

Convergence and Privacy of Decentralized Nonconvex Optimization with Gradient Clipping and Communication Compression

Boyue Li*
Carnegie Mellon University

Yuejie Chi*
Carnegie Mellon University

May 15, 2023

Abstract

Achieving communication efficiency in decentralized machine learning has been attracting significant attention, with communication compression recognized as an effective technique in algorithm design. This paper takes a first step to understand the role of gradient clipping, a popular strategy in practice, in decentralized nonconvex optimization with communication compression. We propose **PORTER**, which considers two variants of gradient clipping added before or after taking a mini-batch of stochastic gradients, where the former variant **PORTER-DP** allows local differential privacy analysis with additional Gaussian perturbation, and the latter variant **PORTER-GC** helps to stabilize training. We develop a novel analysis framework that establishes their convergence guarantees without assuming the stringent bounded gradient assumption. To the best of our knowledge, our work provides the first convergence analysis for decentralized nonconvex optimization with gradient clipping and communication compression, highlighting the trade-offs between convergence rate, compression ratio, network connectivity, and privacy.

Keywords: communication compression, gradient clipping, convergence rate, local differential privacy

1 Introduction

Decentralized machine learning has been attracting significant attention in recent years, which can be often modeled as a nonconvex finite-sum optimization problem:

$$\min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x}) := \frac{1}{n} \sum_{i=1}^n f_i(\mathbf{x}), \quad \text{where } f_i(\mathbf{x}) = \frac{1}{m} \sum_{\mathbf{z} \in \mathcal{Z}_i} \ell(\mathbf{x}; \mathbf{z}), \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^d$ and \mathbf{z} denote the optimization parameter and one data sample, $\ell(\mathbf{x}; \mathbf{z})$ denotes the sample loss function that is nonconvex in \mathbf{x} , and $f_i(\mathbf{x})$ and $f(\mathbf{x})$ denote the local objective function at agent i and the global objective function. In addition, \mathcal{Z}_i denotes the dataset at agent i , $m = |\mathcal{Z}_i|$ denotes the local sample size, and n denotes the number of agents. An undirected communication graph \mathcal{G} is used to model the connectivity between any two agents, where there is an edge between agent i and j only if they can communicate. The goal is to efficiently optimize the global objective function $f(\mathbf{x})$ in a decentralized manner, subject to the network connectivity constraints specified by \mathcal{G} .

Communication efficiency is critical to decentralized optimization algorithms, as communication can quickly become bottleneck of the system as the number of agents and the size of the model increase. This has led to the development of communication compression (or quantization) techniques, which can significantly reduce the communication burden per round by transferring compressed information, especially when the communication bandwidth is limited. Therefore, a number of recent works have focused on designing decentralized nonconvex optimization algorithms with communication compression, including but not limited to [KSJ19, SDGD21, ZLL+22, TLQ+19, HP23, YCC+23].

*Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA. Emails: {boyuel, yuejiec}@andrew.cmu.edu.

Built upon this line of work, the paper aims to understand the role of gradient clipping in decentralized nonconvex optimization algorithms with communication compression. On the one hand, gradient clipping has been used widely in privacy-preserving algorithms [ACG⁺16] to enable (local) differential privacy guarantees [Dwo08]. On the other hand, gradient clipping is also observed to be beneficial in stabilizing neural network training [ZHSJ20a]. However, since gradient clipping necessarily introduces bias, the characterization of the convergence becomes much more challenging compared to their unclipped counterpart. As a result, most of the existing theoretical analyses for stochastic gradient algorithms with clipping — in the context of centralized and server/client settings — make strong assumptions such as the bounded gradient assumption [ACG⁺16, WJEG19, LZLC22] and the uniformly bounded gradient assumption [YZCL22, ZJFW20, ZHSJ20a]. To the best of our knowledge, the convergence of stochastic gradient algorithms with clipping in the decentralized setting has not been investigated before.

