
An Improved Analysis of Gradient Tracking for Decentralized Machine Learning

Anastasia Koloskova
EPFL
anastasia.koloskova@epfl.ch

Tao Lin
EPFL
tao.lin@epfl.ch

Sebastian U. Stich
EPFL*
sebastian.stich@epfl.ch

Abstract

We consider decentralized machine learning over a network where the training data is distributed across n agents, each of which can compute stochastic model updates on their local data. The agent’s common goal is to find a model that minimizes the average of all local loss functions. While gradient tracking (GT) algorithms can overcome a key challenge, namely accounting for differences between workers’ local data distributions, the known convergence rates for GT algorithms are not optimal with respect to their dependence on the mixing parameter p (related to the spectral gap of the connectivity matrix).

We provide a tighter analysis of the GT method in the stochastic strongly convex, convex and non-convex settings. We improve the dependency on p from $\mathcal{O}(p^{-2})$ to $\mathcal{O}(p^{-1}c^{-1})$ in the noiseless case and from $\mathcal{O}(p^{-3/2})$ to $\mathcal{O}(p^{-1/2}c^{-1})$ in the general stochastic case, where $c \geq p$ is related to the negative eigenvalues of the connectivity matrix (and is a constant in most practical applications). This improvement was possible due to a new proof technique which could be of independent interest.

1 Introduction

Methods that train machine learning models on decentralized data offer many advantages over traditional centralized approaches in core aspects such as data ownership, privacy, fault tolerance and scalability [12, 33]. Many current efforts in this direction come under the banner of federated learning [17, 29, 28, 12], where a central entity orchestrates the training and collects aggregate updates from the participating devices. Fully decentralized methods, that do not rely on a central coordinator and that communicate only with neighbors in an arbitrary communication topology, are still in their infancy [24, 18].

The work of Lian et al. [24] on decentralized stochastic gradient descent (D-SGD) has spurred the research on decentralized training methods for machine learning models. This led to improved theoretical analyses [16] and to improved practical schemes, such as support for time-varying topologies [32, 3, 16] and methods with communication compression [45, 51, 15, 47]. One of the most challenging aspects when training over decentralized data is data-heterogeneity, i.e. training data that is in a non-IID fashion distributed over the devices (for instance in data-center training) or generated in non-IID fashion on client devices [21, 13, 22, 23]. For example, the D-SGD method has been shown to be affected by the heterogeneity [16].

In contrast, certain methods can mitigate the impact of heterogeneous data in decentralized optimization. For instance the *gradient tracking* (GT) methods developed by Lorenzo and Scutari [26] and Nedić et al. [34], or the later D^2 method by Tang et al. [46] which is designed for communication topologies that remain fixed and do not change over time.

*Current affiliation: CISPA Helmholtz Center for Information Security.

Table 1: Important advances for Gradient Tracking in the strongly convex case. Our analysis improves upon all prior rates for both with and without the stochastic noise in terms of the graph parameter p .

Reference	rate of convergence to ϵ -accuracy	considered stochastic noise
Nedić et al. [34]	$\mathcal{O}\left(\frac{L^3}{\mu^3 p^2} \log \frac{1}{\epsilon}\right)$	✗
Alghunaim et al. [1]	$\mathcal{O}\left(\frac{L}{\mu} \log \frac{1}{\epsilon} + \frac{1}{p^2} \log \frac{1}{\epsilon}\right)$	✗
Qu and Li [40]	$\mathcal{O}\left(\frac{L^2}{\mu^2 p^2} \log \frac{1}{\epsilon}\right)$	✗
Pu and Nedić [39]	$\tilde{\mathcal{O}}\left(\frac{\sigma^2}{\mu n \epsilon} + \frac{\sqrt{L}\sigma}{\mu\sqrt{pp}\sqrt{\epsilon}} + \frac{C_1}{\sqrt{\epsilon}}\right)^a$	✓
this work	$\tilde{\mathcal{O}}\left(\frac{\sigma^2}{\mu n \epsilon} + \frac{\sqrt{L}\sigma}{\mu\sqrt{pc}\sqrt{\epsilon}} + \frac{L}{\mu pc} \log \frac{1}{\epsilon}\right)$	✓

^a C_1 is a constant that is independent of ϵ , but can depend on other parameters, such as σ, μ, L, p

It is well known that GT methods do not depend on the heterogeneity of the data and that they converge linearly on distributed strongly convex problem instances without stochastic noise [26, 34]. However, when we apply these methods in the context of machine learning, we need to understand how they are affected by stochastic noise and how they behave on non-convex tasks.

In this paper, we develop a new, and improved, analysis of the gradient tracking algorithm with a novel proof technique. Along with the parallel contribution [55] that developed a tighter analysis of the D^2 algorithm, we now have a more accurate understanding of in which setting GT works well and in which ones it does not, and our results allow for a more detailed comparison between the D-SGD, GT and D^2 methods (see Section 5 below).

Our analysis improves over all existing results that analyze the GT algorithm. Specifically, we prove a weaker dependence on the connectivity of the network (spectral gap) which is commonly incorporated into the convergence rates via the standard parameter p . For example, in the strongly convex setting with stochastic noise we prove that GT converges at the rate $\tilde{\mathcal{O}}\left(\frac{\sigma^2}{n\epsilon} + \frac{1}{c} \cdot \left(\frac{\sigma}{\sqrt{p\epsilon}} + \frac{1}{p} \log \frac{1}{\epsilon}\right)\right)$ where σ^2 is an upper bound on the variance of the stochastic noise, and $c \geq p$ a new parameter (often a constant). By comparing this result with the previously best known upper bound, $\tilde{\mathcal{O}}\left(\frac{\sigma^2}{n\epsilon} + \frac{1}{p} \cdot \left(\frac{\sigma}{\sqrt{p\epsilon}} + \frac{1}{p} \log \frac{1}{\epsilon}\right)\right)$, by Pu and Nedić [39], we see that our upper bound improves the last two terms by a factor of $\frac{c}{p} \geq 1$ and that the first term matches with known lower bounds [37]. The D^2 algorithm [46] only converges under the assumption that c is a constant² and the recent upper bound from [55] coincides with our worst case complexity for GT on all topologies where D^2 can be applied. We provide additional comparison of GT convergence rates in the Tables 1 and 2.

Contributions. Our main contributions can be summarized as:

- We prove better complexity estimates for the GT algorithm than known before with a new proof technique (which might be of independent interest).
- In the non-asymptotic regime (of importance in practice), the convergence rate depends on the network topology. By defining new graph parameters, we can give a tighter description of this dependency, explaining why the worst case behavior is rarely observed in practice (see Section 5.1). We verify this dependence in numerical experiments.
- We show that in the presence of stochastic noise, the leading term in the convergence rate of GT is optimal—we are the first to derive this in the non-convex setting—and matching the unimprovable rate of all-reduce mini-batch SGD.

2 Related Work

Decentralized Optimization. Decentralized optimization methods have been studied for decades in the optimization and control community [48, 30, 52, 6]. Many decentralized optimization methods

²In D^2 the smallest eigenvalue of the mixing matrix W must be bounded from below: $\min_i \lambda_i(W) \geq -\frac{1}{3}$.

Table 2: Important advances for Gradient Tracking in the non-convex case. Our result improves upon all existing rates in terms of the graph parameter p .

Reference	rate of convergence to ϵ -accuracy	considered stochastic noise
Lorenzo and Scutari [26]	asymptotic convergence guarantees	\times
Zhang and You [60]	$\mathcal{O}\left(\frac{Ln\sigma^2}{\epsilon^2} + \frac{Ln}{p^3\epsilon}\right)$	\checkmark
Lu et al. [27]	$\mathcal{O}\left(\frac{C_1 + C_2\sigma}{\epsilon^2}\right)^a$	\checkmark
this work	$\tilde{\mathcal{O}}\left(\frac{L\sigma^2}{n\epsilon^2} + \frac{L\sigma}{(\sqrt{pc} + p\sqrt{n})\epsilon^{3/2}} + \frac{L}{pc\epsilon}\right)$	\checkmark

^a C_1 and C_2 are constants that are independent of ϵ , but can depend on other parameters, such as σ, n, L, p .

[30, 11] are based on gossip averaging [14, 53, 4]. Such methods usually also work well on non-convex problems and can be used for training deep neural networks [3, 24, 46]. There exists other methods, such as based on alternating direction method of multipliers (ADMM) [52, 10], dual averaging [6, 31, 41], primal-dual methods [2, 19], block-coordinate methods for generalized linear models [8] or using new gradient propagation mechanisms [50].

Decentralized Optimization with Heterogeneous Objective Functions. There exists several algorithms that are agnostic to data-heterogeneity. Notably, EXTRA [42] and decentralized primal-dual gradient methods [2] do not depend on the data heterogeneity and achieve linear convergence in the strongly convex noiseless setting. However, these algorithms are not designed to be used for non-convex tasks.

D^2 [46, 55] (also known as exact diffusion [56, 57]) and Gradient Tracking (GT) [26] (also known as NEXT [26] or DIGing [34]) are both algorithms that are agnostic to the data heterogeneity level, can tolerate the stochastic noise, and that can be applied to non-convex objectives such as the training of deep neural networks in machine learning. A limitation of the D^2 algorithm is that it is not clear how it can be applied to time-varying topologies, and that it can only be used on constant mixing topologies with negative eigenvalue bounded from below by $-\frac{1}{3}$. Other authors proposed algorithms that perform well on heterogeneous DL tasks [25, 59], but theoretical proofs that these algorithms are independent of the degree of heterogeneity are still pending.

Gradient Tracking. There is a vast literature on the Gradient Tracking method itself. A tracking mechanism was used by Zhu and Martínez [61] as a way to track the average of a distributed continuous process. Lorenzo and Scutari [26] applied this technique to track the gradients, and analyzed its asymptotic behavior in the non-convex setting with a time-varying topologies. Nedić et al. [34] analyze GT (named as DIGing) in the strongly convex noiseless case with a time-varying network. Qu and Li [40] extend the GT analysis to the non-convex, weakly-convex and strongly convex case without stochastic noise. Nedić et al. [35] allow the different stepsizes on different workers. Yuan et al. [58] analyze asymptotic behavior of GT for dynamic optimization. Pu and Nedić [39] studied the GT method on stochastic problems and strongly convex objectives. Further, Xin et al. [54] analyze asymptotic behavior of GT with stochastic noise. For non-convex stochastic functions GT was analyzed by Zhang and You [60] and Lu et al. [27]. Li et al. [20] combine GT with variance reduction to achieve linear convergence in the stochastic case. Tziotis et al. [49] obtain second order guarantees for GT.

3 Setup

We consider optimization problems where the objective function is distributed across n nodes,

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

where $f_i: \mathbb{R}^d \rightarrow \mathbb{R}$ denotes the local function available to the node $i, i \in [n] := \{1, \dots, n\}$. Each f_i is a stochastic function $f_i(\mathbf{x}) = \mathbb{E}_{\xi \sim \mathcal{D}_i} F_i(\mathbf{x}, \xi)$ with access only to stochastic gradients $\nabla F_i(\mathbf{x}, \xi)$. This setting covers empirical risk minimization problems with \mathcal{D}_i being a uniform distribution over the local training dataset. It also covers deterministic optimization when $F_i(\mathbf{x}, \xi) = f_i(\mathbf{x}), \forall \xi$.

We consider optimization over a decentralized network, i.e. when there is an underlying communication graph $G = (V, E)$, $|V| = n$, each of the nodes (e.g. a connected device) can communicate only along the edges E . In decentralized optimization it is convenient to parameterize communication by a mixing matrix $W \in \mathbb{R}^{n \times n}$, where $w_{ij} = 0$ if and only if nodes i and j are not communicating, $(i, j) \notin E$.

Definition 1 (Mixing Matrix). *A matrix with non-negative entries $W \in [0, 1]^{n \times n}$ that is symmetric ($W = W^\top$) and doubly stochastic ($W\mathbf{1} = \mathbf{1}$, $\mathbf{1}^\top W = \mathbf{1}^\top$), where $\mathbf{1}$ denotes the all-one vector in \mathbb{R}^n .*

3.1 Notation

We use the notation $\mathbf{x}_i^{(t)} \in \mathbb{R}^d$, $\mathbf{y}_i^{(t)} \in \mathbb{R}^d$ to denote the iterates and the tracking sequence, respectively, on node i at time step t . For vectors $\mathbf{z}_i \in \mathbb{R}^d$ (\mathbf{z}_i could for instance be $\mathbf{x}_i^{(t)}$ or $\mathbf{y}_i^{(t)}$) defined for $i \in [n]$ we denote by $\bar{\mathbf{z}} = \frac{1}{n} \sum_{i=1}^n \mathbf{z}_i$.

We use both vector and matrix notation whenever it is more convenient. For vectors $\mathbf{z}_i \in \mathbb{R}^d$ defined for $i \in [n]$ we denote by a capital letter the matrix with columns \mathbf{z}_i , formally

$$Z := [\mathbf{z}_1, \dots, \mathbf{z}_n] \in \mathbb{R}^{d \times n}, \quad \bar{Z} := [\bar{\mathbf{z}}, \dots, \bar{\mathbf{z}}] \equiv Z \frac{1}{n} \mathbf{1} \mathbf{1}^\top, \quad \Delta Z := Z - \bar{Z}. \quad (2)$$

We extend this definition to gradients of (1), with $\nabla F(X^{(t)}, \xi^{(t)})$, $\nabla f(X^{(t)}) \in \mathbb{R}^{d \times n}$:

$$\begin{aligned} \nabla F(X^{(t)}, \xi^{(t)}) &:= \left[\nabla F_1(\mathbf{x}_1^{(t)}, \xi_1^{(t)}), \dots, \nabla F_n(\mathbf{x}_n^{(t)}, \xi_n^{(t)}) \right], \\ \nabla f(X^{(t)}) &:= \left[\nabla f(\mathbf{x}_1^{(t)}), \dots, \nabla f(\mathbf{x}_n^{(t)}) \right]. \end{aligned}$$

3.2 Algorithm

The Gradient Tracking algorithm (or NEXT, DIGing) can be written as

$$\begin{pmatrix} X^{(t+1)} \\ \gamma Y^{(t+1)} \end{pmatrix}^\top = \begin{pmatrix} X^{(t)} \\ \gamma Y^{(t)} \end{pmatrix}^\top \begin{pmatrix} W & 0 \\ -W & W \end{pmatrix} + \gamma \begin{pmatrix} 0 \\ \nabla F(X^{(t+1)}, \xi^{(t+1)}) - \nabla F(X^{(t)}, \xi^{(t)}) \end{pmatrix}^\top \quad (\text{GT})$$

in matrix notation. Here and $X^{(t)} \in \mathbb{R}^{d \times n}$ denotes the iterates, $Y^{(t)} \in \mathbb{R}^{d \times n}$, with $Y^{(0)} = \nabla F(X^{(0)}, \xi^{(0)})$ the sequence of tracking variables, and $\gamma > 0$ denotes the stepsize. This update is summarized in Algorithm 1.

Algorithm 1 GRADIENT TRACKING

input Initial values $\mathbf{x}_i^{(0)} \in \mathbb{R}^d$ on each node $i \in [n]$, communication graph $G = ([n], E)$ and mixing matrix W , stepsize γ , initialize $\mathbf{y}_i^{(0)} = \nabla F_i(\mathbf{x}_i^{(0)}, \xi_i^{(0)})$, $\mathbf{g}_i^{(0)} = \mathbf{y}_i^{(0)}$ in parallel for $i \in [n]$.

- 1: **in parallel on all workers** $i \in [n]$, **for** $t = 0, \dots, T - 1$ **do**
 - 2: each node i sends $(\mathbf{x}_i^{(t)}, \mathbf{y}_i^{(t)})$ to its neighbors
 - 3: $\mathbf{x}_i^{(t+1)} = \sum_{j: \{i,j\} \in E} w_{ij} (\mathbf{x}_j^{(t)} - \gamma \mathbf{y}_j^{(t)})$ ▷ update model parameters
 - 4: Sample $\xi_i^{(t+1)}$, compute gradient $\mathbf{g}_i^{(t+1)} = \nabla F_i(\mathbf{x}_i^{(t+1)}, \xi_i^{(t+1)})$
 - 5: $\mathbf{y}_i^{(t+1)} = \sum_{j: \{i,j\} \in E} w_{ij} \mathbf{y}_j^{(t)} + (\mathbf{g}_i^{(t+1)} - \mathbf{g}_i^{(t)})$ ▷ update tracking variable
 - 6: **end parallel for**
-

Each node i stores and updates two variables, the model parameter $\mathbf{x}_i^{(t)}$ and the tracking variable $\mathbf{y}_i^{(t)}$. The model parameters are updated on line 3 with a decentralized SGD update but using $\mathbf{y}_i^{(t)}$ instead of a gradient. Variable $\mathbf{y}_i^{(t)}$ tracks the average of all local gradients on line 5. Intuitively, the algorithm is agnostic to the functions heterogeneity because $\mathbf{y}_i^{(t)}$ is ‘close’ to the full gradient of $f(\mathbf{x})$ (suppose we would replace line 5 with exact averaging in every timestep, then $\mathbf{y}_i^{(t+1)} = \frac{1}{n} \sum_{i=1}^n \mathbf{g}_i^{(t+1)}$). For further discussion of the tracking mechanism refer to [26, 34, 39].

graph/topology	$1/p$	c
ring	$\mathcal{O}(n^2)$	$8/9$
2d-torus	$\mathcal{O}(n)$	$\geq 4/5$
fully connected	$\mathcal{O}(1)$	1

Table 3: Parameters p and c for some common network topologies on n nodes for uniformly averaging W , i.e. $w_{ij} = \frac{1}{\deg(i)} = \frac{1}{\deg(j)}$ for $\{i, j\} \in E$, see e.g. [36].

3.3 Assumptions

We first state an assumption on the mixing matrix.

Assumption 1 (Mixing Matrix). *Let $\lambda_i(W)$, $i \in [n]$, denote the eigenvalues of the mixing matrix W with $1 = \lambda_1(W) > \lambda_2(W) \geq \dots \geq \lambda_n(W) > -1$. With this, we can define the spectral gap $\delta = 1 - \max\{|\lambda_2(W)|, |\lambda_n(W)|\}$, and the mixing parameters*

$$p = 1 - \max\{|\lambda_2(W)|, |\lambda_n(W)|\}^2, \quad c = 1 - \min\{\lambda_n(W), 0\}^2. \quad (3)$$

We assume that $p > 0$ (and consequently $c > 0$).

