-\delta}}\right), \end{aligned}$$

where $C := [F(w_1) - F_*] + \frac{L\sigma^2\gamma^2}{4\delta\beta^{4\delta}} > 0$. This rate converges to $\mathcal{O}\left(\frac{1}{\sqrt{T}}\right)$ as $\delta \downarrow 0$.

Finally, if we use $\alpha := 1/2$ and diminishing step-size, then we have the following result.

Theorem 9 Assume that Assumptions 1(i), 5, and 6 hold for (54). Let $\{w_t\}$ be generated by Algorithm 2 with the step-size $0 < \eta_t := \frac{\gamma}{(t+\beta)^{1/2}} \leq \frac{1}{L}$ for some $\gamma > 0$ and $\beta > 0$, and $C := [F(w_1) - F_*] + \frac{L\sigma^2\gamma^2}{2(1+\beta)} > 0$. Then, the following bound holds:

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla F(w_t)\|^2] &\leq \frac{2(1+\beta)^{1/2} [F(w_1) - F_*]}{\gamma} \cdot \frac{1}{T} + \frac{2C}{\gamma} \cdot \frac{(T-1+\beta)^{1/2}}{T} \\ &\quad + L\gamma\sigma^2 \cdot \frac{(T-1+\beta)^{1/2} \log(T+1+\beta)}{T} + 2L\gamma\sigma^2 \cdot \frac{(T+\beta)^{1/2}}{T} \\ &= \mathcal{O}\left(\frac{\log(T)}{T^{1/2}}\right). \end{aligned} \quad (61)$$

Proof If $\alpha := \frac{1}{2}$, then, by (59), we can easily show that

$$\begin{aligned} \mathbb{E}[F(w_{k+1}) - F_*] &\leq \mathbb{E}[F(w_1) - F_*] + \frac{L\sigma^2}{2} \sum_{t=1}^k \eta_t^2 = [F(w_1) - F_*] + \frac{L\sigma^2}{2} \sum_{t=1}^k \frac{\gamma^2}{(t+\beta)} \\ &\stackrel{(27)}{\leq} [F(w_1) - F_*] + \frac{L\sigma^2\gamma^2}{2(1+\beta)} + \frac{L\sigma^2\gamma^2}{2} \int_{t=1}^k \frac{dt}{(t+\beta)^{2\alpha}}. \end{aligned}$$

Hence, for $k \geq 1$, we have

$$\mathbb{E}[F(w_{k+1}) - F_*] \leq [F(w_1) - F_*] + \frac{L\sigma^2\gamma^2}{2(1+\beta)} + \frac{L\sigma^2\gamma^2}{2} \cdot \log(k+2+\beta). \quad (62)$$

Define $Y_t := \mathbb{E}[F(w_t) - F_*] \geq 0$, $Z_t := \mathbb{E}[\|\nabla F(w_t)\|^2] \geq 0$ for $t \geq 1$, $\rho := \frac{1}{2}$, and $D := \frac{L\sigma^2}{2}$. Then, the estimate (58) becomes

$$Y_{t+1} \leq Y_t - \rho\eta_t Z_t + D\eta_t^2.$$

Let us define $C := [F(w_1) - F_*] + \frac{L\sigma^2\gamma^2}{2(1+\beta)} > 0$, $H := \frac{L\sigma^2\gamma^2}{2} > 0$, and $\theta := 1 + \beta > 0$. Clearly, $1 + \theta - \beta = 2 > \frac{1}{2}e$. By (62), we have $Y_t \leq C + H \log(t + \theta)$ for $t \geq 1$. Applying Lemma 2 with $q = 2$, $m = 1$, and $\alpha = \frac{1}{2}$, we conclude that