1.1 Our contributions

This paper proposes PORTER (cf. Algorithm 1),¹ a communication-efficient decentralized algorithm for nonconvex finite-sum optimization with gradient clipping and communication compression. PORTER is built on BEER [ZLL⁺22] — a fast decentralized algorithm with communication compression proposed recently — by introducing gradient clipping to the local stochastic gradient computation at agents, while inheriting the desirable designs such as error feedback and stochastic gradient tracking that are crucial in enabling the fast convergence of BEER. PORTER considers two variants of gradient clipping, corresponding to adding it before or after taking a mini-batch of stochastic gradients. In particular, the former variant PORTER-DP allows local differential privacy (LDP) analysis with additional Gaussian perturbation, and the latter variant PORTER-GC helps to stabilize training. Assuming a smooth clipping operator (Definition 2) and general compression operators (Definition 3), the highlights of our contributions are as follows.

1. We establish that PORTER-DP (cf. Algorithm 1) achieves (ϵ, δ) -LDP under appropriate Gaussian perturbation. Under the bounded gradient assumption (when gradient clipping can be ignored), PORTER-DP converges in average squared ℓ_2 gradient norm as $\frac{1}{T} \sum_{t \in [T]} \mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2^2 \lesssim \rho^{-\frac{4}{3}} (1 - \alpha)^{-\frac{8}{3}} \phi_m^{-1}$ in $T = \phi_m^{-2}$ iterations, where $\bar{\mathbf{x}}^{(t)}$ is the average parameter, $\phi_m = \frac{\sqrt{d \log(1/\delta)}}{m\epsilon}$ is the baseline utility for a centralized stochastic algorithm to achieve (ϵ, δ) -DP with m data samples [ACG⁺16], $\rho \in (0, 1]$ is the compression ratio, and $\alpha \in [0, 1)$ is the mixing rate of the topology.
2. However, the bounded gradient assumption might be too stringent to hold in practice. Instead we further establish that under the local variance and bounded dissimilarity assumptions, PORTER-DP converges in minimum ℓ_2 gradient norm as $\min_{t \in [T]} \mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 \lesssim \rho^{-\frac{2}{3}} (1 - \alpha)^{-\frac{4}{3}} \phi_m^{-1/2}$ in $T = \phi_m^{-2}$ iterations.
3. We establish that under the local variance and bounded dissimilarity assumptions, by properly choosing the mini-batch size, PORTER-GC converges in minimum ℓ_2 gradient norm as $\min_{t \in [T]} \mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 \lesssim \rho^{-\frac{2}{3}} (1 - \alpha)^{-\frac{4}{3}} T^{-1/2}$, which matches the convergence rate of classical centralized stochastic algorithms.

Our work develops a novel analysis framework that establishes their convergence guarantees without assuming the stringent bounded gradient assumption, highlighting comprehensive trade-offs between convergence rate, compression ratio, network connectivity, and privacy. To the best of our knowledge, our work provides the first private decentralized optimization algorithm with communication compression, and a systematic investigation of gradient clipping in the fully decentralized setting. Table 1 provides a detailed comparison of PORTER-DP with prior art on private server-client algorithms, where the bounded gradient assumption is all in effect except ours.

1.2 Related works

Decentralized optimization algorithms have been extensively studied to solve large-scale optimization problems. We review most closely related works in this section, and refer readers to more comprehensive reviews

¹The name is coined for two reasons: 1) PORTER has strong connection to the prior algorithm BEER (porter is a kind of dark beer), and 2) the authors developed this algorithm in Porter Hall at Carnegie Mellon University.

Table 1: Comparison of final utility upper bounds and communication complexities of different stochastic algorithms that achieve (ϵ, δ) -DP/LDP. The Big- O notation (defined in Section 1.3) is omitted for simplicity. DP-SGD is a single-server optimization algorithm that serves as a baseline, to show the overhead brought in by the distributed setting. DDP-SRM and Soteria-SGD are server/client distributed algorithms, but DDP-SRM doesn’t use communication compression.