The assumption $p > 0$ ensures that the network topology is connected, and that the consensus distance decreases linearly after each averaging step, i.e. $\|XW - \bar{X}\|_F^2 \leq (1-p) \|X - \bar{X}\|_F^2$, $\forall X \in \mathbb{R}^{d \times n}$. The parameter p is closely related to the spectral gap δ as it holds $p = 2\delta - \delta^2$. From this we can conclude that $\delta \leq p \leq 2\delta$ and, asymptotically for $\delta \rightarrow 0$, $p \rightarrow 2\delta$. Assuming a lower bound on p (or equivalently δ) is a standard assumption in the literature.

The parameter c is related to the most negative eigenvalue. From the definition (3) it follows that the auxiliary mixing parameter $c \geq p$ for all mixing matrices W . The parameters p and c are only equal when $|\lambda_n(W)| \geq |\lambda_2(W)|$ and $\lambda_n(W) \leq 0$. Moreover, if the diagonal entries w_{ii} (self-weights) of the mixing matrix are all strictly positive, then c has to be strictly positive.

Remark 1 (Lower bound on c). *Let W be a mixing matrix with diagonal entries (self-weights) $w_{ii} \geq \rho > 0$, for a parameter ρ . Then $\lambda_n(W) \geq 2\rho - 1$ and $c \geq \min\{2\rho, 1\}$.*

This follows from Gershgorin's circle theorem [7] that guarantees $\lambda_n(W) \geq 2\rho - 1$, and hence $c \geq 1 - \min\{2\rho - 1, 0\}^2 \geq \min\{2\rho, 1\}$.

For many choices of W considered in practice, most notably when the graph G has constant node-degree and the weights w_{ij} are chosen by the popular Metropolis-Hastings rule, i.e. $w_{ij} = w_{ji} = \min\{\frac{1}{\deg(i)+1}, \frac{1}{\deg(j)+1}\}$ for $(i, j) \in E$, $w_{ii} = 1 - \sum_{j=1}^n w_{ij} \geq \frac{1}{\max_{j \in [n]} \deg(j)}$, see also [53, 4]. In this case, the parameter c can be bounded by a constant depending on the maximal degree. Moreover, for any given W , considering $\frac{1}{2}(W + I_n)$ instead (i.e. increasing the self-weights), ensures that $c = 1$. However, in contrast to e.g. the analysis in [55] we do not need to pose an explicit bound on c as an assumption. In practice, for many graphs, the parameter c is bounded by a constant (see Table 3).

We further use the following standard assumptions:

Assumption 2 (L -smoothness). *Each function $f_i: \mathbb{R}^d \rightarrow \mathbb{R}$, $i \in [n]$ is differentiable and there exists a constant $L \geq 0$ such that for each $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$:*

$$\|\nabla f_i(\mathbf{y}) - \nabla f_i(\mathbf{x})\| \leq L \|\mathbf{x} - \mathbf{y}\|. \quad (4)$$

Sometimes we will in addition assume that the functions are (strongly) convex.

Assumption 3 (μ -strong convexity). *Each function $f_i: \mathbb{R}^d \rightarrow \mathbb{R}$, $i \in [n]$ is μ -strongly convex for constant $\mu \geq 0$, i.e. for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$:*

$$f_i(\mathbf{x}) - f_i(\mathbf{y}) + \frac{\mu}{2} \|\mathbf{x} - \mathbf{y}\|_2^2 \leq \langle \nabla f_i(\mathbf{x}), \mathbf{x} - \mathbf{y} \rangle. \quad (5)$$

Assumption 4 (Bounded noise). *We assume that there exists constant σ s.t. $\forall \mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^d$*

$$\frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\xi_i} \|\nabla F_i(\mathbf{x}_i, \xi_i) - \nabla f_i(\mathbf{x}_i)\|_2^2 \leq \sigma^2. \quad (6)$$

We discuss possible relaxations of these assumptions in Section 4.1 below.

4 Convergence results

We now present our novel convergence results for GT in Section 4.1 and Section 4.2 below. We provide a proof sketch to explain the key difficulties and technical novelty compared to prior results later in the next Section 6.

4.1 Main theorem—GT convergence in the general case

Theorem 2. Let $\mathbf{x}_i^{(t)}$, $i \in [n]$, $T > \frac{2}{p} \log \left(\frac{50}{p} (1 + \log \frac{1}{p}) \right)$ denote the iterates of the GT Algorithm 1 with a mixing matrix as in Definition 1. If Assumptions 1, 2 and 4 hold, then there exists a stepsize γ such that the optimization error is bounded as follows:

Non-convex: Let $F_0 = f(\bar{\mathbf{x}}^{(0)}) - f^*$ for $f^* \leq \min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x})$. Then it holds

$$\frac{1}{T+1} \sum_{t=0}^T \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2^2 \leq \varepsilon, \text{ after } \tilde{O} \left(\frac{\sigma^2}{n\varepsilon} + \frac{\sigma}{(\sqrt{pc} + p\sqrt{n})\varepsilon^{3/2}} + \frac{1 + L\tilde{R}_0^2 F_0^{-1}}{pc\varepsilon} \right) \cdot LF_0 \text{ iterations.}$$

Strongly-convex: Under the additional Assumption 3 with $\mu > 0$ and weights $w_t \geq 0$, $W_T = \sum_{t=0}^T w_t$, specified in the proof, it holds for $R_{T+1}^2 = \|\bar{\mathbf{x}}^{(T+1)} - \mathbf{x}^*\|^2$:

$$\sum_{t=0}^T \frac{w_t}{W_T} [\mathbb{E} f(\bar{\mathbf{x}}^{(t)}) - f^*] + \frac{\mu}{2} R_{T+1} \leq \varepsilon, \text{ after } \tilde{O} \left(\frac{\sigma^2}{\mu n \varepsilon} + \frac{\sqrt{L}\sigma}{\mu\sqrt{pc}\sqrt{\varepsilon}} + \frac{L}{\mu pc} \log \frac{1}{\varepsilon} \right) \text{ iterations.}$$

General convex: Under the additional Assumption 3 with $\mu \geq 0$, it holds for $R_0^2 = \|\bar{\mathbf{x}}^{(0)} - \mathbf{x}^*\|^2$:

$$\frac{1}{T+1} \sum_{t=0}^T [\mathbb{E} f(\bar{\mathbf{x}}^{(t)}) - f^*] \leq \varepsilon, \text{ after } \tilde{O} \left(\frac{\sigma^2}{n\varepsilon^2} + \frac{\sqrt{L}\sigma}{\sqrt{pc}\varepsilon^{3/2}} + \frac{L(1 + \tilde{R}_0^2 R_0^{-2})}{pc\varepsilon} \right) \cdot R_0^2 \text{ iterations,}$$

where $\tilde{R}_0^2 = \frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i^{(0)} - \bar{\mathbf{x}}^{(0)}\|^2 + \frac{1}{nL^2} \sum_{i=1}^n \|\mathbf{y}_i^{(0)} - \bar{\mathbf{y}}^{(0)}\|^2$.

From these results we see that the leading term in the convergence rate (assuming $\sigma > 0$) is not affected by the graph parameters. Moreover, in this term we see a linear speedup in n , the number of workers. The leading terms of all three results match with the convergence estimates for all-reduce mini-batch SGD [5, 43] and is optimal [37]. This means, that after a sufficiently long transient time, GT achieves a linear speedup in n . This transient time depends on the graph parameters p and c , but not on the data-dissimilarity. We will discuss the dependency of the convergence rate on the graph parameters c, p more carefully below in Sections 5 and 7, and compare the convergence rate to the convergence rates of D-SGD and D².

Possible Relaxations of the Assumptions. Before moving on to the proofs, we mention briefly a few possible relaxations of the assumptions that are possible with only slight adaptations of the proof framework. These extensions can be addressed with known techniques and are omitted for conciseness. We give here the necessary references for completeness.

- **Bounded Gradient Assumption I.** The uniform bound on the stochastic noise in Assumption 4 could be relaxed by allowing the noise to grow with the gradient norm [16, Assumption 3b].
- **Bounded Gradient Assumption II.** In the convex setting it has been observed that σ^2 can be replaced with $\sigma_*^2 := \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\xi_i} \|\nabla F_i(\mathbf{x}^*, \xi_i) - \nabla f_i(\mathbf{x}^*)\|_2^2$, the noise at the optimum. However, this requires smoothness of each $F_i(\mathbf{x}, \xi)$, $\xi \in \mathcal{D}_i$, which is stronger than our Assumption 2. For the technique see e.g. [38].
- **Different mixing for X and Y .** In Algorithm 1, both the \mathbf{x} and \mathbf{y} iterates are averaged on the same communication topology (the same mixing matrix). This can be relaxed by allowing for two separate matrices. This follows from inspecting our proof below.
- **Local Steps.** It is possible to extend Algorithm 1 and our analysis in Theorem 2 to allow for local computation steps. Mixing matrix would alternate between identity matrix I (no communication, local steps) and W (communication steps).

However, it is non trivial to extend our analysis to the general time-varying graphs, as the product of two arbitrary mixing matrices $W_1 W_2$ might be non symmetric.

4.2 Faster convergence on consensus functions

We now state an additional result, which improves Theorem 2 on the consensus problem, defined as

$$\min \left[f(\mathbf{x}) = \frac{1}{n} \sum_{i=1}^n [f_i(\mathbf{x}) := \frac{1}{2} \|\mathbf{x} - \boldsymbol{\mu}_i\|^2] \right], \quad (7)$$

for vectors $\boldsymbol{\mu}_i \in \mathbb{R}^d$, $i \in [n]$ and optimal solution $\mathbf{x}^* = \frac{1}{n} \sum_{i=1}^n \boldsymbol{\mu}_i$. Note that this is a special case of the general problem (1) without stochastic noise ($\sigma = 0$). For this function, we can improve the complexity estimate that would follow from Theorem 2 by proving a convergence rate that does not depend on c .

Theorem 3. *Let f be as in (7) let Assumption 1 hold. Then there exists a stepsize $\gamma \leq p$ such that it holds $\frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i^{(T)} - \mathbf{x}^*\|^2 \leq \epsilon$, for the iterates GT 1 and any $\epsilon > 0$, after at most $T = \tilde{\mathcal{O}}(p \log \frac{1}{\epsilon})$ iterations.*

5 Discussion

We now provide a discussion of these results.

5.1 Parameter c

The convergence rate in Theorem 2 depends on the parameter c , that in the worst case could be as small as p . In this case our theoretical result does not improve over existing results for the strongly convex case. However, for many graphs in practice parameter c is bounded by a constant (see Table 3 and discussion below Assumption 1).

While we show in Theorem 3 that it is possible to remove the dependency on c entirely from the convergence rate in special cases, it is still an open question if the parameter c in Theorem 2 is tight in general.

5.2 Comparison to prior GT literature

Tables 1 and 2 compare our theoretical convergence rates in strongly convex and non convex settings. Our result tightens all existing prior work.

5.3 Comparison to other methods.

We now compare our complexity estimate of GT to D-SGD and D^2 in the strongly convex case. Analogous observations hold for the other cases too.

Comparison to D-SGD. A popular algorithm for decentralized optimization is D-SGD [24] that converges as [16]:

$$\tilde{\mathcal{O}} \left(\frac{\sigma^2}{\mu n \epsilon} + \frac{\sqrt{L} (\zeta + \sqrt{p} \sigma)}{\mu p \sqrt{\epsilon}} + \frac{L}{\mu p} \log \frac{1}{\epsilon} \right). \quad (\text{D-SGD})$$

While GT is agnostic to data-heterogeneity, here the convergence estimate depends on the data-heterogeneity, measured by a constant ζ^2 that satisfies:

$$\frac{1}{n} \sum_{i=1}^n \|\nabla f_i(\mathbf{x}^*) - \nabla f(\mathbf{x}^*)\|_2^2 \leq \zeta^2. \quad (8)$$

Comparing with Theorem 2, GT completely removes dependence on data heterogeneity level ζ . Moreover, even in the homogeneous case when $\zeta = 0$, GT enjoys the same rate as D-SGD for many practical graphs when c is bounded by a constant.

Comparison to D^2 . Similarly to GT, D^2 also removes the dependence on functions heterogeneity. The convergence rate of D^2 holds under assumption that $\lambda_{\min}(W) > -\frac{1}{3}$ and it is equal to [55]:

$$\mathcal{O} \left(\frac{\sigma^2}{\mu n \epsilon} + \frac{\sqrt{L} \sigma}{\mu \sqrt{p} \sqrt{\epsilon}} + \frac{L}{\mu p} \log \frac{1}{\epsilon} \right). \quad (\text{D}^2)$$

Under the assumption $\lambda_{\min}(W) > -\frac{1}{3}$ the parameter c is a constant, and the GT rate estimated in Theorem 2 matches (D^2).

6 Proof sketch of the main theorem

Here we give a proof sketch for Theorem 2, for the special case of strongly convex objectives. We give all proof details in the appendix and highlight the main technical difficulties and novel techniques.

Key Lemma. It is very common—and useful—to write the iterates in the form $X^{(t)} = \bar{X}^{(t)} + (X^{(t)} - \bar{X}^{(t)})$, where $\bar{X}^{(t)}$ denotes the matrix with the average over the nodes. We can then separately analyze $\bar{X}^{(t)}$ and the consensus difference $\Delta X^{(t)} := (X^{(t)} - \bar{X}^{(t)})$ (and $\Delta Y^{(t)} := (Y^{(t)} - \bar{Y}^{(t)})$). Define $\tilde{W} = W - \frac{\mathbf{1}\mathbf{1}^\top}{n}$. From the update equation (GT) we see that

$$\begin{pmatrix} \Delta X^{(t+1)} \\ \gamma \Delta Y^{(t+1)} \end{pmatrix}^\top = \underbrace{\begin{pmatrix} \Delta X^{(t)} \\ \gamma \Delta Y^{(t)} \end{pmatrix}^\top}_{=: \Psi_t} \underbrace{\begin{pmatrix} \tilde{W} & 0 \\ -\tilde{W} & \tilde{W} \end{pmatrix}}_{=: J} + \gamma \underbrace{\begin{pmatrix} 0 \\ (\nabla F(X^{t+1}, \xi^{t+1}) - \nabla F(X^t, \xi^t)) (I - \frac{\mathbf{1}\mathbf{1}^\top}{n}) \end{pmatrix}^\top}_{=: E_t},$$

in short, by using the notation Ψ_t , J , and E_t as introduced above,

$$\Psi_{t+1} = \Psi_t J + \gamma E_t. \quad (9)$$

We could immediately adapt the proof technique from [16] if it would hold that the spectral radius of J is smaller than one. However, this is not the case, and in general $\|J\| > 1$.

Note that for any integer $i \geq 0$:

$$J^i = \begin{pmatrix} \tilde{W}^i & 0 \\ -i\tilde{W}^i & \tilde{W}^i \end{pmatrix} \quad \|J^i\|^2 = \|\tilde{W}^i\|^2 + i^2 \|\tilde{W}^i\|^2 \leq (1-p)^i + i^2(1-p)^i, \quad (10)$$

by Assumption 1. With this observation we can now formulate a key lemma:

Lemma 4 (Contraction). *For any integer $\tau \geq \frac{2}{p} \log \left(\frac{50}{p} (1 + \log \frac{1}{p}) \right)$ it holds that $\|J^\tau\|^2 \leq \frac{1}{2}$.*

While the constants in this lemma are chosen to ease the presentation, most important for us is that after $\tau = \tilde{\Theta}(\frac{1}{p})$ communication rounds, old parameter values (from τ steps ago) get discounted and averaged by a constant factor. We can alternatively write the statement of Lemma 4 as

$$\|ZJ^\tau - \bar{Z}\|_F^2 \leq \frac{1}{2} \|Z - \bar{Z}\|_F^2, \quad \forall Z \in \mathbb{R}^{2d \times n}.$$

This resembles [16, Assumption 4] and the proof now follows the same pattern. A few crucial differences remain, as the result in [16] depends on a data-dissimilarity parameter which we can avoid by carefully estimating the tracking errors. For completeness, we sketch the outline and give all details in the appendix.

Average Sequence. First, we consider the average sequences $\bar{X}^{(t)}$ and $\bar{Y}^{(t)}$. As all columns of these matrices are equal, we can equivalently consider a single column only: $\bar{\mathbf{x}}^{(t)}$ and $\bar{\mathbf{y}}^{(t)}$.

Lemma 5 (Average). *It holds that*

$$\bar{\mathbf{y}}^{(t)} = \frac{1}{n} \sum_{i=1}^n \nabla F_i(\mathbf{x}_i^{(t)}, \xi_i^{(t)}), \quad \bar{\mathbf{x}}^{(t+1)} = \bar{\mathbf{x}}^{(t)} - \gamma \frac{1}{n} \sum_{i=1}^n \nabla F_i(\mathbf{x}_i^{(t)}, \xi_i^{(t)}). \quad (11)$$

This follows directly from the update (GT) and the fact that $\bar{X} = \bar{X}W$ for doubly stochastic mixing matrices. The update of $\bar{\mathbf{x}}^{(t)}$ in (11) is almost identical to one step of mini-batch SGD (on a complete graph). The average sequence behaves almost as a SGD sequence:

Lemma 6 (Descent lemma, [16, Lemma 8]). *Under the Assumptions of Theorem 2 for the convex functions, the averages $\bar{\mathbf{x}}^{(t)} := \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i^{(t)}$ of the iterates of Algorithm 1 with the stepsize $\gamma \leq \frac{1}{12L}$ satisfy*

$$\mathbb{E} \|\bar{\mathbf{x}}^{(t+1)} - \mathbf{x}^*\|^2 \leq \left(1 - \frac{\gamma\mu}{2}\right) \mathbb{E} \|\bar{\mathbf{x}}^{(t)} - \mathbf{x}^*\|^2 + \frac{\gamma^2 \sigma^2}{n} - \gamma e_t + \frac{3\gamma L}{n} \sum_{i=1}^n \mathbb{E} \|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2, \quad (12)$$

where $e_t = \mathbb{E} f(\bar{\mathbf{x}}^{(t)}) - f^*$, for $f^* = \min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x})$.