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbb{E}[\|\nabla F(w_t)\|^2] &\leq \frac{2(1+\beta)^{1/2}[F(w_1)-F_*]}{\gamma} \cdot \frac{1}{T} + \frac{2C}{\gamma} \cdot \frac{(T-1+\beta)^{1/2}}{T} \\ &\quad + L\gamma\sigma^2 \cdot \frac{(T-1+\beta)^{1/2} \log(T+1+\beta)}{T} + 2L\gamma\sigma^2 \cdot \frac{(T+\beta)^{1/2}}{T} \\ &= \mathcal{O}\left(\frac{\log(T)}{T^{1/2}}\right), \end{aligned}$$

which proves (61). ■

References

- K. Ahn, C. Yun, and S. Sra. SGD with Shuffling: Optimal Rates without Component Convexity and Large Epoch Requirements. In *Advances in Neural Information Processing Systems*, volume 33, pages 17526–17535, 2020.
- D. P. Bertsekas. Incremental Gradient, Subgradient, and Proximal Methods for Convex Optimization: A Survey, 2015.
- L. Bottou. Curiously Fast Convergence of Some Stochastic Gradient Descent Algorithms. In *Proceedings of the Symposium on Learning and Data Science, Paris*, volume 8, pages 2624–2633, 2009.
- L. Bottou. Stochastic Gradient Descent Tricks. In *Neural networks: Tricks of the trade*, pages 421–436. Springer, 2012.
- L. Bottou, F. E. Curtis, and J. Nocedal. Optimization Methods for Large-Scale Machine Learning. *SIAM Review*, 60(2):223–311, 2018.

- C.-C. Chang and C.-J. Lin. LIBSVM: A Library for Support Vector Machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27, 2011. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- A. Defazio, F. Bach, and S. Lacoste-Julien. SAGA: A Fast Incremental Gradient Method with Support for Non-Strongly Convex Composite Objectives. In *Advances in Neural Information Processing Systems*, pages 1646–1654, 2014a.
- A. Defazio, T. Caetano, and J. Domke. Finito: A Faster, Permutable Incremental Gradient Method for Big Data Problems. In *International Conference on Machine Learning*, pages 1125–1133, 2014c.
- J. Duchi, E. Hazan, and Y. Singer. Adaptive Subgradient Methods for Online Learning and Stochastic Optimization. *Journal of Machine Learning Research*, 12:2121–2159, 2011.
- S. Ghadimi and G. Lan. Stochastic First- and Zeroth-Order Methods for Nonconvex Stochastic Programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- M. Gürbüzbalaban, A. Ozdaglar, and P. Parrilo. Convergence Rate of Incremental Gradient and Newton Methods. *SIAM Journal on Optimization*, volume 29, number 4, pages 2542–2565, 2019.
- M. Gürbüzbalaban, A. Ozdaglar, and P. Parrilo. Why Random Reshuffling Beats Stochastic Gradient Descent? *Mathematical Programming*, volume 186, number 1, pages 49–84, 2021.
- J. Z. HaoChen and S. Sra. Random Shuffling Beats SGD After Finite Epochs. In *International Conference on Machine Learning*, pages 2624–2633. PMLR, 2019.
- K. Hiroyuki. SGDLibrary: A MATLAB Library for Stochastic Optimization Algorithms. *Journal of Machine Learning Research*, 18(215):1–5, 2018.
- R. Johnson and T. Zhang. Accelerating Stochastic Gradient Descent Using Predictive Variance Reduction. In *Advances in Neural Information Processing Systems*, pages 315–323, 2013.
- H. Karimi, J. Nutini, and M. Schmidt. Linear Convergence of Gradient and Proximal-Gradient Methods under the Polyak-Lojasiewicz Condition. In P. Frasconi, N. Landwehr, G. Manco, and J. Vreeken, editors, *Machine Learning and Knowledge Discovery in Databases*, pages 795–811, Cham, 2016. Springer International Publishing.
- D. P. Kingma and J. Ba. ADAM: A Method for Stochastic Optimization. In *International Conference on Learning Representations*, 2015.
- A. Krizhevsky and G. Hinton. Learning Multiple Layers of Features from Tiny Images. Technical report, Citeseer, 2009.
- Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-Based Learning Applied to Document Recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