⁽¹⁾ $\theta = (1 - \omega)^{-3/2} n^{-1/2}$, where ω is the parameter for unbiased compression.

Algorithm	Privacy	Compression operator	Bounded gradient	Utility	Communication rounds
DP-SGD [ACG ⁺ 16]	DP	-	✓	ϕ_m	-
DDP-SRM [WJEG19]	DP	-	✓	$\frac{1}{n}\phi_m$	$n^2 d\phi_m^{-1}$
Soteria-SGD ⁽¹⁾ [LZLC22]	LDP	Unbiased	✓	$(1 + \theta^{1/2})\left(\frac{1+\omega}{n}\right)^{1/2}\phi_m$	$(1 + \theta^{1/2})\left(\frac{n}{1+\omega}\right)^{2/3} d\phi_m^{-1}$
PORTER-DP (Algorithm 1)	LDP	General	✓	$\frac{1}{(1-\alpha)^{8/3}\rho^{4/3}}\phi_m$	ϕ_m^{-2}
PORTER-DP (Algorithm 1)	LDP	General	✗	$\frac{1}{(1-\alpha)^{16/3}\rho^{8/3}}\phi_m$	ϕ_m^{-2}

in [NRB20, XPNK20].

Decentralized stochastic nonconvex optimization. Decentralized stochastic algorithms have been a actively researched area in recent years. Various algorithms have been proposed by directly adapting existing centralized algorithms, e.g., [KDG03, XB04, Sha07, BJ13, LZZ⁺17, ALBR19, WJ21]. However, the simple adaptations usually fail to achieve better convergence rates. Gradient tracking [ZM10], originally proposed by the control theory community, can be used to track the global gradient at each agent, which leads to a simple systematic framework for extending existing centralized algorithms to the decentralized setting. Gradient tracking can be used for both deterministic optimization algorithms, e.g., [DLS16, NOS17, QL18, LCCC20], and stochastic algorithms, e.g., [SLH20, XKK22b, XKK22a, LLC22, HSZ⁺22, LY22].

Communication compression. In [DSZOR15, AGL⁺17], gradient compression was adopted to create a server/client distributed SGD algorithm, however, the large variance of compressed gradients leads to a sub-optimal convergence rate. [SFD⁺14] first proposed the use of error feedback to compensate for the variance induced by compression. [SCJ18, AHJ⁺18, MGTR19, LKQR20, GBLR21, LR21] all equipped similar mechanism to improve convergence for server/client distributed optimization algorithms, and [RSF21, FSG⁺21] formalized the error feedback mechanism and reaches an $O(1/T)$ convergence rate for smooth nonconvex objective functions. [TGZ⁺18, KSJ19, SDGD21, TMHP20, ZLL⁺22, YCC⁺23, LLP23, ZBLR21] further extended communication compression schemes to the decentralized setting.

Private optimization algorithms. The concern of leaking sensitive data has been increasing with the rapid development of large-scale machine learning systems. To address this concern, the concept of differential privacy is widely adopted [DMNS06, Dwo08], where a popular approach to protect privacy is adding noise to the model or gradients. This approach is first adopted in the single server setting to design differentially private optimization algorithms [ACG⁺16, WYX17, INS⁺19, FKT20, CWH20, WDX20], while [HDJ⁺20, ACCÖ21, NBD22, DLC⁺22, LZLC22, MS23, ZCH⁺22] considered differential privacy for the server/client distributed setting.