Consensus Distance. The main difficulty comes from estimating the consensus distance $\|\Psi_t\|^2$, in the notation introduced in (9). Note that

$$\|\Psi_t\|^2 = \frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i^{(t)} - \bar{\mathbf{x}}^{(t)}\|_2^2 + \frac{\gamma^2}{n} \sum_{i=1}^n \|\mathbf{y}_i^{(t)} - \bar{\mathbf{y}}^{(t)}\|_2^2.$$

By unrolling (9) for $\tau \leq k \leq 2\tau$, $\tau = \frac{2}{p} \log \left(\frac{50}{p} (1 + \log \frac{1}{p}) \right) + 1$ steps,

$$\Psi_{t+k} = \Psi_t J^k + \gamma \sum_{j=1}^{k-1} E_{t+j-1} J^{k-j}. \quad (13)$$

By taking the Frobenius norm, and carefully estimating the norm of the error term $\left\| \sum_{j=1}^{\tau-1} E_{t+j-1} J^{\tau-j} \right\|_F^2$, and using Lemma 4 we can derive a recursion for the consensus distance.

Lemma 7 (Consensus distance recursion). *There exists absolute constants $B_1, B_2, B_3 > 0$ such that for a stepsize $\gamma < \frac{c}{B_3 L \tau}$*

$$\mathbb{E} \|\Psi_{t+k}\|_F^2 \leq \frac{7}{8} \mathbb{E} \|\Psi_t\|_F^2 + \frac{1}{128\tau} \sum_{j=0}^{k-1} \|\Psi_{t+j}\|_F^2 + \frac{B_1 \tau L \gamma^2}{c^2} \sum_{j=0}^{k-1} n e_{t+j} + \frac{B_2 \tau \gamma^2}{c^2} n \sigma^2. \quad (14)$$

This lemma allows to replace p with c in the final convergence rate. This is achieved by grouping same gradients in the sum $\left\| \sum_{j=1}^{k-1} E_{t+j-1} J^{k-j} \right\|_F^2$ and estimating the norm with Lemma 13.

An additional technical difficulty comes when unrolling consensus recursion (14). As iteration matrix J is not contractive, i.e. $\|J\| > 1$, then $\|\Psi_{t+j}\|_F^2$ for $j < \tau$ can be larger than $\|\Psi_t\|_F^2$ (up to $\approx \frac{1}{p^2}$ times as $\|J^i\|^2 \leq \mathcal{O}\left(\frac{1}{p^2}\right) \forall i$). We introduce an additional term in the recursion that is provably non-increasing

$$\Phi_{t+\tau} := \frac{1}{\tau} \sum_{j=0}^{\tau-1} \|\Psi_{t+j}\|_F^2.$$

With this we unroll consensus recursion.

Lemma 8 (Unrolling recursion). *For $\gamma < \frac{c}{\sqrt{7} B_1 L \tau} \leq \frac{1}{2L\tau}$ it holds,*

$$\mathbb{E} \|\Psi_t\|_F^2 \leq \left(1 - \frac{1}{64\tau}\right)^t A_0 + \frac{22B_1 \tau L \gamma^2}{c^2} \sum_{j=0}^{t-1} \left(1 - \frac{1}{64\tau}\right)^{t-j} n e_j + \frac{20B_2 \tau \gamma^2}{c^2} n \sigma^2 \quad (15)$$

where $e_j = \mathbb{E}[f(\bar{\mathbf{x}}^{(j)}) - f(\mathbf{x}^*)]$, $A_0 = 16\|\Delta X^{(0)}\|_F^2 + \frac{24\gamma^2}{p^2} \|\Delta Y^{(0)}\|_F^2$.

It remains to combine (14) and (15) using technique from [16]. \square

Proof sketch of Theorem 3. Using the matrix notation introduced above, the iterations of GT on problem (7) can be written in a simple form:

$$\begin{pmatrix} \Delta X^{(t+1)} \\ \gamma \Delta Y^{(t+1)} \end{pmatrix}^\top = \begin{pmatrix} \Delta X^{(t)} \\ \gamma \Delta Y^{(t)} \end{pmatrix}^\top \underbrace{\begin{pmatrix} \tilde{W} & \gamma(W - I) \\ -\tilde{W} & (1 - \gamma)\tilde{W} \end{pmatrix}}_{J'}.$$

Similar as above, also the matrix J' is not a contraction operator, but in contrast to J it is diagonalizable: $J' = Q\Lambda Q^{-1}$ for some Q and diagonal Λ . It follows that $\|(J')^t\|^2 = \|Q\Lambda^t Q^{-1}\|^2$ is decreasing as $(1-p)^t \|Q\|^2 \|Q^{-1}\|^2$. With this observation, the proof simplifies. \square

7 Experiments

In this section we investigate the tightness of parameters c and p in our theoretical result.

Setup. We consider simple quadratic functions defined as $f_i(\mathbf{x}) = \|\mathbf{x}\|^2$, and $\mathbf{x}^{(0)}$ is randomly initialized from a normal distribution $\mathcal{N}(0, 1)$. We add artificially stochastic noise to gradients as

$\nabla F_i(\mathbf{x}, \xi) = \nabla f_i(\mathbf{x}) + \xi$, where $\xi \sim \mathcal{N}(0, \frac{\sigma^2}{d} I)$ so that Assumption 4 is satisfied. We elaborate the details as well as results under other problem setups in Appendix C.

We verify the dependence on graph parameters p and c for the stochastic noise term. We fix the stepsize γ to be constant, vary p and c and measure the value of $f(\bar{\mathbf{x}}^{(t)}) - f^*$ that GT reaches after a large number of steps. According to the theory, GT converges to the level $\mathcal{O}\left(\frac{\gamma\sigma^2}{n} + \frac{\gamma^2\sigma^2}{pc^2}\right)$ in a linear number of steps (to reach higher accuracy, smaller stepsizes must be used). To decouple the second term we need to ensure that the first term is small enough. For that, we take the number of nodes n to be large. In all experiments we ensure that the first term is at least by order of magnitude smaller than the second by comparing the noise level with GT on a fully-connected topology.

The effect of p . First, in Figure 1 we verify the expected $\mathcal{O}\left(\frac{1}{p}\right)$ dependence when c is a constant. For a fixed $n = 300$ number of nodes with $d = 100$ we vary the value of a parameter p by interpolating the ring topology (with uniform weights) with the fully-connected graph. The loss value $f(\mathbf{x}^{(\infty)})$ scales linearly in $\frac{1}{p}$ as can be observed in Figure 1 and the dependency on p can thus not further be improved.




Figure 1: Impact of p on convergence with the stochastic noise $\sigma^2 = 1$, when c and γ are kept constant. We see a linear scaling in $\frac{1}{p}$ that verifies the $\mathcal{O}\left(\frac{1}{p}\right)$, dependence rather than prior predicted $\mathcal{O}\left(\frac{1}{p^2}\right)$.

The effect of c . In Figure 2 we aim to examine the dependence of the term $\mathcal{O}\left(\frac{1}{pc^2}\right)$ on the parameter c , in terms of $1/(pc^2)$ and $1/(cp)$. We take the ring topology on a fixed number of $n = 300$ nodes and reduce the self-weights to achieve different values of c (see appendix for details). Otherwise the setup is as above. The current numerical results may suggest the existence of a potentially better theoretical dependence of the term c (as discussed in Section 4.2); we leave the study for future work.




Figure 2: Impact of c on the convergence with the stochastic noise $\sigma^2 = 1$, when p and γ are kept constant. We see a near linear scaling in $\mathcal{O}\left(\frac{1}{pc}\right)$ while the estimate $\mathcal{O}\left(\frac{1}{pc^2}\right)$ appears to be too conservative on this problem.

8 Conclusion

We have derived improved complexity bounds for the GT method, that improve over all previous results. We verify the tightness of the second term in the convergence rate in numerical experiments. Our analysis identifies that the smallest eigenvalue of the mixing matrix has a strong impact on the performance of GT, however the smallest eigenvalue can often be controlled in practice by choosing large enough self-weights (w_{ii}) on the nodes.

Our proof technique might be of independent interest in the community and might lead to improved analyses for other gossip based methods where the mixing matrix is not contracting (for e.g. in directed graphs, or using row- or column-stochastic matrices).

Acknowledgments and Disclosure of Funding

This project was supported by SNSF grant 200020_200342, EU project DIGIPREDICT, and a Google PhD Fellowship. The authors thank Martin Jaggi for his support.

References

- [1] Sulaiman Alghunaim, Ernest Ryu, Kun Yuan, and Ali Sayed. Decentralized proximal gradient algorithms with linear convergence rates. *IEEE Transactions on Automatic Control*, 66(6), 2021.
- [2] Sulaiman A. Alghunaim and Ali H. Sayed. Linear convergence of primal–dual gradient methods and their performance in distributed optimization. *Automatica*, 117:109003, 2020.
- [3] Mahmoud Assran, Nicolas Loizou, Nicolas Ballas, and Michael Rabbat. Stochastic gradient push for distributed deep learning. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*. PMLR, 2019.
- [4] Stephen Boyd, Arpita Ghosh, Balaji Prabhakar, and Devavrat Shah. Randomized gossip algorithms. *IEEE/ACM Trans. Netw.*, 14(SI):2508–2530, 2006.
- [5] Ofer Dekel, Ran Gilad-Bachrach, Ohad Shamir, and Lin Xiao. Optimal distributed online prediction using mini-batches. *Journal of Machine Learning Research (JMLR)*, 13(1):165–202, 2012.
- [6] J. C. Duchi, A. Agarwal, and M. J. Wainwright. Dual averaging for distributed optimization: Convergence analysis and network scaling. *IEEE Transactions on Automatic Control*, 57(3):592–606, 2012.
- [7] S. Gerschgorin. Über die Abgrenzung der Eigenwerte einer Matrix. *Bulletin de l’Académie des Sciences de l’URSS. Classe des sciences mathématiques et na*, 6:749–754, 1931.
- [8] Lie He, An Bian, and Martin Jaggi. COLA: Decentralized linear learning. In *Advances in Neural Information Processing Systems 31 (NeurIPS)*, pages 4541–4551. Curran Associates, Inc., 2018.
- [9] Wolfram Research, Inc. Mathematica, Version 12.3, 2021. Champaign, IL.
- [10] Franck Iutzeler, Pascal Bianchi, Philippe Ciblat, and Walid Hachem. Asynchronous distributed optimization using a randomized alternating direction method of multipliers. In *Proceedings of the 52nd IEEE Conference on Decision and Control, CDC*, pages 3671–3676. IEEE, 2013.
- [11] B. Johansson, M. Rabi, and M. Johansson. A randomized incremental subgradient method for distributed optimization in networked systems. *SIAM Journal on Optimization*, 20(3):1157–1170, 2010.
- [12] Peter Kairouz, H. Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Keith Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, Rafael G. L. D’Oliveira, Hubert Eichner, Salim El Rouayheb, David Evans, Josh Gardner, Zachary Garrett, Adrià Gascón, Badih Ghazi, Phillip B. Gibbons, Marco Gruteser, Zaid Harchaoui, Chaoyang He, Lie He, Zhouyuan Huo, Ben Hutchinson, Justin Hsu, Martin Jaggi, Tara Javidi, Gauri Joshi, Mikhail Khodak, Jakub Konečný, Aleksandra Korolova, Farinaz Koushanfar, Sanmi Koyejo, Tancrède Lepoint, Yang Liu, Prateek Mittal, Mehryar Mohri, Richard Nock, Ayfer Özgür, Rasmus Pagh, Mariana Raykova, Hang Qi, Daniel Ramage, Ramesh Raskar, Dawn Song, Weikang Song, Sebastian U. Stich, Ziteng Sun, Ananda Theertha Suresh, Florian Tramèr, Praneeth Vepakomma, Jianyu Wang, Li Xiong, Zheng Xu, Qiang Yang, Felix X. Yu, Han Yu, and Sen Zhao. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2):1–210, 2021.
- [13] Sai P. Karimireddy, Satyen Kale, Mehryar Mohri, Sashank J. Reddi, Sebastian U. Stich, and Ananda T. Suresh. SCAFFOLD: Stochastic controlled averaging for on-device federated learning. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*. PMLR, 2019.
- [14] David Kempe, Alin Dobra, and Johannes Gehrke. Gossip-based computation of aggregate information. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. IEEE Computer Society, 2003.
- [15] Anastasia Koloskova, Sebastian Stich, and Martin Jaggi. Decentralized stochastic optimization and gossip algorithms with compressed communication. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, volume 97, pages 3478–3487. PMLR, 2019.

- [16] Anastasia Koloskova, Nicolas Loizou, Sadra Boreiri, Martin Jaggi, and Sebastian U. Stich. A unified theory of decentralized sgd with changing topology and local updates. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*. PMLR, 2020.
- [17] Jakub Konečný, H. Brendan McMahan, Daniel Ramage, and Peter Richtárik. Federated optimization: Distributed machine learning for on-device intelligence. *arXiv preprint arXiv:1610.02527*, 2016.
- [18] Lingjing Kong, Tao Lin, Anastasia Koloskova, Martin Jaggi, and Sebastian U. Stich. Consensus control for decentralized deep learning. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, volume 139, pages 5686–5696. PMLR, 2021.
- [19] Dmitry Kovalev, Anastasia Koloskova, Martin Jaggi, Peter Richtarik, and Sebastian U. Stich. A linearly convergent algorithm for decentralized optimization: Sending less bits for free! In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 130, pages 4087–4095. PMLR, 2021.
- [20] Boyue Li, Shicong Cen, Yuxin Chen, and Yuejie Chi. Communication-efficient distributed optimization in networks with gradient tracking and variance reduction. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 108, pages 1662–1672. PMLR, 2020.
- [21] Tian Li, Anit Kumar Sahu, Maziar Sanjabi, Manzil Zaheer, Ameet Talwalkar, and Virginia Smith. On the convergence of federated optimization in heterogeneous networks. *arXiv preprint arXiv:1812.06127*, 2018.
- [22] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Fed-DANE: A federated Newton-type method. *arXiv preprint arXiv:2001.01920*, 2020.
- [23] Xiang Li, Kaixuan Huang, Wenhao Yang, Shusen Wang, and Zhihua Zhang. On the convergence of FedAvg on non-IID data. *International Conference on Learning Representations (ICLR)*, 2020.
- [24] Xiangru Lian, Ce Zhang, Huan Zhang, Cho-Jui Hsieh, Wei Zhang, and Ji Liu. Can decentralized algorithms outperform centralized algorithms? a case study for decentralized parallel stochastic gradient descent. In *Advances in Neural Information Processing Systems 30 (NIPS)*, pages 5330–5340. Curran Associates, Inc., 2017.
- [25] Tao Lin, Sai Praneeth Karimireddy, Sebastian Stich, and Martin Jaggi. Quasi-global momentum: Accelerating decentralized deep learning on heterogeneous data. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, volume 139, pages 6654–6665. PMLR, 2021.
- [26] Paolo Di Lorenzo and Gesualdo Scutari. NEXT: In-network nonconvex optimization. *IEEE Transactions on Signal and Information Processing over Networks*, 2(2):120–136, 2016.
- [27] Songtao Lu, Xinwei Zhang, Haoran Sun, and Mingyi Hong. GNSD: a gradient-tracking based nonconvex stochastic algorithm for decentralized optimization. In *IEEE Data Science Workshop (DSW)*, pages 315–321, 2019.
- [28] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1273–1282, 2017.
- [29] H. Brendan McMahan, Eider Moore, Daniel Ramage, and Blaise Agüera y Arcas. Federated learning of deep networks using model averaging. *arXiv preprint arXiv:1602.05629*, 2016.
- [30] A. Nedić and A. Ozdaglar. Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on Automatic Control*, 54(1):48–61, 2009.
- [31] A. Nedić, S. Lee, and M. Raginsky. Decentralized online optimization with global objectives and local communication. In *2015 American Control Conference (ACC)*, pages 4497–4503, 2015.
- [32] Angelia Nedić and Alex Olshevsky. Distributed optimization over time-varying directed graphs. *IEEE Transactions on Automatic Control*, 60(3):601–615, 2014.
- [33] Angelia Nedić. Distributed gradient methods for convex machine learning problems in networks: Distributed optimization. *IEEE Signal Processing Magazine*, 37(3):92–101, 2020.
- [34] Angelia Nedić, Alex Olshevsky, and Wei Shi. Achieving geometric convergence for distributed optimization over time-varying graphs. *SIAM Journal on Optimization*, 27, 07 2016.

- [35] Angelia Nedić, Alex Olshevsky, Wei Shi, and César A. Uribe. Geometrically convergent distributed optimization with uncoordinated step-sizes. In *2017 American Control Conference (ACC)*, pages 3950–3955, 2017.
- [36] Angelia Nedić, Alex Olshevsky, and Michael G. Rabbat. Network topology and communication-computation tradeoffs in decentralized optimization. *Proceedings of the IEEE*, 106(5):953–976, 2018.
- [37] Arkadii Semenovich Nemirovsky and David Borisovich Yudin. *Problem complexity and method efficiency in optimization*. Wiley, 1983.
- [38] Lam M. Nguyen, Phuong Ha Nguyen, Peter Richtárik, Katya Scheinberg, Martin Takáč, and Marten van Dijk. New convergence aspects of stochastic gradient algorithms. *Journal of Machine Learning Research*, 20(176):1–49, 2019.
- [39] Shi Pu and Angelia Nedić. Distributed stochastic gradient tracking methods. *Math. Program.*, 187:409–457, 2020.
- [40] Guannan Qu and Na Li. Harnessing smoothness to accelerate distributed optimization. *IEEE Transactions on Control of Network Systems*, 2017.
- [41] M. Rabbat. Multi-agent mirror descent for decentralized stochastic optimization. In *IEEE 6th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, pages 517–520, 2015.
- [42] Wei Shi, Qing Ling, Gang Wu, and Wotao Yin. EXTRA: An exact first-order algorithm for decentralized consensus optimization. *SIAM Journal on Optimization*, 25(2):944–966, 2015.
- [43] Sebastian U. Stich. Unified optimal analysis of the (stochastic) gradient method. *arXiv preprint arXiv:1907.04232*, 2019.
- [44] Sebastian U. Stich and Sai P. Karimireddy. The error-feedback framework: Better rates for SGD with delayed gradients and compressed communication. *Journal of Machine Learning Research*, 21(237):1–36, 2020.
- [45] Hanlin Tang, Shaoduo Gan, Ce Zhang, Tong Zhang, and Ji Liu. Communication compression for decentralized training. In *Advances in Neural Information Processing Systems 31 (NeurIPS)*, pages 7663–7673. Curran Associates, Inc., 2018.
- [46] Hanlin Tang, Xiangru Lian, Ming Yan, Ce Zhang, and Ji Liu. D^2 : Decentralized training over decentralized data. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, volume 80, pages 4848–4856. PMLR, 2018.
- [47] Hanlin Tang, Xiangru Lian, Shuang Qiu, Lei Yuan, Ce Zhang, Tong Zhang, and Ji Liu. Deepsqueeze: Decentralization meets error-compensated compression. *arXiv preprint arXiv:1907.07346*, 2019.
- [48] John N. Tsitsiklis. *Problems in decentralized decision making and computation*. PhD thesis, Massachusetts Institute of Technology, 1984.
- [49] Isidoros Tziotis, Constantine Caramanis, and Aryan Mokhtari. Second order optimality in decentralized non-convex optimization via perturbed gradient tracking. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 21162–21173. Curran Associates, Inc., 2020.
- [50] Thijs Vogels, Lie He, Anastasia Koloskova, Tao Lin, Sai Praneeth Karimireddy, Sebastian U. Stich, and Martin Jaggi. Relaysum for decentralized deep learning on heterogeneous data. In *Advances in Neural Information Processing Systems (NeurIPS)*. Curran Associates, Inc., 2021.
- [51] Jianyu Wang, Anit Kumar Sahu, Zhouyi Yang, Gauri Joshi, and Soumya Kar. MATCHA: Speeding up decentralized SGD via matching decomposition sampling. In *2019 Sixth Indian Control Conference (ICC)*, pages 299–300, 2019.
- [52] E. Wei and A. Ozdaglar. Distributed alternating direction method of multipliers. In *IEEE 51st IEEE Conference on Decision and Control (CDC)*, pages 5445–5450, 2012.
- [53] Lin Xiao and Stephen Boyd. Fast linear iterations for distributed averaging. *Systems & Control Letters*, 53(1):65–78, 2004.
- [54] Ran Xin, Anit Kumar Sahu, Usman A. Khan, and Soumya Kar. Distributed stochastic optimization with gradient tracking over strongly-connected networks. In *IEEE 58th Conference on Decision and Control (CDC)*, pages 8353–8358, 2019.