- X. Li, Z. Zhu, A. So, and J. D. Lee. Incremental Methods for Weakly Convex Optimization. In *OPT2020: 12th Annual Workshop on Optimization for Machine Learning*, 2020.
- Q. Meng, W. Chen, Y. Wang, Z.-M. Ma, and T.-Y. Liu. Convergence Analysis of Distributed Stochastic Gradient Descent with Shuffling. *Neurocomputing*, 337:46–57, 2019.
- K. Mishchenko, A. Khaled Ragab Bayoumi, and P. Richtárik. Random Reshuffling: Simple Analysis with Vast Improvements. In *Advances in Neural Information Processing Systems*, 33, 2020.
- D. Nagaraj, P. Jain, and P. Netrapalli. SGD without Replacement: Sharper Rates for General Smooth Convex Functions. In *International Conference on Machine Learning*, pages 4703–4711, 2019.
- A. Nedić and D. Bertsekas. Convergence Rate of Incremental Subgradient Algorithms. In *Stochastic optimization: Algorithms and applications*, pages 223–264. Springer, 2001.
- A. Nedic and D. P. Bertsekas. Incremental Subgradient Methods for Nondifferentiable Optimization. *SIAM Journal on Optimization*, 12(1):109–138, 2001.
- A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust Stochastic Approximation Approach to Stochastic Programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.
- Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*, volume 87 of *Applied Optimization*. Kluwer Academic Publishers, 2004.
- Y. Nesterov and B. T. Polyak. Cubic Regularization of Newton Method and Its Global Performance. *Mathematical Programming*, 108(1):177–205, 2006.
- L. M. Nguyen, J. Liu, K. Scheinberg, and M. Takáč. SARAH: A Novel Method for Machine Learning Problems Using Stochastic Recursive Gradient. In *International Conference on Machine Learning*, 2017.
- L. M. Nguyen, H. P. Nguyen, M. van Dijk, P. Richtárik, K. Scheinberg, and M. Takáč. SGD and Hogwild! Convergence Without the Bounded Gradients Assumption. In *International Conference on Machine Learning*, pages 3747–3755, 2018.
- N. H. Pham, L. M. Nguyen, D. T. Phan, and Q. Tran-Dinh. ProxSARAH: An Efficient Algorithmic Framework for Stochastic Composite Nonconvex Optimization. *Journal of Machine Learning Research*, 21:1–48, 2020.
- B. Polyak and A. Juditsky. Acceleration of Stochastic Approximation by Averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855, 1992.
- Boris T. Polyak. Some Methods of Speeding Up the Convergence of Iteration Methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5):1–17, 1964.
- S. Rajput, A. Gupta, and D. Papailiopoulos. Closing the Convergence Gap of SGD Without Replacement. In *International Conference on Machine Learning*, pages 7964–7973. PMLR, 2020.

- S. J. Reddi, A. Hefny, S. Sra, B. Póczos, and A. J. Smola. Stochastic Variance Reduction for Nonconvex Optimization. In *International Conference on Machine Learning*, pages 314–323, 2016.
- H. Robbins and S. Monro. A Stochastic Approximation Method. *The Annals of Mathematical Statistics*, 22(3):400–407, 1951.
- N. Le Roux, M. Schmidt, and F. Bach. A Stochastic Gradient Method with an Exponential Convergence Rate for Finite Training Sets. In *Advances in Neural Information Processing Systems*, pages 2663–2671, 2012.
- I. Safran and O. Shamir. How Good is SGD with Random Shuffling? In *Conference on Learning Theory*, pages 3250–3284, 2020.
- O. Shamir. Without-Replacement Sampling for Stochastic Gradient Methods. In *Advances in Neural Information Processing Systems*, pages 46–54, 2016.
- S. Sra, S. Nowozin, and S. J. Wright. *Optimization for Machine Learning*. MIT Press, 2012.
- Q. Tran-Dinh, N. H. Pham, D. T. Phan, and L. M. Nguyen. A Hybrid Stochastic Optimization Framework for Stochastic Composite Nonconvex Optimization. *Mathematical Programming*, pages 1–67, 2021.
- B. Ying, K. Yuan, and A. H. Sayed. Variance-Reduced Stochastic Learning Under Random Reshuffling. *IEEE Transactions on Signal Processing*, 68:1390–1408, 2020.