Gradient clipping. Understanding gradient clipping has gained significant attention in recent years. Earlier works, e.g. [PMB13, BBP13, KLL16, KFI17, YGG17], used gradient clipping as a pure heuristic to solve gradient vanishing/exploding problems without theoretical understandings. Then, [ZHSJ20b, ZKV⁺20, ZJFW20, RLDJ23] introduced theoretical analyses to understand its impact on the convergence rate and

training performance. This question is also investigated in [CWH20, ZCH⁺22, DKX⁺22, FLFL23], which applies gradient clipping to limit the magnitude of each of the sample gradients, so that the variance of privacy perturbation can be decided without the bounded gradient assumption. While finishing up this paper, we became aware of [KHS23], which also develops convergence guarantees on the minimum ℓ_2 gradient norm of clipped stochastic gradient algorithms in the centralized setting with a piece-wise linear clipping operator. In contrast, our focus is on the decentralized setting with a smooth clipping operator, where extra care is taken to deal with the discrepancy between the local and global objective functions.

1.3 Paper organization and notation

Section 2 introduces preliminary concepts, Section 3 describes the algorithm development, Section 4 shows the theoretical performance guarantees for PORTER, Section 5 provides numerical evidence to support the analysis, and Section 6 concludes the paper. Proofs are postponed to appendices.

Throughout this paper, we use uppercase and lowercase boldface letters to represent matrices and vectors, respectively. We use $\|\cdot\|_{\text{op}}$ for matrix operator norm, $\|\cdot\|_{\text{F}}$ for Frobenius norm, \mathbf{I}_n for the n -dimensional identity matrix, $\mathbf{1}_n$ for the n -dimensional all-one vector and $\mathbf{0}_{d \times n}$ for the $(d \times n)$ -dimensional zero matrix. For two real functions $f(\cdot)$ and $g(\cdot)$ defined on \mathbb{R}^+ , we say $f(x) = O(g(x))$ or $f(x) \lesssim g(x)$ if there exists some universal constant $M > 0$ such that $f(x) \leq Mg(x)$. The notation $f(x) = \Omega(g(x))$ or $f(x) \gtrsim g(x)$ means $g(x) = O(f(x))$.

2 Preliminaries

Mixing. The information mixing between agents is conducted by updating the local information via a weighted sum of information from neighbors, which is characterized by a mixing (gossiping) matrix. Concerning this matrix is an important quantity called the mixing rate, defined in Definition 1.

Definition 1 (Mixing matrix and mixing rate). *The mixing matrix is a matrix $\mathbf{W} = [w_{ij}] \in \mathbb{R}^{n \times n}$, such that $w_{ij} = 0$ if agent i and j are not connected according to the communication graph \mathcal{G} . Furthermore, $\mathbf{W}\mathbf{1}_n = \mathbf{1}_n$ and $\mathbf{W}^\top \mathbf{1}_n = \mathbf{1}_n$. The mixing rate of a mixing matrix \mathbf{W} is defined as*

$$\alpha := \left\| \mathbf{W} - \frac{1}{n} \mathbf{1}_n \mathbf{1}_n^\top \right\|_{\text{op}}. \quad (2)$$

The mixing rate describes the connectivity of a communication graph and the speed of information sharing. Generally, a better connected graph leads to a smaller mixing rate, for example, \mathbf{W} can be the averaging matrix for a fully connected communication network, which results in $\alpha = 0$. A comprehensive list of bounds on $1 - \alpha$ is provided by [NOR18, Proposition 5]. Our analysis does not require the mixing matrix to be doubly stochastic, while allows us to use a non-symmetric matrix with negative values as the mixing matrix, such as the FDLA matrix [XB04], which has a smaller mixing rate under the same connectivity pattern.

Gradient clipping. In practice, gradient clipping is frequently adopted to ensure the gradients are within a predetermined region, so that the variance of privacy perturbation can be decided accordingly. The clipping operator we adopt is a smooth clipping operator [YZCL22] defined in Definition 2, which scales a vector into a ball of radius τ centered at the origin.