- [55] Kun Yuan and Sulaiman A. Alghunaim. Removing data heterogeneity influence enhances network topology dependence of decentralized SGD. *arXiv preprint arXiv:2105.08023*, 2021.
- [56] Kun Yuan, Bicheng Ying, Xiaochuan Zhao, and Ali H. Sayed. Exact diffusion for distributed optimization and learning - part I: algorithm development. *IEEE Trans. Signal Process.*, 67(3):708–723, 2019.
- [57] Kun Yuan, Bicheng Ying, Xiaochuan Zhao, and Ali H. Sayed. Exact diffusion for distributed optimization and learning - part II: convergence analysis. *IEEE Trans. Signal Process.*, 67(3):724–739, 2019.
- [58] Kun Yuan, Wei Xu, and Qing Ling. Can primal methods outperform primal-dual methods in decentralized dynamic optimization? *IEEE Transactions on Signal Processing*, 68:4466–4480, 2020.
- [59] Kun Yuan, Yiming Chen, Xinmeng Huang, Yingya Zhang, Pan Pan, Yinghui Xu, and Wotao Yin. DecentLaM: Decentralized momentum SGD for large-batch deep training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3029–3039, 2021.
- [60] Jiaqi Zhang and Keyou You. Decentralized stochastic gradient tracking for non-convex empirical risk minimization. *arXiv preprint arXiv:1909.02712*, 2020.
- [61] Minghui Zhu and Sonia Martínez. Discrete-time dynamic average consensus. *Automatica*, 46(2):322–329, 2010. ISSN 0005-1098.

A Proof of Theorem 3 — Consensus Functions

We consider functions $f_i(\mathbf{x}) = \frac{1}{2} \|\mathbf{x} - \boldsymbol{\mu}_i\|^2$, where $\mathbf{x}, \boldsymbol{\mu}_i \in \mathbb{R}^d$. Then $\nabla f_i(\mathbf{x}) = \mathbf{x} - \boldsymbol{\mu}_i$. In matrix notation, the GT algorithm in this special case is equivalent to

$$\begin{pmatrix} X^{(t+1)} \\ \gamma Y^{(t+1)} \end{pmatrix}^\top = \begin{pmatrix} X^{(t)} \\ \gamma Y^{(t)} \end{pmatrix}^\top \begin{pmatrix} W & -W \\ 0 & W \end{pmatrix} + \gamma \begin{pmatrix} 0 \\ X^{(t+1)} - X^{(t)} \end{pmatrix}^\top = \begin{pmatrix} X^{(t)} \\ \gamma Y^{(t)} \end{pmatrix}^\top \begin{pmatrix} W & -W \\ \gamma(W - I) & (1 - \gamma)W \end{pmatrix}.$$

The optimal point $\mathbf{x}^* = \bar{\boldsymbol{\mu}} = \frac{1}{n} \sum_{i=1}^n \boldsymbol{\mu}_i$. Denote $X^* = [\mathbf{x}^*, \dots, \mathbf{x}^*] \in \mathbb{R}^{d \times n}$. We decompose the error as

$$\|X^{(t)} - X^*\|_F^2 = \underbrace{\|X^{(t)} - \bar{X}^{(t)}\|_F^2}_{\text{consensus error}} + \underbrace{\|\bar{X}^{(t)} - X^*\|_F^2}_{\text{optimization error}}.$$

For the optimization part, notice that $\bar{Y}^{(t)} = \bar{X}^{(t)} - X^*$. That is because

$$\bar{Y}^{(0)} = \nabla f(X^{(0)}) \frac{1}{n} \mathbf{1}\mathbf{1}^\top = \bar{X}^{(0)} - X^*, \quad \bar{Y}^{(t+1)} = \bar{Y}^{(t)} + \bar{X}^{(t+1)} - \bar{X}^{(t)}.$$

Therefore, the optimization error is equal to

$$\begin{aligned} \|\bar{X}^t - X^*\|_F^2 &= \|\bar{X}^{(t-1)} - \gamma \bar{Y}^{(t-1)} - X^*\|_F^2 = \|(1 - \gamma) (\bar{X}^{(t-1)} - X^*)\|_F^2 \\ &= (1 - \gamma)^{2t} \|\bar{X}^{(0)} - X^*\|_F^2. \end{aligned}$$

For the consensus part, denoting, $\tilde{W} = W - \frac{\mathbf{1}\mathbf{1}^\top}{n}$, $\Delta X^{(t)} = X^{(t)} - \bar{X}^{(t)}$, $\Delta Y^{(t)} = Y^{(t)} - \bar{Y}^{(t)}$,

$$\begin{pmatrix} \Delta X^{(t)} \\ \gamma \Delta Y^{(t)} \end{pmatrix}^\top = \begin{pmatrix} \Delta X^{(0)} \\ \gamma \Delta Y^{(0)} \end{pmatrix}^\top \underbrace{\begin{pmatrix} \tilde{W} & -\tilde{W} \\ \gamma(W - I) & (1 - \gamma)\tilde{W} \end{pmatrix}^t}_{J'}.$$

Taking the norm,

$$\|\Delta X^{(t)}\|_F^2 + \gamma^2 \|\Delta Y^{(t)}\|_F^2 \leq \|J^t\|_2^2 \left(\|\Delta X^{(0)}\|_F^2 + \gamma^2 \|\Delta Y^{(0)}\|_F^2 \right).$$

Lets analyze spectral properties of matrix J^t . Let the eigenvalue decomposition of W be $W = U\Lambda U^\top$, the eigenvalue decomposition of \tilde{W} is $\tilde{W} = U\tilde{\Lambda}U^\top$ for diagonal $\tilde{\Lambda}$.

We can decompose

$$J' = \begin{pmatrix} U & 0 \\ 0 & U \end{pmatrix} \underbrace{\begin{pmatrix} \tilde{\Lambda} & -\tilde{\Lambda} \\ \gamma(\Lambda - I) & (1 - \gamma)\tilde{\Lambda} \end{pmatrix}}_{=:M} \begin{pmatrix} U^\top & 0 \\ 0 & U^\top \end{pmatrix}.$$

And,

$$\|J^t\|_2^2 = \left\| \begin{pmatrix} U & 0 \\ 0 & U \end{pmatrix} \begin{pmatrix} \tilde{\Lambda} & -\tilde{\Lambda} \\ \gamma(\Lambda - I) & (1 - \gamma)\tilde{\Lambda} \end{pmatrix}^t \begin{pmatrix} U^\top & 0 \\ 0 & U^\top \end{pmatrix} \right\|_2^2 = \left\| \begin{pmatrix} \tilde{\Lambda} & -\tilde{\Lambda} \\ \gamma(\Lambda - I) & (1 - \gamma)\tilde{\Lambda} \end{pmatrix}^t \right\|_2^2,$$

where the last equality is due to unitary property of U .

Lemma 9. To diagonalize a block-diagonal matrix

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix},$$

where $A = \text{diag}(a_0, \dots, a_n) \in R^{n \times n}$, $B = \text{diag}(b_0, \dots, b_n)$, $C = \text{diag}(c_0, \dots, c_n)$, $D = \text{diag}(d_0, \dots, d_n)$. Assume that each of the 2×2 matrices

$$\begin{pmatrix} a_i & b_i \\ c_i & d_i \end{pmatrix}$$

are diagonalizable with

$$\begin{pmatrix} a_i & b_i \\ c_i & d_i \end{pmatrix} = \begin{pmatrix} q_i^{(1)} & q_i^{(2)} \\ q_i^{(3)} & q_i^{(4)} \end{pmatrix} \cdot \begin{pmatrix} d_i^{(1)} & 0 \\ 0 & d_i^{(2)} \end{pmatrix} \cdot \begin{pmatrix} q_i^{(-1)} & q_i^{(-2)} \\ q_i^{(-3)} & q_i^{(-4)} \end{pmatrix}$$

Then the original matrix is diagonalizable and its diagonalization is equal to

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix} = \begin{pmatrix} Q_1 & Q_2 \\ Q_3 & Q_4 \end{pmatrix} \cdot \begin{pmatrix} D_1 & 0 \\ 0 & D_2 \end{pmatrix} \cdot \begin{pmatrix} Q_{-1} & Q_{-2} \\ Q_{-3} & Q_{-4} \end{pmatrix},$$

where each $Q_l = \text{diag}(q_1^{(l)}, \dots, q_n^{(l)})$, $D_l = \text{diag}(d_1^{(l)}, \dots, d_n^{(l)})$.

We need to show that the following 2×2 matrices are diagonalizable.

$$M_i := \begin{pmatrix} \lambda_i & -\lambda_i \\ \gamma(\lambda_i - 1) & (1 - \gamma)\lambda_i \end{pmatrix},$$

where the λ_i are eigenvalues of the matrix \tilde{W} . The eigenvalues of M_i are

$$\lambda(M_i) = \left\{ \lambda_i - \frac{\gamma\lambda_i}{2} - \frac{1}{2}\sqrt{\gamma\lambda_i\sqrt{4 + (\gamma - 4)\lambda_i}}, \lambda_i - \frac{\gamma\lambda_i}{2} + \frac{1}{2}\sqrt{\gamma\lambda_i\sqrt{4 + (\gamma - 4)\lambda_i}} \right\},$$

which are distinct for $\gamma > 0$, therefore the matrix is diagonalizable (over \mathbb{C}).

If λ_i is positive, then by choosing $\gamma \leq 1 - \lambda_i$,

$$|\lambda(M_i)| \leq \frac{1}{3}\lambda_i + \frac{2}{3}.$$

If λ_i is negative, then, then by choosing $\gamma \leq 1 - |\lambda_i|$,

$$|\lambda(M_i)| \leq \frac{1}{3}|\lambda_i| + \frac{2}{3}.$$

We do not give the full formal prove of these two bounds. First we note that $|\lambda(M_i)|$ is monotone in γ , i.e. the absolute value increases in γ . Therefore it is enough to check that it holds $|\lambda(M_i)| \leq \frac{1}{3}|\lambda_i| + \frac{2}{3}$ for $\gamma = 1 - |\lambda_i|$. We visualize these upper bounds with Mathematica [9] in Figure 3.




Figure 3: The upper bound $\frac{1}{3}|\lambda_i| + \frac{2}{3}$ (yellow) vs. the true $|\lambda(M_i)|$ for the choice $\gamma = 1 - |\lambda_i|$.

This concludes the proof.

B Proof of Theorem 2 — General Case

We first re-state theorem 2 in terms of number of iterations T

Theorem 10. *For GT algorithm 1 with a mixing matrix as in Definition 1, under Assumptions 1, 2, 4, after T iterations, if $T > \frac{2}{p} \log \left(\frac{50}{p} (1 + \log \frac{1}{p}) \right)$, there exists a constant stepsize $\gamma_t = \gamma$ such that the error is bounded as*

Non-convex:

$$\frac{1}{T+1} \sum_{t=0}^T \left\| \nabla f(\bar{\mathbf{x}}^{(t)}) \right\|_2^2 \leq \tilde{\mathcal{O}} \left(\sqrt{\frac{LF_0\sigma^2}{nT}} + \left(\frac{\sigma LF_0}{(\sqrt{pc} + p\sqrt{n})T} \right)^{2/3} + \frac{L(F_0 + L\tilde{R}_0^2)}{pcT} \right),$$

Strongly-convex: Under additional Assumption 3 with $\mu > 0$, it holds

$$\sum_{t=0}^T \frac{w_t}{W_T} \left[\mathbb{E} f(\bar{\mathbf{x}}^{(t)}) - f^* \right] + \frac{\mu}{2} R_{T+1} \leq \tilde{\mathcal{O}} \left(\frac{\sigma^2}{\mu n T} + \frac{L\sigma^2}{\mu^2 pc^2 T^2} + \frac{L(R_0^2 + \frac{L}{\mu} \tilde{R}_0^2)}{pc} \exp \left[-\frac{\mu pc T}{L} \right] \right),$$

Weakly-convex: Under Assumptions 3 with $\mu \geq 0$, it holds

$$\frac{1}{T+1} \sum_{t=0}^T \left[\mathbb{E} f(\bar{\mathbf{x}}^{(t)}) - f^* \right] \leq \tilde{\mathcal{O}} \left(\sqrt{\frac{R_0^2\sigma^2}{nT}} + \left(\frac{\sigma\sqrt{L}R_0^2}{\sqrt{pc}T} \right)^{2/3} + \frac{L(R_0^2 + \tilde{R}_0^2)}{pcT} \right),$$

where $F_0 = f(\bar{\mathbf{x}}^{(0)}) - f^*$, $R_t = \|\mathbf{x}^{(t)} - \mathbf{x}^*\|$, $t \in \{0, T+1\}$, $\tilde{R}_0^2 = \frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i^{(0)} - \bar{\mathbf{x}}^{(0)}\|^2 + \frac{1}{n} \sum_{i=1}^n \|\mathbf{y}_i^{(0)} - \bar{\mathbf{y}}^{(0)}\|^2$.

B.1 Useful Inequalities

Proof of Lemma 4. By monotonicity, it suffices to check the inequality for $i = \tau$. By using $(1-p)^i \leq e^{-ip}$ and plugging τ into (10) it follows:

$$\begin{aligned} \|J^i\|^2 &\leq e^{-\tau p} (1 + \tau^2) \leq \frac{p^2}{50^2 (1 + \log \frac{1}{p})^2} \left(1 + \frac{(2(\log(50) + \log(\frac{1}{p}(1 + \log \frac{1}{p})))^2)}{p^2} \right) \\ &\leq \frac{1}{50^2} + \frac{1}{10} + \frac{1}{4} \end{aligned}$$

with $\log(\frac{1}{p}(1 + \log \frac{1}{p})) \leq \log \frac{1}{p} + \log \log \frac{1}{p} \leq 2 \log \frac{1}{p}$, then $(\log(4) + 2 \log \frac{1}{p})^2 \leq 2 \log 4 + 8 \log \frac{1}{p}$, and $(4 \log 50 + 16 \log \frac{1}{p})^2 \leq (128 + 512 \log \frac{1}{p})$. \square

Lemma 11. *Let $\lambda \in (-1, 1)$ with $|\lambda| = 1 - \alpha$, for $0 < \alpha < 1$. Then $|i\lambda^i| \leq \frac{1}{\alpha}$ for all $i \geq 0$.*

Proof. $|i\lambda^i| \leq i(1 - \alpha)^i \leq \sum_{j=1}^i (1 - \alpha)^j \leq \frac{1 - \alpha}{\alpha}$. \square

Lemma 12 (fact). *Let W be a symmetric matrix with eigenvalues $\lambda_1(W) \geq \dots \geq \lambda_n(W)$. Then $\|W\|^2 = \max_i \lambda_i^2(W)$.*

Lemma 13. *It holds $\left\| (i+1)\tilde{W}^{i+1} - i\tilde{W}^i \right\|^2 \leq \frac{4}{\alpha^2} \leq \frac{16}{c^2}$ for all $i \geq 0$, where $\alpha = 1 - |\lambda_n(W)|$ and c as defined in (3).*

Proof. The eigenvalues of $(i+1)\tilde{W}^{i+1} - i\tilde{W}^i$ have the form $(i+1)\lambda^{i+1} - i\lambda^i$, for $\lambda \in \Lambda := \{\lambda_1(\tilde{W}), \dots, \lambda_n(\tilde{W})\}$, the eigenvalues of \tilde{W} . By Lemma 12, it holds

$$\left\| (i+1)\tilde{W}^{i+1} - i\tilde{W}^i \right\|^2 = \max_{\lambda \in \Lambda} ((i+1)\lambda^{i+1} - i\lambda^i)^2.$$

If the maximum is attained for a positive $\lambda > 0$, we conclude

$$\begin{aligned} ((i+1)\lambda^{i+1} - i\lambda^i)^2 &= (\lambda^{i+1} - i\lambda^i(1-\lambda))^2 \\ &\leq 2(\lambda^{i+1})^2 + 2(1-\lambda)^2(i\lambda^i)^2 \\ &\leq 2(\lambda^{i+1})^2 + 2\frac{(1-\lambda)^2}{(1-\lambda)^2} \\ &\leq 4 \end{aligned}$$

with Lemma 11 for the first estimate and using $\lambda \leq 1$ on the last line. If the maximum is attained for a negative $\lambda < 0$ with $\lambda = -1 + \beta$, for $\beta > 0$, then

$$\begin{aligned} ((i+1)\lambda^{i+1} - i\lambda^i)^2 &\leq 2((i+1)\lambda^{i+1})^2 + 2(i\lambda^i)^2 \\ &\leq \frac{2}{\beta^2} + \frac{2}{\beta^2} \leq \frac{4}{\alpha^2} \end{aligned}$$

with Lemma 11 and $\alpha \leq \beta$.

Note that $c = 1 - (1 - \alpha)^2 = 2\alpha - \alpha^2 \geq \alpha$, since $\alpha(1 - \alpha) \geq 0$ and that $c \leq 2\alpha$. \square

Lemma 14. It holds $\|i\tilde{W}^i\|^2 \leq \frac{1}{\alpha^2} \leq \frac{4}{p^2}$.

Proof. $\|i\tilde{W}^i\|^2 = \left(i\|\tilde{W}^i\|\right)^2$, and the proof follows with Lemma 11 and 12 from above. \square

Lemma 15. It holds $\|\Psi^0 J^t\|_F^2 \leq 2\|\Delta X^{(0)}\|_F^2 + \frac{3\gamma^2}{p^2}\|\Delta Y^{(0)}\|_F^2$ for all $t \geq 0$, where p is defined in (3).

Proof. Starting from (10) and using Lemma 11 with $\delta = 1 - \lambda_2$

$$\|\Psi^0 J^i\|_F^2 = \left\| \begin{pmatrix} \Delta X^{(0)} \tilde{W}^i - i\gamma \Delta Y^{(0)} \tilde{W}^i \\ \gamma \Delta Y^{(0)} \tilde{W}^i \end{pmatrix}^\top \right\|_F^2 \leq 2\|\Delta X^{(0)}\|_F^2 + \frac{3\gamma^2}{p^2}\|\Delta Y^{(0)}\|_F^2.$$

\square

Lemma 16. For arbitrary set of n vectors $\{\mathbf{a}_i\}_{i=1}^n$, $\mathbf{a}_i \in \mathbb{R}^d$

$$\left\| \sum_{i=1}^n \mathbf{a}_i \right\|^2 \leq n \sum_{i=1}^n \|\mathbf{a}_i\|^2. \quad (16)$$

Lemma 17. For given two vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^d$

$$2\langle \mathbf{a}, \mathbf{b} \rangle \leq \gamma \|\mathbf{a}\|^2 + \gamma^{-1} \|\mathbf{b}\|^2, \quad \forall \gamma > 0. \quad (17)$$

Lemma 18. For given two vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^d$

$$\|\mathbf{a} + \mathbf{b}\|^2 \leq (1 + \alpha) \|\mathbf{a}\|^2 + (1 + \alpha^{-1}) \|\mathbf{b}\|^2, \quad \forall \alpha > 0. \quad (18)$$

This inequality also holds for the sum of two matrices $A, B \in \mathbb{R}^{n \times d}$ in Frobenius norm.

Lemma 19. For $A \in \mathbb{R}^{d \times n}$, $B \in \mathbb{R}^{n \times n}$

$$\|AB\|_F \leq \|A\|_F \|B\|_2. \quad (19)$$

B.2 Convex Cases

Proof of Lemma 7 We first state auxiliary lemma about consensus recursion.

Lemma 20. There exists absolute constants $C_1 = 440$, $C_2 = 380$ such that iterates of Algorithm 1 satisfy,

$$\mathbb{E} \|\Psi_{t+k}\|_F^2 \leq \frac{3}{4} \mathbb{E} \|\Psi_t\|_F^2 + \gamma^2 \frac{C_1 \tau}{c^2} \sum_{j=0}^{k-1} \mathbb{E} \|\nabla f(X^{t+j}) - \nabla f(X^*)\|_F^2 + \gamma^2 \frac{C_2 \tau}{c^2} n \sigma^2. \quad (20)$$

where $\tau \leq k \leq 2\tau$, $\tau = \frac{2}{p} \log \left(\frac{50}{p} (1 + \log \frac{1}{p}) \right) + 1$, p and c are defined in (3), $\Psi_t = (\Delta X^{(t)}, \gamma \Delta Y^{(t)})$ and is defined in (9).

Proof. We start from the recursion (13) given in the main text

$$\Psi_{t+k} = \Psi_t J^k + \gamma \sum_{j=1}^k E_{t+j-1} J^{k-j}.$$

Taking the norm,

$$\|\Psi_{t+k}\|_F^2 \stackrel{(18), \alpha=\frac{1}{4}, (19)}{\leq} \left(1 + \frac{1}{4}\right) \|J^k\|_2^2 \|\Psi_t\|_F^2 + 5\gamma^2 \left\| \sum_{j=1}^k E_{t+j-1} J^{k-j} \right\|_F^2$$

Using the key Lemma 4, the first term can be estimated as

$$\left(1 + \frac{1}{4}\right) \|J^k\|_2^2 \|\Psi_t\|_F^2 \leq \frac{3}{4} \|\Psi_t\|_F^2.$$

Lets estimate separately the second term. Denoting $G^{(t)} = \nabla F(X^{(t)}, \xi^{(t)})$,

$$\begin{aligned} \left\| \sum_{j=1}^k E_{t+j-1} J^{k-j} \right\|_F^2 &= \left\| \begin{pmatrix} -\sum_{j=1}^k (G^{(t+j)} - G^{(t+j-1)}) (k-j) \tilde{W}^{k-j} (I - \frac{\mathbf{1}\mathbf{1}^\top}{n}) \\ \sum_{j=1}^k (G^{(t+j)} - G^{(t+j-1)}) \tilde{W}^{k-j} (I - \frac{\mathbf{1}\mathbf{1}^\top}{n}) \end{pmatrix} \right\|_F^2 \\ &\stackrel{(19)}{\leq} \underbrace{\left\| \sum_{j=1}^k (G^{(t+j)} - G^{(t+j-1)}) (k-j) \tilde{W}^{k-j} \right\|_F^2}_{=: T_1} \\ &\quad + \underbrace{\left\| \sum_{j=1}^k (G^{(t+j)} - G^{(t+j-1)}) \tilde{W}^{k-j} \right\|_F^2}_{=: T_2}, \end{aligned}$$

where we used the definition of the Frobenius norm and $\|I - \frac{\mathbf{1}\mathbf{1}^\top}{n}\| \leq 1$. We now give upper bounds for T_1 and T_2 .

The second term T_2 . We firstly separate the stochastic noise by adding and subtracting the full gradient,

$$\begin{aligned} T_2 &\stackrel{(18)}{\leq} 3 \left\| \sum_{j=1}^k (\nabla f(X^{(t+j)}) - \nabla f(X^{(t+j-1)})) \tilde{W}^{k-j} \right\|_F^2 \\ &\quad + 3 \left\| \sum_{j=1}^k (G^{(t+j)} - \nabla f(X^{(t+j)})) \tilde{W}^{k-j} \right\|_F^2 + 3 \left\| \sum_{j=1}^k (G^{(t+j-1)} - \nabla f(X^{(t+j-1)})) \tilde{W}^{k-j} \right\|_F^2. \end{aligned}$$

Note that

$$\begin{aligned} \mathbb{E} \left\| \sum_{j=1}^k (G^{(t+j)} - \nabla f(X^{(t+j)})) \tilde{W}^{k-j} \right\|_F^2 &= \sum_{j=1}^k \mathbb{E} \left\| (G^{(t+j)} - \nabla f(X^{(t+j)})) \tilde{W}^{k-j} \right\|_F^2 \\ &\leq \sum_{j=1}^k \mathbb{E} \left\| G^{(t+j)} - \nabla f(X^{(t+j)}) \right\|_F^2, \end{aligned}$$

where we used the martingale property $\mathbb{E}_j [G^{(j)} - \nabla f(X^{(j)}) \mid X^{(j)}] = 0$ for all $j \leq t$. It follows

$$\mathbb{E}[T_2] \stackrel{(6)}{\leq} 3 \mathbb{E} \left\| \sum_{j=1}^k (\nabla f(X^{(t+j)}) - \nabla f(X^{(t+j-1)})) \tilde{W}^{k-j} \right\|_F^2 + 6kn\sigma^2.$$

We expand further by adding and subtracting $\nabla f(X^*)$ to the first norm, and bounding stochastic noise by (6) in the other terms

$$\begin{aligned} \mathbb{E}[T_2] &\stackrel{(18),(6)}{\leq} 6 \mathbb{E} \left\| \sum_{j=1}^k (\nabla f(X^{(t+j)}) - \nabla f(X^*)) \tilde{W}^{k-j} \right\|_F^2 + 6 \mathbb{E} \left\| \sum_{j=1}^k (\nabla f(X^{(t+j-1)}) - \nabla f(X^*)) \tilde{W}^{k-j} \right\|_F^2 + 6kn\sigma^2 \\ &\stackrel{(16),(19)}{\leq} 12k \sum_{j=0}^k \mathbb{E} \left\| \nabla f(X^{(t+j)}) - \nabla f(X^*) \right\|_F^2 + 6kn\sigma^2. \end{aligned}$$

The first term T_1 . First, we separate the stochastic noise similarly as above. Defining $Z^{(t)} = G^{(t)} - \nabla f(X^{(t)})$,

$$T_1 \stackrel{(18)}{\leq} 2 \left\| \sum_{j=1}^k [\nabla f(X^{(t+j)}) - \nabla f(X^{(t+j-1)})] (k-j) \tilde{W}^{k-j} \right\|_F^2 + 2 \left\| \sum_{j=1}^k (Z^{(t+j)} - Z^{(t+j-1)}) (k-j) \tilde{W}^{k-j} \right\|_F^2.$$

Next, we add and subtract $\nabla f(X^*)$ in the first term $k-1$ times and temporarily denote $D^{(j)} = \nabla f(X^{(j)}) - \nabla f(X^*)$

$$T_1 \leq 2 \left\| \sum_{j=1}^k (D^{(t+j)} - D^{(t+j-1)}) (k-j) \tilde{W}^{k-j} \right\|_F^2 + 2 \left\| \sum_{j=1}^k (Z^{(t+j)} - Z^{(t+j-1)}) (k-j) \tilde{W}^{k-j} \right\|_F^2.$$

Next, we re-group the sums by the gradient index.

$$\begin{aligned} T_1 &\leq 2 \left\| D^{(t+k-1)} \tilde{W} - (k-1) D^{(t)} \tilde{W}^{k-1} + \sum_{j=1}^{k-2} D^{(t+j)} [(k-j) \tilde{W}^{k-j} - (k-j-1) \tilde{W}^{k-j-1}] \right\|_F^2 \\ &\quad + 2 \left\| Z^{(t+k-1)} \tilde{W} - (k-1) Z^{(t)} \tilde{W}^{k-1} + \sum_{j=1}^{k-2} Z^{(t+j)} [(k-j) \tilde{W}^{k-j} - (k-j-1) \tilde{W}^{k-j-1}] \right\|_F^2 \\ &\stackrel{(16),(19)}{\leq} 2k \left[\left\| D^{(t+k-1)} \right\|_F^2 + \left\| D^{(t)} (k-1) \tilde{W}^{k-1} \right\|_F^2 + \sum_{j=1}^{k-2} \left\| D^{(t+j)} [(k-j) \tilde{W}^{k-j} - (k-j-1) \tilde{W}^{k-j-1}] \right\|_F^2 \right] \\ &\quad + 2 \left[\left\| Z^{(t+k-1)} \right\|_F^2 + \left\| Z^{(t)} (k-1) \tilde{W}^{k-1} \right\|_F^2 + \sum_{j=1}^{k-2} \left\| Z^{(t+j)} [(k-j) \tilde{W}^{k-j} - (k-j-1) \tilde{W}^{k-j-1}] \right\|_F^2 \right] \end{aligned}$$

where for splitting Z we used martingale property $\mathbb{E}_j [G^{(j)} - \nabla f(X^{(j)}) | X^{(j)}] = 0$ for all $j \leq t$.

Next, we use Lemma 13 to estimate the norm $\left\| (k-j) \tilde{W}^{k-j} - (k-j-1) \tilde{W}^{k-j-1} \right\|_2^2 \leq \frac{16}{c^2}$; and using (10) we estimate $\left\| (k-1) \tilde{W}^{k-1} \right\|_2^2 \leq \|J^{k-1}\|_2^2 \leq \frac{1}{2}$ due to our choice of $k \geq \tau$ and a key Lemma 4

$$T_1 \stackrel{(19)}{\leq} \frac{32k}{c^2} \sum_{j=0}^{k-1} \left\| D^{(t+j)} \right\|_F^2 + \frac{32}{c^2} \sum_{j=0}^{k-1} \left\| Z^{(t+j)} \right\|_F^2$$

Taking expectation over the stochastic noise,

$$\mathbb{E}[T_1] \stackrel{(6)}{\leq} \frac{32k}{c^2} \sum_{j=0}^{k-1} \left\| D^{(t+j)} \right\|_F^2 + \frac{32kn\sigma^2}{c^2}$$

Summing up T_1 and T_2 and estimating $k \leq 2\tau$ we conclude the proof

$$\mathbb{E} \|\Psi_{t+k}\|_F^2 \leq \frac{3}{4} \mathbb{E} \|\Psi_t\|_F^2 + \gamma^2 \frac{440\tau}{c^2} \sum_{j=0}^k \mathbb{E} \left\| \nabla f(X^{(t+j)}) - \nabla f(X^*) \right\|_F^2 + \gamma^2 \frac{380\tau}{c^2} n\sigma^2. \quad \square$$

We will proof Lemma 7 with $B_1 = 28C_1$, $B_2 = 4C_2$, $B_3 = \sqrt{515 \cdot 2C_1}$, where $C_1 = 220$ and $C_2 = 190$ are constants from Lemma 20.

Proof of Lemma 7. Observe, for any t ,

$$\begin{aligned} \left\| \nabla f(X^{(t)}) - \nabla f(X^*) \right\|_F^2 &\stackrel{(18)}{\leq} 2 \left\| \nabla f(X^{(t)}) - \nabla f(\bar{X}^{(t)}) \right\|_F^2 + 2 \left\| \nabla f(\bar{X}^{(t)}) - \nabla f(X^*) \right\|_F^2 \\ &\stackrel{(4)}{\leq} 2L^2 \underbrace{\left\| X^{(t)} - \bar{X}^{(t)} \right\|_F^2}_{\leq \|\Psi_t\|_F^2} + 2 \left\| \nabla f(\bar{X}^{(t)}) - \nabla f(X^*) \right\|_F^2 \end{aligned}$$

With Lemma 20

$$\begin{aligned} \mathbb{E} \|\Psi_{t+k}\|_F^2 &\stackrel{(20)}{\leq} \frac{3}{4} \mathbb{E} \|\Psi_t\|_F^2 + \gamma^2 \frac{\tau C_1}{c^2} \sum_{j=0}^k \mathbb{E} \left\| \nabla f(X^{(t+j)}) - \nabla f(X^*) \right\|_F^2 + \gamma^2 \frac{\tau C_2}{c^2} n \sigma^2 \\ &\leq \frac{3}{4} \mathbb{E} \|\Psi_t\|_F^2 + \gamma^2 \frac{2C_1 \tau L^2}{c^2} \sum_{j=0}^k \mathbb{E} \|\Psi_{t+j}\|_F^2 + \gamma^2 \frac{2C_1 \tau}{c^2} \sum_{j=0}^k \mathbb{E} \left\| \nabla f(\bar{X}^{(t+j)}) - \nabla f(X^*) \right\|_F^2 + \gamma^2 \frac{\tau C_2}{c^2} n \sigma^2 \\ &\stackrel{\gamma < \frac{c}{\sqrt{512 \cdot 2C_1 L \tau}}}{\leq} \frac{3}{4} \mathbb{E} \|\Psi_t\|_F^2 + \frac{1}{512\tau} \sum_{j=0}^k \mathbb{E} \|\Psi_{t+j}\|_F^2 + \gamma^2 \frac{2C_1 \tau}{c^2} \sum_{j=0}^k \mathbb{E} \left\| \nabla f(\bar{X}^{(t+j)}) - \nabla f(X^*) \right\|_F^2 + \gamma^2 \frac{\tau C_2}{c^2} n \sigma^2 \end{aligned}$$

Next, we estimate the third term by smoothness for $j < k$

$$\left\| \nabla f(\bar{X}^{(t+j)}) - \nabla f(X^*) \right\|_F^2 \leq 2Ln \left(f(\bar{\mathbf{x}}^{(t+j)}) - f(\mathbf{x}^*) \right).$$

And for $j = k$, the index is $t + k$ and it should appear only in LHS. Thus we estimate

$$\begin{aligned} \left\| \nabla f(\bar{X}^{(t+k)}) - \nabla f(X^*) \right\|_F^2 &\stackrel{(18)}{\leq} 2 \left\| \nabla f(\bar{X}^{(t+k)}) - \nabla f(\bar{X}^{(t+k-1)}) \right\|_F^2 + 2 \left\| \nabla f(\bar{X}^{(t+k-1)}) - \nabla f(X^*) \right\|_F^2 \\ &\stackrel{(4)}{\leq} 2L^2 \left\| \bar{X}^{(t+k)} - \bar{X}^{(t+k-1)} \right\|_F^2 + 4Ln \left(f(\bar{\mathbf{x}}^{(t+k-1)}) - f(\mathbf{x}^*) \right) \end{aligned}$$

Next we use (11), that is equivalent to $\bar{X}^{(t+k)} = \bar{X}^{(t+k-1)} - \gamma \nabla F(X^{(t+k-1)}, \xi^{(t+k-1)}) \frac{\mathbf{1}\mathbf{1}^\top}{n}$. Taking expectation

$$\begin{aligned} \mathbb{E} \left\| \bar{X}^{(t+k)} - \bar{X}^{(t+k-1)} \right\|_F^2 &\stackrel{(11),(6)}{\leq} \gamma^2 \left\| \nabla f(X^{(t+k-1)}) \frac{\mathbf{1}\mathbf{1}^\top}{n} \right\|_F^2 + \gamma^2 \sigma^2 \\ &\leq 2\gamma^2 \left\| \nabla f(X^{(t+k-1)}) \frac{\mathbf{1}\mathbf{1}^\top}{n} - \nabla \bar{f}(\bar{X}^{(t+k-1)}) \right\|_F^2 + 2\gamma^2 \left\| \nabla \bar{f}(\bar{X}^{(t+k-1)}) - \nabla \bar{f}(X^*) \right\|_F^2 + \gamma^2 \sigma^2 \\ &\stackrel{(18),(4)}{\leq} 2\gamma^2 L^2 \left\| X^{(t+k-1)} - \bar{X}^{(t+k-1)} \right\|_F^2 + 4\gamma^2 Ln \left(f(\bar{\mathbf{x}}^{(t+k-1)}) - f(\mathbf{x}^*) \right) + \gamma^2 \sigma^2 \end{aligned}$$

where on the second line we used $\nabla f(\bar{X}) \frac{\mathbf{1}\mathbf{1}^\top}{n} = \nabla \bar{f}(\bar{X})$, and $\nabla \bar{f}(X^*) = 0$. As $\gamma \leq \frac{c}{\sqrt{512 \cdot 2C_1 L \tau}}$

$$\left\| \nabla f(\bar{X}^{(t+k)}) - \nabla f(X^*) \right\|_F^2 \leq L^2 \|\Psi_{t+k-1}\|_F^2 + 5Ln \left(f(\bar{\mathbf{x}}^{(t+k-1)}) - f(\mathbf{x}^*) \right)$$