Definition 2 (Smooth clipping operator). *For $\mathbf{x} \in \mathbb{R}^d$, the clipping operator is defined as*

$$\text{Clip}_\tau(\mathbf{x}) = \frac{\tau}{\tau + \|\mathbf{x}\|_2} \mathbf{x}.$$

For $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n] \in \mathbb{R}^{d \times n}$, the distributed clipping operator is defined as

$$\text{Clip}_\tau(\mathbf{X}) = [\text{Clip}_\tau(\mathbf{x}_1), \dots, \text{Clip}_\tau(\mathbf{x}_n)].$$




Figure 1: Illustration of input norm and clipped norm for the smooth clipping operator (Definition 2) and piece-wise linear clipping operator, where τ is the clipping parameter.

Remark 1. Another widely used clipping operator is the piece-wise linear clipping operator, which scales inputs whenever its gradient norm is larger than τ and does nothing otherwise, defined by

$$\text{Clip}_\tau(\mathbf{x}) = \mathbf{x} \min \{1, \tau/\|\mathbf{x}\|_2\}.$$

Figure 1 plots the norm of a vector before and after clipping for these two clipping operators, which show that they behave quite similarly.

Compression operators. Following [RSF21, FSG⁺21], Definition 3 defines a randomized general compression operator that only guarantees the expected compression error $\mathbb{E}\|\mathcal{C}(\mathbf{x}) - \mathbf{x}\|_2^2$ is less than the magnitude of original message $\|\mathbf{x}\|_2^2$.

Definition 3 (General compression operator). A randomized map $\mathcal{C} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a ρ -compression operator if $\forall \mathbf{x} \in \mathbb{R}^d$ and some $\rho \in [0, 1]$, the following inequality holds:

$$\mathbb{E}\|\mathcal{C}(\mathbf{x}) - \mathbf{x}\|_2^2 \leq (1 - \rho)\|\mathbf{x}\|_2^2.$$

Many widely used compression schemes can be modeled as special cases, for example, random sparsification and top- k compression.

Example 1 (Random sparsification). Random sparsification keeps an element from a d -dimensional vector with probability $\frac{k}{d}$. Let $\mathbf{u} \in \mathbb{R}^d$ where $u_i \sim B(\frac{k}{d})$, then random sparsification is defined as $\text{random}_k(\mathbf{x}) = \mathbf{u} \odot \mathbf{x}$, which satisfies Definition 3 with $\rho = \frac{k}{d}$.

Example 2 (top $_k$). top $_k$ [AHJ⁺18, SCJ18] keeps k elements that have the largest absolute values and sets other elements to 0, which is defined as $\text{top}_k(\mathbf{x}) := \mathbf{x} \odot \mathbf{u}(\mathbf{x})$, where $[\mathbf{u}(\mathbf{x})]_i = 1$ if the absolute value of the i -th element is one of the k -largest absolute values, otherwise $[\mathbf{u}(\mathbf{x})]_i = 0$. It follows that top $_k$ satisfies Definition 3 with $\rho = k/d$.

Local differential privacy. In decentralized learning systems, all agents share information with their neighbors that are potentially sensitive. If some agents are exploited by adversaries, the system will face a risk of privacy leakage even when the system-level privacy is protected. Therefore, we introduce local differential privacy (LDP) — defined in Definition 4 — following [DJW13, CABP13, XYD19], which protects each agent’s privacy from leaking to other agents.

Definition 4 (Local differential privacy (LDP)). A randomized mechanism $\mathcal{M} : \mathcal{Z} \rightarrow \mathcal{R}$ with domain \mathcal{Z} and range \mathcal{R} satisfies (ϵ, δ) -local differential privacy for client i , if for any two neighboring dataset $\mathbf{Z}_i, \mathbf{Z}'_i \in \mathcal{Z}$ at client i and for any subset of outputs $\mathbf{R} \subseteq \mathcal{R}$, it holds that

$$\mathbb{P}(\mathcal{M}(\mathbf{Z}_i) \in \mathbf{R}) \leq e^\epsilon \mathbb{P}(\mathcal{M}(\mathbf{Z}'_i) \in \mathbf{R}) + \delta. \quad (3)$$

The two datasets \mathbf{Z}_i and \mathbf{Z}'_i are called neighboring if they are only different by one data point at client i .