Coming back to recursion for $\|\Psi_{t+k}\|_F^2$ and using that $\frac{2C_1 L^2 \tau}{c^2} \gamma^2 \leq \frac{1}{512\tau}$ by our choice of γ ,

$$\mathbb{E} \|\Psi_{t+k}\|_F^2 \leq \frac{3}{4} \mathbb{E} \|\Psi_t\|_F^2 + \frac{1}{256\tau} \sum_{j=0}^k \mathbb{E} \|\Psi_{t+j}\|_F^2 + \gamma^2 \frac{C_1 \tau}{c^2} 14Ln \sum_{j=0}^{k-1} \mathbb{E} \left(f(\bar{\mathbf{x}}^{(t+j)}) - f(\mathbf{x}^*) \right) + \gamma^2 \frac{2C_2 \tau}{c^2} n \sigma^2$$

It is only left to get rid of $\|\Psi_{t+k}\|_F^2$ from RHS. For that we move the term with $\|\Psi_{t+k}\|_F^2$ to LHS and divide the whole equation by $(1 - \frac{1}{256\tau})$. We use that $(1 - \frac{1}{256\tau})^{-1} \leq 1 + \frac{1}{128\tau} \leq 1 + \frac{1}{256} < 2$, and that $(1 - \frac{1}{4})(1 + \frac{1}{128\tau}) \leq (1 - \frac{1}{4})(1 + \frac{1}{128}) \leq (1 - \frac{1}{8})$. We thus arrive to the Lemma's statement

$$\mathbb{E} \|\Psi_{t+k}\|_F^2 \leq \frac{7}{8} \mathbb{E} \|\Psi_t\|_F^2 + \frac{1}{128\tau} \sum_{j=0}^{k-1} \mathbb{E} \|\Psi_{t+j}\|_F^2 + \gamma^2 \frac{28C_1 \tau}{c^2} Ln \sum_{j=0}^{k-1} \mathbb{E} \left(f(\bar{\mathbf{x}}^{(t+j)}) - f(\mathbf{x}^*) \right) + \gamma^2 \frac{4C_2 \tau}{c^2} n \sigma^2$$

□

Proof of Lemma 8.

Proof. Define $\alpha = 28C_1 \frac{\tau}{c^2} Ln$, $\beta = 4C_2 \frac{\tau}{c^2} \sigma^2 n$ for simplicity. Then inequality (14) takes the form

$$\mathbb{E} \|\Psi_{t+k}\|_F^2 \leq \left(1 - \frac{1}{8}\right) \mathbb{E} \|\Psi_t\|_F^2 + \frac{1}{128\tau} \sum_{j=0}^{k-1} \mathbb{E} \|\Psi_{t+j}\|_F^2 + \alpha\gamma^2 \sum_{j=0}^{k-1} \mathbb{E} e_{t+j} + \beta\gamma^2 \quad (21)$$

A new quantity. We define a new quantity that has non-increasing properties even for $k < \tau$ in contrast to $\mathbb{E} \|\Psi_{t+k}\|_F^2$. For $t \geq 0$ we define

$$\Phi_{t+\tau} := \frac{1}{\tau} \sum_{j=0}^{\tau-1} \mathbb{E} \|\Psi_{t+j}\|_F^2 \quad E_{t+\tau} := \alpha \sum_{j=0}^{\tau-1} \mathbb{E} e_{t+j}$$

Non-increasing property for $k < \tau$ (but $t+k \geq \tau$).

$$\Phi_{t+k} = \frac{1}{\tau} \left(\sum_{i=k}^{\tau-1} \mathbb{E} \|\Psi_{t-\tau+i}\|_F^2 + \sum_{i=0}^{k-1} \mathbb{E} \|\Psi_{t+i}\|_F^2 \right)$$

Applying (21) to the second sum,

$$\begin{aligned} \Phi_{t+k} &\leq \frac{1}{\tau} \sum_{i=k}^{\tau-1} \mathbb{E} \|\Psi_{t-\tau+i}\|_F^2 + \frac{1}{\tau} \sum_{i=0}^{k-1} \left[\left(1 - \frac{1}{8}\right) \mathbb{E} \|\Psi_{t-\tau+i}\|_F^2 + \frac{1}{128} \Phi_{t+i} + \gamma^2 E_{t+i} + \beta\gamma^2 \right] \\ \Phi_{t+k} &\leq \Phi_t + \frac{1}{128\tau} \sum_{i=0}^{k-1} \Phi_{t+i} + \frac{1}{\tau} \gamma^2 \sum_{i=0}^{k-1} E_{t+i} + \frac{k}{\tau} \beta\gamma^2, \end{aligned} \quad (22)$$

where we used that $\Theta_t \geq 0 \forall t$ and that $\tau \geq k$.

Contraction property for $\tau \leq k \leq 2\tau$. Using (21) and a definition of Φ_{t+k} ,

$$\Phi_{t+k} = \frac{1}{\tau} \sum_{j=k-\tau}^{k-1} \mathbb{E} \|\Psi_{t+j}\|_F^2 \leq \left(1 - \frac{1}{8}\right) \underbrace{\frac{1}{\tau} \sum_{j=k-\tau}^{k-1} \mathbb{E} \|\Psi_{t+j-\tau}\|_F^2}_{\Phi_{t+k-\tau}} + \frac{1}{128\tau} \sum_{j=k-\tau}^{k-1} \Phi_{t+j} + \gamma^2 \frac{1}{\tau} \sum_{i=k-\tau}^{k-1} E_{t+i} + \beta\gamma^2$$

Combining with (22) we get contraction for Φ_{t+k}

$$\Phi_{t+k} \stackrel{(22)}{\leq} \left(1 - \frac{1}{8}\right) \Phi_t + \frac{1}{128\tau} \sum_{j=0}^{k-1} \Phi_{t+j} + \gamma^2 \frac{1}{\tau} \sum_{i=0}^{k-1} E_{t+i} + 2\beta\gamma^2 \quad (23)$$

Simplifying contraction property. First, we substitute (22) into the second term of (23)

$$\begin{aligned} \Phi_{t+k} &\leq \left(1 - \frac{1}{8}\right) \Phi_t + \frac{1}{128\tau} \sum_{i=0}^{k-2} \Phi_{t+i} + \frac{1}{128\tau} \left[\Phi_t + \frac{1}{128\tau} \sum_{i=0}^{k-2} \Phi_{t+i} + \gamma^2 \frac{1}{\tau} \sum_{i=0}^{k-2} E_{t+i} + 2\beta\gamma^2 \right] + \gamma^2 \frac{1}{\tau} \sum_{i=0}^{k-1} E_{t+i} + 2\beta\gamma^2 \\ &\leq \left(1 - \frac{1}{8}\right) \left(1 + \frac{1}{64\tau}\right) \Phi_t + \left(1 + \frac{1}{128\tau}\right) \left[\frac{1}{128\tau} \sum_{i=0}^{k-2} \Phi_{t+i} + \gamma^2 \frac{1}{\tau} \sum_{i=0}^{k-2} E_{t+i} + 2\beta\gamma^2 \right] + \gamma^2 \frac{1}{\tau} E_{t+k-1} \end{aligned}$$

where we used that $\frac{1}{128\tau} = (1 - \frac{1}{2}) \frac{1}{64\tau} \leq (1 - \frac{1}{8}) \frac{1}{64\tau}$. Similarly applying (22) to the rest of Φ_{t+i} ,

$$\Phi_{t+k} \leq \left(1 - \frac{1}{8}\right) \left(1 + \frac{1}{64\tau}\right)^k \Phi_t + \gamma^2 \frac{1}{\tau} \sum_{i=0}^{k-1} \left(1 + \frac{1}{128\tau}\right)^{t+k-1-i} E_{t+i} + \left(1 + \frac{1}{128\tau}\right)^k 2\beta\gamma^2$$

We further use $(1 + \frac{1}{64\tau})^k \leq (1 + \frac{1}{64\tau})^{2\tau} \leq \exp(\frac{1}{32}) \leq 1 + \frac{1}{16}$ and $(1 - \frac{1}{8}) (1 + \frac{1}{64\tau})^k \leq (1 - \frac{1}{16})$; and that $(1 + \frac{1}{128\tau})^k \leq 1 + \frac{1}{32} \leq 2$. Therefore,

$$\Phi_{t+k} \leq \left(1 - \frac{1}{16}\right) \Phi_t + 2\gamma^2 \frac{1}{\tau} \sum_{i=0}^{k-1} E_{t+i} + 4\beta\gamma^2 \quad (24)$$

Simplifying non-increasing property (22). Similarly as above we substitute recursively (22) into the second term of (22), for $0 < k < \tau$

$$\begin{aligned}\Phi_{t+k} &\leq \left(1 + \frac{1}{128\tau}\right) \Phi_t + \left(1 + \frac{1}{128\tau}\right) \left[\frac{1}{128\tau} \sum_{i=0}^{k-2} \Phi_{t+i} + \gamma^2 \frac{1}{\tau} \sum_{i=0}^{k-2} E_{t+i} + \beta\gamma^2 \right] + \gamma^2 \frac{1}{\tau} E_{t+\tau-1} \\ &\leq \left(1 + \frac{1}{128\tau}\right)^\tau \Phi_t + \gamma^2 \frac{1}{\tau} \sum_{i=0}^{\tau-1} \left(1 + \frac{1}{128\tau}\right)^{t+\tau-1-i} E_{t+i} + \left(1 + \frac{1}{128\tau}\right)^\tau \beta\gamma^2\end{aligned}$$

Using now that $\left(1 + \frac{1}{128\tau}\right)^\tau \leq 2$ we get

$$\Phi_{t+k} \leq 2\Phi_t + 2\gamma^2 \frac{1}{\tau} \sum_{i=0}^{\tau-1} E_{t+i} + 2\beta\gamma^2 \quad (25)$$

Obtaining recursion for $\mathbb{E} \|\Psi_t\|_F^2 + \Phi_t$. As our final goal is to obtain inequality for $\mathbb{E} \|\Psi_t\|_F^2$, we start modifying (21), for $\tau \leq k \leq 2\tau$

$$\begin{aligned}\mathbb{E} \|\Psi_{t+k}\|_F^2 &\leq \left(1 - \frac{1}{8}\right) \mathbb{E} \|\Psi_t\|_F^2 + \frac{1}{128} (\Phi_{t+k} + \Phi_{t+\tau}) + \alpha\gamma^2 \sum_{j=0}^{k-1} e_{t+j} + \beta\gamma^2 \\ &\stackrel{(25)}{\leq} \left(1 - \frac{1}{8}\right) \mathbb{E} \|\Psi_t\|_F^2 + \frac{1}{128} \left[4\Phi_t + 4\gamma^2 \frac{1}{\tau} \sum_{j=0}^{k-1} E_{t+j} \right] + \alpha\gamma^2 \sum_{j=0}^{k-1} e_{t+j} + 2\beta\gamma^2 \\ &\leq \left(1 - \frac{1}{8}\right) \mathbb{E} \|\Psi_t\|_F^2 + \frac{1}{32} \Phi_t + \frac{\gamma^2}{32} \frac{1}{\tau} \sum_{j=0}^{k-1} E_{t+j} + \alpha\gamma^2 \sum_{j=0}^{k-1} e_{t+j} + 2\beta\gamma^2\end{aligned}$$

Summing up the last inequality and (24) we get

$$\mathbb{E} \|\Psi_{t+k}\|_F^2 + \Phi_{t+k} \leq \left(1 - \frac{1}{32}\right) \left[\mathbb{E} \|\Psi_t\|_F^2 + \Phi_t \right] + 3\gamma^2 \frac{1}{\tau} \sum_{j=0}^{k-1} E_{t+j} + \gamma^2 \alpha \sum_{j=0}^{k-1} e_{t+j} + 6\beta\gamma^2$$

Unrolling recursion up to τ . For a given $t \geq \tau$, lets define $m = \lfloor t/\tau \rfloor - 1$. Then

$$\mathbb{E} \|\Psi_t\|_F^2 + \Phi_t \leq \left(1 - \frac{1}{32}\right) \left[\mathbb{E} \|\Psi_{m\tau}\|_F^2 + \Phi_{m\tau} \right] + 3\gamma^2 \frac{1}{\tau} \sum_{j=m\tau}^{t-1} E_j + \gamma^2 \alpha \sum_{j=m\tau}^{t-1} e_j + 6\beta\gamma^2$$

Unrolling this recursively up to τ we get,

$$\mathbb{E} \|\Psi_t\|_F^2 + \Phi_t \leq \left(1 - \frac{1}{32}\right)^m \left[\mathbb{E} \|\Psi_\tau\|_F^2 + \Phi_\tau \right] + \gamma^2 \sum_{j=\tau}^{t-1} \left(1 - \frac{1}{32}\right)^{\lfloor (t-j)/\tau \rfloor} \left[3\frac{1}{\tau} E_j + \alpha e_j \right] + 6\beta\gamma^2 \sum_{j=0}^{m-1} \left(1 - \frac{1}{32}\right)^j \quad (26)$$

Initial conditions. Inequality above work for $t \geq \tau$. Here, we focus on $t < \tau$. Using similar calculations as in Lemma 7 replacing estimation of $\|\Psi^0 J^t\|_F^2$ by Lemma 15, we get that

$$\mathbb{E} \|\Psi_t\|_F^2 \leq 2 \underbrace{\left\| \Delta X^{(0)} \right\|_F^2 + \frac{3\gamma^2}{p^2} \left\| \Delta Y^{(0)} \right\|_F^2}_{:=\tilde{\Theta}_0} + \frac{1}{128\tau} \sum_{j=0}^{t-1} \mathbb{E} \|\Psi_j\|_F^2 + \alpha\gamma^2 \sum_{j=0}^{t-1} e_j + \beta\gamma^2 \quad (27)$$

Recursively applying (27) to the second term of (27), similarly as above, we get

$$\mathbb{E} \|\Psi_t\|_F^2 \leq 2\tilde{\Theta}_0 + 2\alpha\gamma^2 \frac{1}{\tau} \sum_{j=0}^{t-1} e_j + 2\beta\gamma^2 \quad (28)$$

And therefore,

$$\Phi_\tau = \frac{1}{\tau} \sum_{j=0}^{\tau-1} \mathbb{E} \|\Psi_j\|_F^2 \leq 2\tilde{\Theta}_0 + 2\alpha\gamma^2 \frac{1}{\tau} \sum_{j=0}^{\tau-1} e_j + 2\beta\gamma^2 \quad (29)$$

Final recursion. Finally we apply (28), (29) to the first term of (26)

$$\mathbb{E} \|\Psi_t\|_F^2 + \Phi_t \leq \left(1 - \frac{1}{32}\right)^m 4\tilde{\Theta}_0 + \gamma^2 \sum_{j=\tau}^{t-1} \left(1 - \frac{1}{32}\right)^{\lfloor (t-j)/\tau \rfloor} \left[3\frac{1}{\tau} E_j + 5\alpha e_j \right] + 10\beta\gamma^2 \sum_{j=0}^{m-1} \left(1 - \frac{1}{32}\right)^j$$

- For the last term we estimate $\sum_{j=0}^{m-1} \left(1 - \frac{1}{32}\right)^j \leq 2$.
- For the terms with e_j and E_j we estimate, similar to [16],

$$\begin{aligned} \left(1 - \frac{1}{32}\right)^{1/\tau} &\leq \exp\left(-\frac{1}{32\tau}\right) \leq 1 - \frac{1}{64\tau} && \text{and thus} \\ \left(1 - \frac{1}{32}\right)^{\lfloor (t-j)/\tau \rfloor} &\leq \left(1 - \frac{1}{64\tau}\right)^{\tau \lfloor (t-j)/\tau \rfloor} \leq \left(1 - \frac{1}{64\tau}\right)^{t-j} \left(1 - \frac{1}{64\tau}\right)^{-\tau} \leq 2 \left(1 - \frac{1}{64\tau}\right)^{t-j} \end{aligned}$$

where as $\frac{1}{64\tau} \leq \frac{1}{2}$ we estimated $\left(1 - \frac{1}{64\tau}\right)^{-\tau} \leq \left(\frac{1}{1 - \frac{1}{64\tau}}\right)^\tau \leq \left(1 + \frac{1}{32\tau}\right)^\tau \leq \exp\left(\frac{1}{32}\right) < 2$.

- Similarly, for $\tilde{\Theta}_0$ term we estimate $\left(1 - \frac{1}{32}\right)^m = \left(1 - \frac{1}{32}\right)^{\lfloor \frac{t-\tau}{\tau} \rfloor} \leq \left(1 - \frac{1}{64\tau}\right)^t \left(1 - \frac{1}{64\tau}\right)^{-2\tau} \leq 2 \left(1 - \frac{1}{64\tau}\right)^t$.
- For the terms with E_j we additionally estimate

$$2 \left(1 - \frac{1}{64\tau}\right)^{t-j} E_j = 2 \left(1 - \frac{1}{64\tau}\right)^{t-j} \sum_{i=j-\tau}^{j-1} e_i = 2 \sum_{i=j-\tau}^{j-1} \left(1 - \frac{1}{64\tau}\right)^{t-i} \left(1 - \frac{1}{64\tau}\right)^{i-j} e_i$$

Further, $-\tau < i - j < -1$, and thus $\left(1 - \frac{1}{64\tau}\right)^{i-j} \leq 2$ for all such $-\tau < i - j < -1$.

Therefore we obtain

$$\mathbb{E} \|\Psi_t\|_F^2 + \Phi_t \leq \left(1 - \frac{1}{64\tau}\right)^t 8\tilde{\Theta}_0 + 22\gamma^2 \alpha \sum_{j=0}^{t-1} \left(1 - \frac{1}{64\tau}\right)^{t-j} e_j + 20\beta\gamma^2$$

This brings us to the statement of the lemma. □

The rest of the proof follows closely [16].

B.2.1 τ -slow Sequences

Definition 2 (τ -slow sequences [44]). *The sequence $\{a_t\}_{t \geq 0}$ of positive values is τ -slow decreasing for parameter $\tau > 0$ if*

$$a_{t+1} \leq a_t, \quad \forall t \geq 0 \quad \text{and,} \quad a_{t+1} \left(1 + \frac{1}{2\tau}\right) \geq a_t, \quad \forall t \geq 0.$$

The sequence $\{a_t\}_{t \geq 0}$ is τ -slow increasing if $\{a_t^{-1}\}_{t \geq 0}$ is τ -slow decreasing.

Proposition 21 (Examples).

1. The sequence $\{\eta_t^2\}_{t \geq 0}$ with $\eta_t = \frac{a}{b+t}$, $b \geq 32\tau$ is 4τ -slow decreasing.
2. The sequence of constant stepsizes $\{\eta_t^2\}_{t \geq 0}$ with $\eta_t = \eta$ is τ -slow decreasing for any τ .
3. The sequence $\{w_t\}_{t \geq 0}$ with $w_t = (b+t)^2$, $b \geq 84\tau$ is 8τ -slow increasing.
4. The sequence of constant weights $\{w_t\}_{t \geq 0}$ with $w_t = 1$ is τ -slow increasing for any τ .

B.2.2 The Main Recursion

Lemma 22 (The main recursion). *Let $\{w_t\}_{t \geq 0}$ be 64τ -slow increasing sequence, $W_t = \frac{1}{T+1} \sum_{t=0}^T w_t$, with $\gamma \leq \frac{c}{582C_1\tau L}$ it holds that*

$$\sum_{t=0}^T w_t \mathbb{E} \|\Psi_t\|_F^2 \leq \sum_{t=0}^T w_t \left(1 - \frac{1}{64\tau}\right)^t 8\tilde{\Theta}_0 + \frac{n}{6L} \sum_{t=0}^T e_t w_t + 40C_2 \frac{\tau}{c^2} \sigma^2 n \gamma^2 W_T, \quad (30)$$

where $e_t = f(\bar{\mathbf{x}}^{(t)}) - f^*$, $\tilde{\Theta}_0 = 2 \|\Delta X^{(0)}\|_F^2 + \frac{3\gamma^2}{p^2} \|\Delta Y^{(0)}\|_F^2$, $C_1 = 440$, $C_2 = 380$.

Proof. We start by averaging (15) with weights w_t . Define $W_T = \sum_{t=0}^T w_t$, $\alpha = 28C_1 \frac{\tau}{c^2} Ln$, $\beta = 4C_2 \frac{\tau}{c^2} \sigma^2 n$,

$$\sum_{t=0}^T w_t \mathbb{E} \|\Psi_t\|_F^2 \leq \sum_{t=0}^T w_t \left(1 - \frac{1}{64\tau}\right)^t 8\tilde{\Theta}_0 + 22\gamma^2 \alpha \underbrace{\sum_{t=0}^T w_t \sum_{j=0}^{t-1} \left(1 - \frac{1}{64\tau}\right)^{t-j} e_j}_{:=T_1} + 20\beta\gamma^2 W_T$$

For the middle term T_1 we use that w_t are 64τ -slow increasing sequences, i.e. $w_t \leq w_j \left(1 + \frac{1}{128\tau}\right)^{t-j}$, we get

$$\begin{aligned} T_1 &= \sum_{t=0}^T \sum_{j=0}^{t-1} \left(1 - \frac{1}{64\tau}\right)^{t-j} \left(1 + \frac{1}{128\tau}\right)^{t-j} e_j w_j \leq \sum_{t=0}^T \sum_{j=0}^{t-1} \left(1 - \frac{1}{128\tau}\right)^{t-j} e_j w_j \\ &\leq \sum_{j=0}^T e_j w_j \sum_{t=j+1}^T \left(1 - \frac{1}{128\tau}\right)^{t-j} \leq \sum_{j=0}^T e_j w_j \sum_{t=0}^{\infty} \left(1 - \frac{1}{128\tau}\right)^{t-j} \leq 128\tau \sum_{t=0}^T e_t w_t \end{aligned}$$

Therefore,

$$\sum_{t=0}^T w_t \Theta_t \leq \sum_{t=0}^T w_t \left(1 - \frac{1}{64\tau}\right)^t 8\tilde{\Theta}_0 + 2816\gamma^2 \alpha \tau \sum_{t=0}^T e_t w_t + 20\beta\gamma^2 W_T$$

Now using that $\gamma \leq \frac{c}{582C_1\tau L}$ and that $\alpha = 20C_1 \frac{\tau}{c^2} Ln$, $\beta = 2C_2 \frac{\tau}{c^2} \sigma^2 n$.

$$\sum_{t=0}^T w_t \Theta_t \leq \sum_{t=0}^T w_t \left(1 - \frac{1}{64\tau}\right)^t 8\tilde{\Theta}_0 + \frac{n}{6L} \sum_{t=0}^T e_t w_t + 40C_2 \frac{\tau}{c^2} \sigma^2 n \gamma^2 W_T$$

□

B.2.3 Combining with the Descent Lemma 6

Lemma 23. *Define $D = \frac{\sigma^2}{n}$, $a = \frac{\mu}{2}$, $A = 24L \frac{1}{n} \tilde{\Theta}_0$, $\tilde{\Theta}_0 = 2 \|\Delta X^{(0)}\|_F^2 + \frac{3\gamma^2}{p^2} \|\Delta Y^{(0)}\|_F^2$, $B = 120C_2 L \frac{\tau}{c^2} \sigma^2$, $C_1 = 440$, $C_2 = 380$. Then with $\gamma \leq \frac{c}{582C_1\tau L}$ it holds that*

$$\frac{1}{2W_T} \sum_{t=0}^T w_t e_t \leq \frac{1}{W_T} \sum_{t=0}^T \left(\frac{(1-\gamma a)}{\gamma} w_t r_t - \frac{w_t}{\gamma} r_{t+1} \right) + D\gamma + \frac{A}{W_T} \sum_{t=0}^T w_t \left(1 - \frac{1}{64\tau}\right)^t + B\gamma^2 \quad (31)$$

Proof. First, define $W_T = \sum_{t=0}^T w_t$, $r_t = \|\bar{\mathbf{x}}^{(t)} - \mathbf{x}^*\|^2$, $\Theta_t = \sum_{i=1}^n \mathbb{E} \|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2$. In this notation, (12) writes as

$$r_{t+1} \leq \left(1 - \frac{\gamma\mu}{2}\right) r_t + \frac{\gamma^2\sigma^2}{n} - \gamma e_t + \gamma \frac{3L}{n} \Theta_t,$$

We rearrange (12) by multiplying by w_t and dividing by γ

$$w_t e_t \leq \frac{(1 - \frac{\gamma\mu}{2})}{\gamma} w_t r_t - \frac{w_t}{\gamma} r_{t+1} + \frac{\sigma^2}{n} w_t \gamma + \frac{3L}{n} w_t \Theta_t,$$

Now summing up, dividing by W_T , using that $\Theta_t \leq \mathbb{E} \|\Psi_t\|_F^2$, and using (30)

$$\begin{aligned} \frac{1}{W_T} \sum_{t=0}^T w_t e_t &\leq \frac{1}{W_T} \sum_{t=0}^T \left(\frac{(1 - \frac{\gamma\mu}{2})}{\gamma} w_t r_t - \frac{w_t}{\gamma} r_{t+1} \right) + \frac{\sigma^2}{n} \gamma + \frac{1}{W_T} \sum_{t=0}^T w_t 24L \left(1 - \frac{1}{64\tau}\right)^t \frac{1}{n} \tilde{\Theta}_0 \\ &\quad + \frac{1}{2} \frac{1}{W_T} \sum_{t=0}^T e_t w_t + 120C_2 L \frac{\tau}{c^2} \sigma^2 \gamma^2 \end{aligned}$$

Putting the fourth term to LHS we get the statement of the lemma. \square

Now similar to [16, Lemma 15] we obtain the rates of Theorem 2 for the strongly convex case, and similar to [16, Lemma 16] for the weakly convex case.

B.2.4 Strongly Convex Case

Lemma 24. *If non-negative sequences $\{r_t\}_{t \geq 0}, \{e_t\}_{t \geq 0}$ satisfy (31) for some constants $a > 0, D, A, B \geq 0$, then there exists a constant stepsize $\gamma < \frac{1}{b}$ with $b \geq 128a\tau$ such that for weights $w_t = (1 - a\gamma)^{-(t+1)}$ and $W_T := \sum_{t=0}^T w_t$ it holds:*

$$\frac{1}{2W_T} \sum_{t=0}^T e_t w_t + ar_{T+1} \leq \tilde{\mathcal{O}} \left((r_0 + A/2a)b \exp \left[-\frac{a(T+1)}{b} \right] + \frac{D}{aT} + \frac{B}{a^2 T^2} \right),$$

where $\tilde{\mathcal{O}}$ hides polylogarithmic factors.

Proof. Starting from (31) and using that that $\frac{w_t(1-a\gamma)}{\gamma} = \frac{w_{t-1}}{\gamma}$ we obtain a telescoping sum,

$$\frac{1}{2W_T} \sum_{t=0}^T w_t e_t \leq \frac{1}{W_T \gamma} ((1 - a\gamma)w_0 r_0 - w_T r_{T+1}) + D\gamma + B\gamma^2 + \frac{A}{W_T} \sum_{t=0}^T w_t \left(1 - \frac{1}{64\tau}\right)^t,$$

And hence,

$$\frac{1}{2W_T} \sum_{t=0}^T w_t e_t + \frac{w_T r_{T+1}}{W_T \gamma} \leq \frac{r_0}{W_T \gamma} + D\gamma + B\gamma^2 + \frac{A}{W_T} \sum_{t=0}^T w_t \left(1 - \frac{1}{64\tau}\right)^t,$$

Now we estimate the last term. We use that $2\gamma a \leq \frac{1}{64\tau}$ and thus $(1 - \frac{1}{64\tau})^t \leq (1 - a\gamma)^{2t}$

$$\frac{1}{W_T} \sum_{t=0}^T (1 - a\gamma)^{-(t+1)} \left(1 - \frac{1}{64\tau}\right)^t \leq \frac{1}{W_T} \sum_{t=0}^T (1 - a\gamma)^{t-1} \leq \frac{1}{W_T} \frac{1}{2a\gamma}$$

where we used that $\frac{1}{1-a\gamma} \leq \frac{1}{2}$. Thus,

$$\frac{1}{2W_T} \sum_{t=0}^T w_t e_t + \frac{w_T r_{T+1}}{W_T \gamma} \leq \frac{1}{W_T \gamma} \left(r_0 + \frac{A}{2a} \right) + D\gamma + B\gamma^2,$$

Using that $W_T \leq \frac{w_T}{a\gamma}$ and $W_T \geq w_T = (1 - a\gamma)^{-(T+1)}$ we can simplify

$$\frac{1}{2W_T} \sum_{t=0}^T w_t e_t + ar_{T+1} \leq (1 - a\gamma)^{T+1} \frac{1}{\gamma} \left(r_0 + \frac{A}{2a} \right) + D\gamma + B\gamma^2 \leq \frac{1}{\gamma} \left(r_0 + \frac{A}{2a} \right) \exp[-a\gamma(T+1)] + D\gamma + B\gamma^2,$$

Now lemma follows by tuning γ the same way as in [43].

- If $\frac{1}{b} \geq \frac{\ln(\max\{2, a^2(r_0 + \frac{A}{2a})T^2/D\})}{aT}$ then we choose $\eta = \frac{\ln(\max\{2, a^2(r_0 + \frac{A}{2a})T^2/D\})}{aT}$ and get that

$$\begin{aligned} &\tilde{\mathcal{O}} \left(a(r_0 + A/2a)T \exp \left[-\ln(\max\{2, a^2(r_0 + A/2a)T^2/D\}) \right] \right) + \tilde{\mathcal{O}} \left(\frac{D}{aT} \right) + \tilde{\mathcal{O}} \left(\frac{B}{a^2 T^2} \right) \\ &= \tilde{\mathcal{O}} \left(\frac{D}{aT} \right) + \tilde{\mathcal{O}} \left(\frac{B}{a^2 T^2} \right), \end{aligned}$$

- Otherwise $\frac{1}{b} \leq \frac{\ln(\max\{2, a^2(r_0 + \frac{A}{2a})T^2/D\})}{aT}$ we pick $\eta = \frac{1}{b}$ and get that

$$\begin{aligned} & \tilde{\mathcal{O}} \left((r_0 + A/2a)b \exp \left[-\frac{a(T+1)}{b} \right] + \frac{D}{b} + \frac{B}{b^2} \right) \\ & \leq \tilde{\mathcal{O}} \left((r_0 + A/2a)b \exp \left[-\frac{a(T+1)}{b} \right] + \frac{D}{aT} + \frac{B}{a^2T^2} \right). \quad \square \end{aligned}$$

B.3 Weakly Convex and Non Convex Cases

Lemma 25. *If non-negative sequences $\{r_t\}_{t \geq 0}, \{e_t\}_{t \geq 0}$ satisfy (31) with $a = 0, D, A, B \geq 0$, then there exists a constant stepsize $\gamma < \frac{1}{b}$ with $b \geq 128a\tau$ such that for weights $\{w_t = 1\}_{t \geq 0}$ it holds that:*

$$\frac{1}{(T+1)} \sum_{t=0}^T e_t \leq \mathcal{O} \left(2 \left(\frac{cr_0}{T+1} \right)^{\frac{1}{2}} + 2B^{1/3} \left(\frac{r_0}{T+1} \right)^{\frac{2}{3}} + \frac{br_0 + A\tau}{T+1} \right).$$

Proof. With $a = 0$, constant stepsizes $\eta_t = \eta$ and weights $\{w_t = 1\}_{t \geq 0}$ (31) is equivalent to

$$\begin{aligned} \frac{1}{2(T+1)} \sum_{t=0}^T e_t & \leq \frac{1}{(T+1)\gamma} \sum_{t=0}^T (r_t - r_{t+1}) + D\gamma + B\gamma^2 + \frac{A}{T+1} \sum_{t=0}^T \left(1 - \frac{1}{64\tau} \right)^t \\ & \leq \frac{r_0}{(T+1)\gamma} + D\gamma + B\gamma^2 + \frac{64A\tau}{T+1}. \end{aligned}$$

To conclude the proof we tune the stepsize for the first three terms using Lemma 26. □

Lemma 26 (Tuning the stepsize). *For any parameters $r_0 \geq 0, b \geq 0, e \geq 0, d \geq 0$ there exists constant stepsize $\eta \leq \frac{1}{b}$ such that*

$$\Psi_T := \frac{r_0}{\gamma(T+1)} + D\eta + B\eta^2 \leq 2 \left(\frac{Dr_0}{T+1} \right)^{\frac{1}{2}} + 2B^{1/3} \left(\frac{r_0}{T+1} \right)^{\frac{2}{3}} + \frac{br_0}{T+1}$$

Proof. Choosing $\eta = \min \left\{ \left(\frac{r_0}{D(T+1)} \right)^{\frac{1}{2}}, \left(\frac{r_0}{B(T+1)} \right)^{\frac{1}{3}}, \frac{1}{b} \right\} \leq \frac{1}{b}$ we have three cases

- $\eta = \frac{1}{b}$ and is smaller than both $\left(\frac{r_0}{D(T+1)} \right)^{\frac{1}{2}}$ and $\left(\frac{r_0}{B(T+1)} \right)^{\frac{1}{3}}$, then

$$\Psi_T \leq \frac{br_0}{T+1} + \frac{D}{b} + \frac{B}{b^2} \leq \left(\frac{Dr_0}{T+1} \right)^{\frac{1}{2}} + \frac{br_0}{T+1} + B^{1/3} \left(\frac{r_0}{T+1} \right)^{\frac{2}{3}}$$

- $\eta = \left(\frac{r_0}{D(T+1)} \right)^{\frac{1}{2}} < \left(\frac{r_0}{B(T+1)} \right)^{\frac{1}{3}}$, then

$$\Psi_T \leq 2 \left(\frac{r_0 D}{T+1} \right)^{\frac{1}{2}} + B \left(\frac{r_0}{D(T+1)} \right) \leq 2 \left(\frac{r_0 D}{T+1} \right)^{\frac{1}{2}} + B^{\frac{1}{3}} \left(\frac{r_0}{(T+1)} \right)^{\frac{2}{3}},$$

- The last case, $\eta = \left(\frac{r_0}{B(T+1)} \right)^{\frac{1}{3}} < \left(\frac{r_0}{D(T+1)} \right)^{\frac{1}{2}}$

$$\Psi_T \leq 2B^{\frac{1}{3}} \left(\frac{r_0}{(T+1)} \right)^{\frac{2}{3}} + D \left(\frac{r_0}{B(T+1)} \right)^{\frac{1}{3}} \leq 2B^{\frac{1}{3}} \left(\frac{r_0}{(T+1)} \right)^{\frac{2}{3}} + \left(\frac{Dr_0}{T+1} \right)^{\frac{1}{2}}. \quad \square$$

B.4 Non-convex Case

First, we state the descent Lemma for non-convex cases. Due to Lemma 5, it holds that

Lemma 27 (Descent lemma for non-convex case, Lemma 11 from [16]). *Under Assumptions as in Theorem 2, the averages $\bar{\mathbf{x}}^{(t)} := \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i^{(t)}$ of the iterates of Algorithm 1 with the constant stepsize $\gamma < \frac{1}{4L(M+1)}$ satisfy*

$$\mathbb{E}_{t+1} f(\bar{\mathbf{x}}^{(t+1)}) \leq f(\bar{\mathbf{x}}^{(t)}) - \frac{\gamma}{4} \left\| \nabla f(\bar{\mathbf{x}}^{(t)}) \right\|_2^2 + \frac{\gamma L^2}{n} \sum_{i=1}^n \left\| \bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)} \right\|_2^2 + \frac{L}{n} \gamma^2 \sigma^2. \quad (32)$$

Similarly as for the convex cases we prove the following recursion

Lemma 28 (Consensus distance recursion). *There are exists absolute constants $C_1, C_2 > 0$ such that*

$$\mathbb{E} \|\Psi_{t+k}\|_F^2 \leq \frac{3}{4} \|\Psi_t\|_F^2 + \frac{1}{128\tau} \sum_{j=0}^{k-1} \mathbb{E} \|\Psi_{t+j}\|_F^2 + C_1 \gamma^2 \tau n \sum_{j=0}^{k-1} e_{t+j} + C_2 \gamma^2 \left(\frac{\tau n}{c^2} + \tau^2 \right) \sigma^2 \quad (33)$$

where $e_j = \left\| \nabla f(\bar{\mathbf{x}}^{(j)}) \right\|^2$, $\tau \leq k \leq 2\tau$, $\tau = \frac{2}{p} \log \left(\frac{50}{p} (1 + \log \frac{1}{p}) \right) + 1$, p and c are defined in (3), $\Psi_t = (\Delta X^{(t)}, \gamma \Delta Y^{(t)})$ and is defined in (9).