Definition 4 is a stricter privacy guarantee because it can imply general differential privacy (DP). Consequently, LDP requires a larger perturbation variance than general DP. To identify the impact of the decentralized LDP setting compared to centralized DP setting, we define the baseline utility

$$\phi_m = \frac{\sqrt{d \log(1/\delta)}}{m\epsilon}, \quad (4)$$

which can be understood as the final utility of a centralized system with m data samples that guarantees (ϵ, δ) -DP. For typical private problems, the local sample size m has to be large enough for the privacy perturbation to deliver meaning guarantees, we impose a mild assumption that $\phi_m < 1$ to simplify presentation. For example, the problem defined in (1) has in total mn data samples, running an (ϵ, δ) -DP algorithm on one server that can access all data will achieve $\frac{1}{n}\phi_m$ utility in $n\phi_m^{-1}$ iterations.

3 Proposed algorithm

We propose PORTER, a novel stochastic decentralized optimization algorithm for finding first-order stationary points of nonconvex finite-sum problems with gradient clipping and communication compression; the details are described in Algorithm 1. On a high level, PORTER is composed of local stochastic gradient updates and neighboring information sharing, following a similar structure as BEER [ZLL⁺22], in terms of the use of error feedback [RSF21], which accelerates the convergence with biased compression operators, and stochastic gradient tracking to track the global gradient locally at each agent. A key difference, however, is the use of gradient clipping. Motivated by efficient training and privacy preserving, we consider two variants, corresponding to clipping *before* the mini-batch with privacy perturbation (PORTER-DP), and *after* taking the mini-batch (PORTER-GC), respectively.

Algorithm 1 PORTER

- 1: **input:** $\bar{\mathbf{x}}^{(0)}$, gradient stepsize η , consensus stepsize γ , clipping threshold τ , mini-batch size b , perturbation noise σ_p , number of iterations T
 - 2: **initialize:** $\mathbf{V}^{(0)} = \mathbf{Q}_v^{(0)} = \mathbf{G}_p^{(0)} = \mathbf{0}_{d \times n}$, $\mathbf{Q}_x^{(0)} = \mathbf{X}^{(0)} = \bar{\mathbf{x}}^{(0)} \mathbf{1}_n^\top$
 - 3: **for** $t = 1, \dots, T$ **do**
 - 4: Draw the local mini-batch of size b uniformly at random $\mathcal{Z}^{(t)} = \{\mathcal{Z}_i^{(t)}\}_{i=1}^n$
 - 5: **Option I: PORTER-DP (differentially-private SGD)**
 - 6: $\mathbf{G}_\tau^{(t)} = \frac{1}{b} \sum_{\mathbf{Z} \in \mathcal{Z}^{(t)}} \text{Clip}_\tau(\nabla \ell(\mathbf{X}^{(t-1)}; \mathbf{Z}))$
 - 7: $\mathbf{G}_p^{(t)} = \mathbf{G}_\tau^{(t)} + \mathbf{E}^{(t)}$, where $\mathbf{e}_i^{(t)} \sim \mathcal{N}(\mathbf{0}_d, \sigma_p^2 \mathbf{I}_d)$
 - 8: **Option II: PORTER-GC (SGD with gradient clipping)**
 - 9: $\mathbf{G}^{(t)} = \frac{1}{b} \sum_{\mathbf{Z} \in \mathcal{Z}^{(t)}} \nabla \ell(\mathbf{X}^{(t-1)}; \mathbf{Z})$
 - 10: $\mathbf{G}_p^{(t)} = \mathbf{G}_\tau^{(t)} = \text{Clip}_\tau(\mathbf{G}^{(t)})$
 - 11: $\mathbf{Q}_v^{(t)} = \mathbf{Q}_v^{(t-1)} + \mathcal{C}(\mathbf{V}^{(t-1)} - \mathbf{Q}_v^{(t-1)})$ ▷ Communication
 - 12: $\mathbf{V}^{(t)} = \mathbf{V}^{(t-1)} + \gamma \mathbf{Q}_v^{(t)} (\mathbf{W} - \mathbf{I}_n) + \mathbf{G}_p^{(t)} - \mathbf{G}_p^{(t-1)}$
 - 13: $\mathbf{Q}_x^{(t)} = \mathbf{Q}_x^{(t-1)} + \mathcal{C}(\mathbf{X}^{(t-1)} - \mathbf{Q}_x^{(t-1)})$ ▷ Communication
 - 14: $\mathbf{X}^{(t)} = \mathbf{X}^{(t-1)} + \gamma \mathbf{Q}_x^{(t)} (\mathbf{W} - \mathbf{I}_n) - \eta \mathbf{V}^{(t)}$
 - 15: **end for**
 - 16: **output:** $\mathbf{x}_{out} \sim \text{Uniform}(\{\mathbf{x}_i^{(t)} \mid i \in [n], t \in [T]\})$.
-