Proof. The proof starts exactly the same as in the convex cases, Lemma 20. The difference comes when estimating terms T_1 and T_2 .

The second term T_2 . After splitting the stochastic noise,

$$\begin{aligned} \mathbb{E}[T_2] &\leq 3 \mathbb{E} \left\| \sum_{j=1}^k \left(\nabla f(X^{(t+j)}) - \nabla f(X^{(t+j-1)}) \right) \tilde{W}^{\tau-j} \right\|_F^2 + 6kn\sigma^2 \\ &\stackrel{(16)}{\leq} 3k \sum_{j=1}^k \mathbb{E} \left\| \nabla f(X^{(t+j)}) - \nabla f(X^{(t+j-1)}) \right\|_F^2 + 6kn\sigma^2 \end{aligned}$$

Estimating separately

$$\begin{aligned} \mathbb{E} \left\| \nabla f(X^{(t+j)}) - \nabla f(X^{(t+j-1)}) \right\|_F^2 &\stackrel{(16)}{\leq} 3 \mathbb{E} \left\| \nabla f(X^{(t+j)}) - \nabla f(\bar{X}^{(t+j)}) \right\|_F^2 + 3 \left\| \nabla f(\bar{X}^{(t+j-1)}) - \nabla f(X^{(t+j-1)}) \right\|_F^2 \\ &\quad + 3 \left\| \nabla f(\bar{X}^{(t+j)}) - \nabla f(\bar{X}^{(t+j-1)}) \right\|_F^2 \\ &\stackrel{(4)}{\leq} 3L^2 \mathbb{E} \left\| X^{(t+j)} - \bar{X}^{(t+j)} \right\|_F^2 + 3L^2 \left\| \bar{X}^{(t+j-1)} - X^{(t+j-1)} \right\|_F^2 \\ &\quad + 3L^2 \left\| \bar{X}^{(t+j)} - \bar{X}^{(t+j-1)} \right\|_F^2 \end{aligned}$$

And for the last term we estimate

$$\begin{aligned} \mathbb{E} \left\| \bar{\mathbf{x}}^{(t+j)} - \bar{\mathbf{x}}^{(t+j-1)} \right\|_2^2 &\leq \gamma^2 \left\| \frac{1}{n} \sum_{i=1}^n \nabla f_i(\mathbf{x}_i^{(t+j-1)}) \right\|_2^2 + \gamma^2 \frac{\sigma^2}{n} \\ &\leq 2\gamma^2 \left\| \frac{1}{n} \sum_{i=1}^n \nabla f_i(\mathbf{x}_i^{(t+j-1)}) - \frac{1}{n} \sum_{i=1}^n \nabla f_i(\bar{\mathbf{x}}^{(t+j-1)}) \right\|_2^2 + 2\gamma^2 \left\| \nabla f(\bar{\mathbf{x}}^{(t+j-1)}) \right\|^2 + \gamma^2 \frac{\sigma^2}{n} \\ &\leq 2\gamma^2 L^2 \frac{1}{n} \sum_{i=1}^n \left\| \mathbf{x}_i^{(t+j-1)} - \bar{\mathbf{x}}^{(t+j-1)} \right\|^2 + 2\gamma^2 \left\| \nabla f(\bar{\mathbf{x}}^{(t+j-1)}) \right\|^2 + \gamma^2 \frac{\sigma^2}{n} \end{aligned}$$

Thus, using that $\gamma < \frac{1}{24L\tau}$, $k \leq 2\tau$

$$\mathbb{E}[T_2] \leq \tau \sum_{j=0}^{k-1} n \mathbb{E} \left\| \nabla f(\bar{\mathbf{x}}^{(t+j)}) \right\|^2 + 21L^2 \tau \sum_{j=0}^{k-1} \mathbb{E} \left\| X^{(t+j)} - \bar{X}^{(t+j)} \right\|_F^2 + 7\tau n \sigma^2.$$

Term T_1 . Similarly, after separating the stochastic noise with $Z^{(t)} = G^{(t)} - \nabla f(X^{(t)})$,

$$T_1 \stackrel{(18)}{\leq} 2 \left\| \sum_{j=1}^k \left[\nabla f(X^{(t+j)}) - \nabla f(X^{(t+j-1)}) \right] (k-j) \tilde{W}^{k-j} \right\|_F^2 + 2 \left\| \sum_{j=1}^k \left(Z^{(t+j)} - Z^{(t+j-1)} \right) (k-j) \tilde{W}^{k-j} \right\|_F^2.$$

We add and subtract $\nabla f(\bar{X}^{t+j}), \nabla f(\bar{X}^{t+j-1})$ in the first term and denote $D^{(j)} = \nabla f(X^{(j)}) - \nabla f(\bar{X}^{(j)})$.

$$T_1 \leq 4 \left\| \sum_{j=1}^k \left(D^{(t+j)} - D^{(t+j-1)} \right) (k-j) \tilde{W}^{k-j} \right\|_F^2 + 4 \left\| \sum_{j=1}^k \left[\nabla f(\bar{X}^{t+j}) - \nabla f(\bar{X}^{t+j-1}) \right] (k-j) \tilde{W}^{k-j} \right\|_F^2 + 2 \left\| \sum_{j=1}^k \left(Z^{(t+j)} - Z^{(t+j-1)} \right) (k-j) \tilde{W}^{k-j} \right\|_F^2.$$

Terms with D and Z we estimate exactly the same as in the convex case, thus getting

$$\mathbb{E}[T_1] \stackrel{(6)}{\leq} \frac{64k}{c^2} \sum_{j=0}^{k-1} \|D^{(t+j)}\|_F^2 + \frac{32kn\sigma^2}{c^2} + 4 \underbrace{\left\| \sum_{j=1}^k \left[\nabla f(\bar{X}^{t+j}) - \nabla f(\bar{X}^{t+j-1}) \right] (k-j) \tilde{W}^{k-j} \right\|_F^2}_{T_3}$$

It is only left to estimate the last term. For that we use Lemma 14, and $\frac{1}{p} \leq \tau$ due to our choice of τ ,

$$\begin{aligned} T_3 &\stackrel{(16)}{\leq} k \sum_{j=1}^k \left\| \left[\nabla f(\bar{X}^{t+j}) - \nabla f(\bar{X}^{t+j-1}) \right] (k-j) \tilde{W}^{k-j} \right\|_F^2 \stackrel{L. 14}{\leq} 4k\tau^2 \sum_{j=1}^k \left\| \nabla f(\bar{X}^{t+j}) - \nabla f(\bar{X}^{t+j-1}) \right\|_F^2 \\ &\leq 4k\tau^2\gamma^2 \sum_{j=1}^k \left[2L^2 \left\| X^{(t+j-1)} - \bar{X}^{(t+j-1)} \right\|_F^2 + 2n \left\| \nabla f(\bar{\mathbf{x}}^{(t+j-1)}) \right\|^2 + \sigma^2 \right] \end{aligned}$$

Where the last inequality was obtained while estimating Term T_2 . Using that $k \leq 2\tau$, $\gamma \leq \frac{1}{24L\tau}$ and that $\|D^{(t+j)}\|_F^2 \leq L^2 \|X^{(t+j)} - \bar{X}^{(t+j)}\|_F^2$ by smoothness

$$\mathbb{E}[T_1] \stackrel{(6)}{\leq} \frac{129\tau}{c^2} L^2 \sum_{j=0}^{k-1} \left\| X^{(t+j)} - \bar{X}^{(t+j)} \right\|_F^2 + \tau \sum_{j=0}^{k-1} n \left\| \nabla f(\bar{\mathbf{x}}^{(t+j)}) \right\|^2 + \left(\frac{64\tau n}{c^2} + \tau^2 \right) \sigma^2$$

Summing T_1 and T_2 together, and using that $\gamma \leq \frac{c}{310\tau L}$

$$\mathbb{E} \|\Psi_{t+k}\|_F^2 \leq \frac{3}{4} \|\Psi_t\|_F^2 + \frac{1}{128\tau} \sum_{j=0}^{k-1} \mathbb{E} \|\Psi_{t+j}\|_F^2 + \gamma^2 10\tau n \sum_{j=0}^{k-1} \left\| \nabla f(\bar{\mathbf{x}}^{(t+j)}) \right\|^2 + 5\gamma^2 \left(\frac{64\tau n}{c^2} + \tau^2 \right) \sigma^2$$

□

Next, we unroll this recursion with Lemma 8.

For $\gamma < \frac{c}{\sqrt{7}B_1L\tau} \leq \frac{1}{2L\tau}$, and with some positive absolute constants $B_1, B_2 > 0$ it holds,

$$\mathbb{E} \|\Psi_t\|_F^2 \leq \left(1 - \frac{1}{64\tau} \right)^t A_0 + B_1\tau\gamma^2 \sum_{j=0}^{t-1} \left(1 - \frac{1}{64\tau} \right)^{t-j} n e_j + B_2\gamma^2 \left(\frac{\tau n}{c^2} + \tau^2 \right) \sigma^2 \quad (34)$$

where $e_j = \left\| \nabla f(\bar{\mathbf{x}}^{(j)}) \right\|^2$, $A_0 = 16 \|\Delta X^{(0)}\|_F^2 + \frac{24\gamma^2}{p^2} \|\Delta Y^{(0)}\|_F^2$.

The rest of proof consists of combining (34) with the descent lemma for non-convex case (32) in similar fashion as in Lemmas 22, 23; and further using Lemma 25 to obtain the final rate.

C Experimental Setup and Additional Plots

We illustrate the dependence of the convergence rate on the parameters c and p .

In these experiments, we vary p and c (by changing the mixing matrix) and measure the value of $f(\bar{\mathbf{x}}^{(t)}) - f^*$ that GT reaches after a large number of steps t , when using a constant stepsize γ (chosen small enough so that none of the runs diverges). According to our theoretical results, GT converges to the level $\mathcal{O}\left(\frac{\gamma\sigma^2}{n} + \frac{\gamma^2\sigma^2}{pc^2}\right)$ in a linear number of steps (to reach higher accuracy, smaller stepsizes must be used). Thus, for n large enough, this term is dominated by $\mathcal{O}\left(\frac{\gamma^2\sigma^2}{pc^2}\right)$, which we aim to measure. In all experiments we ensure that the first term is at least by order of magnitude smaller than the second by comparing the noise level with GT on a fully-connected topology.

C.1 Problem Instances

We used $n = 300$, $d = 100$.

Setup A (Gaussian Noise). We consider quadratic functions defined as $f_i(\mathbf{x}) = \|\mathbf{x}\|^2$, and $\mathbf{x}^{(0)}$ is randomly initialized from a normal distribution $\mathcal{N}(0, 1)$. We add artificially stochastic noise to gradients as $\nabla F_i(\mathbf{x}, \xi) = \nabla f_i(\mathbf{x}) + \xi$, where $\xi \sim \mathcal{N}(0, \frac{\sigma^2}{d}I)$.

Setup B (Structured Noise). We consider quadratic functions defined as $f_i(\mathbf{x}) = \|\mathbf{x}\|^2$, and $\mathbf{x}^{(0)}$ is randomly initialized from a normal distribution $\mathcal{N}(0, 1)$. We add artificially stochastic noise to gradients as $\nabla F(X, \xi) = \nabla f(X) + \text{diag}(\xi)V$, where $\xi \sim \mathcal{N}(0, \frac{\sigma^2}{d}I)$ is a d -dimensional Gaussian noise vector, $\text{diag}(\xi)$ a matrix with ξ on the diagonal, and $V \in \mathbb{R}^{d \times n}$ is a matrix with half of the rows equal to $\mathbf{v} \in \mathbb{R}^n$, and half of the rows equal to $\mathbf{u} \in \mathbb{R}^n$, where \mathbf{v}, \mathbf{u} are eigenvectors of the mixing matrix, $W\mathbf{v} = \lambda_n(W)\mathbf{v}$, i.e. corresponding to the smallest eigenvalue of W , and $W\mathbf{u} = \lambda_2(W)\mathbf{u}$, i.e. corresponding to the second largest eigenvalue of W .

This is motivated by the observations in Lemma 13, where we noted that components in the eigenspace corresponding to the smallest eigenvalue of W get amplified the most.

C.2 Graph Topologies and Mixing Matrices

Interpolated Ring (between uniform weights and interpolate with a fully-connected topology).

We consider the ring topology W_{ring} on n nodes, where each node i has self weight $w_{ii} = \frac{1}{3}$ and $w_{i,1+(i \bmod n)} = w_{i,(i-2 \bmod n)+1} = \frac{1}{3}$ for its neighbors. We interpolate this uniform weight ring topology with a fully-connected topology, $W_{\text{complete}} = \frac{1}{n}\mathbf{1}\mathbf{1}^\top$, that is, $W_\alpha := \alpha W_{\text{ring}} + (1 - \alpha)W_{\text{complete}}$. The eigenvalues of W_{ring} are $\lambda(W_{\text{ring}}) \in [-\frac{1}{3}, 1]$, and $\lambda(W_{\text{complete}}) \in [0, 1]$, and therefore c of W_α is also a constant.

Ring with smaller self weight. We consider the ring topology W_w on n nodes, where each node i has self weight $w_{ii} = w \leq \frac{1}{3}$ and $w_{i,1+(i \bmod n)} = w_{i,1+(i-2 \bmod n)} = \frac{1-w}{2}$ for its neighbors. The eigenvalues of W_w are $\lambda(W_w) \in [2w - 1, 1]$, and therefore c can become small by choosing w (note that the $\lambda_n(W_w)$, while decreasing for smaller w , is not equal to $2w - 1$ in general, except when $w = \frac{1}{3}$). We measure the exact value $\lambda_n(W_w)$ when reporting c below.

C.3 Additional Plots for Setup A

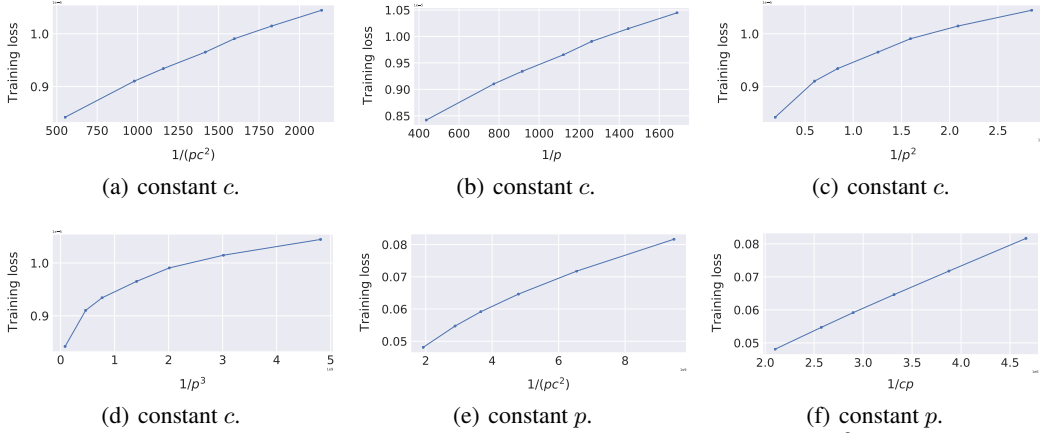


Figure 4: Impact of c and p on the convergence with the Gaussian stochastic noise $\sigma^2 = 1$. The first four subfigures illustrate the impact of p on convergence when c is kept constant; showing a linear scaling of the loss compared to $\frac{1}{p}$. The last subfigure varies c in the graph while keeping p as a constant, and we see a linear scaling compared to $\frac{1}{c^2}$.

C.4 Additional Plots for Setup B

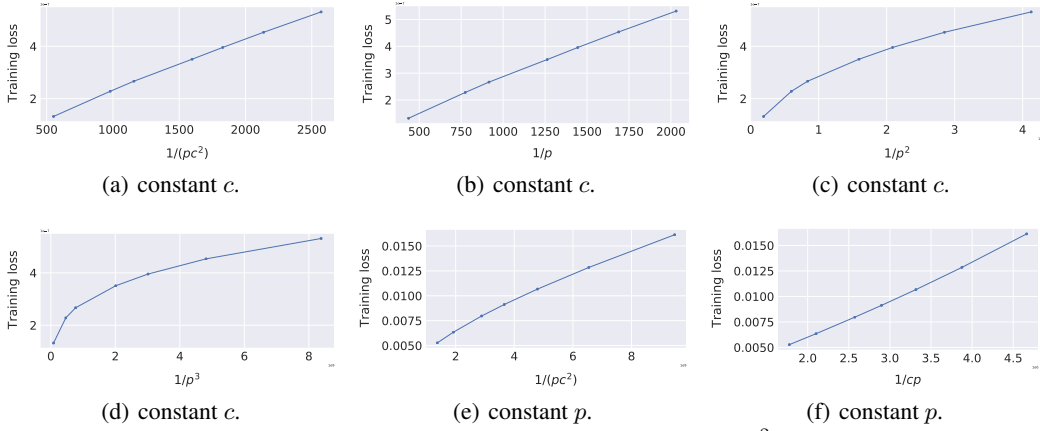


Figure 5: Impact of c and p on convergence with the structured stochastic noise $\sigma^2 = 1$. The first four subfigures illustrate the impact of p on convergence when c is kept constant; showing a linear scaling of the loss compared to $\frac{1}{p}$. The last subfigure varies c in the graph while keeping p as a constant, and we can see a linear scaling compared to $\frac{1}{c^2}$.

In Figures 4 and 5 we study the impact of c and p on the convergence. These findings support the $\mathcal{O}\left(\frac{\gamma^2 \sigma^2}{pc^2}\right)$ scaling predicted by theory—however, cannot replace a formal proof. We leave this for future work.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes] See section 6 and 8.
 - (c) Did you discuss any potential negative societal impacts of your work? [N/A]
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes] See Sections 3–4.
 - (b) Did you include complete proofs of all theoretical results? [Yes] See appendix.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Supplementary.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] Main text and appendix.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] Appendix.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [N/A]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]