Before proceeding, we introduce some notation convenient for describing decentralized algorithms. Let $\mathbf{x}_i \in \mathbb{R}^d$ be the optimization variable at agent i , we define the collection of all optimization variables as a matrix $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n] \in \mathbb{R}^{d \times n}$, and the average as $\bar{\mathbf{x}} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i$. The gradient estimates \mathbf{V} , stochastic gradients \mathbf{G} , perturbation noise \mathbf{E} , compressed surrogate \mathbf{Q}_x and \mathbf{Q}_v , and their corresponding agent-wise values are defined analogously. The distributed gradient is defined as $\nabla F(\mathbf{X}) = [\nabla f_1(\mathbf{x}_1), \nabla f_2(\mathbf{x}_2), \dots, \nabla f_n(\mathbf{x}_n)] \in \mathbb{R}^{d \times n}$.

To provide more detail, PORTER initializes gradient-related variables to $\mathbf{0}_d$ and other variables to the same random value $\bar{\mathbf{x}}^{(0)}$, which improves stability in early iterations and simplifies analysis, but has no impact on the convergence rates. Within each iteration, PORTER is consisted of 3 major steps.

- (1) **Computing clipped stochastic gradients.** We consider two options. The first option PORTER-DP corresponds to differentially-private SGD (Line 6-Line 7), where Line 6 computes a batch of clipped stochastic gradient on each agent, and then Line 7 adds Gaussian noise to ensure privacy. The second option PORTER-GC corresponds to SGD with gradient clipping (Line 9-Line 10), where Line 9 computes a batch stochastic gradient on each agent, and Line 10 applies clipping to each batch stochastic gradient.
- (2) **Updating gradient estimates.** Line 11 updates the auxiliary variable $\mathbf{Q}_v^{(t)}$, which is a compressed surrogate of $\mathbf{V}^{(t-1)}$, by adding the compressed difference to itself $\mathcal{C}(\mathbf{V}^{(t-1)} - \mathbf{Q}_v^{(t-1)})$. Meanwhile, each agent i sends its compressed difference $\mathcal{C}(\mathbf{v}_i^{(t-1)} - \mathbf{q}_{v,i}^{(t-1)})$ to all of its neighbors, so that every neighbor can reconstruct the auxiliary variable $\mathbf{q}_{v,i}^{(t)}$ by accumulating this difference. Line 12 then adds a correction term $\gamma \mathbf{Q}_v^{(t)}(\mathbf{W} - \mathbf{I}_n)$, and applies stochastic gradient tracking to update gradient estimates.
- (3) **Updating variable estimates.** Similar to updating gradient estimates, Line 13 updates the auxiliary variable $\mathbf{Q}_x^{(t)}$, which is a compressed surrogate of variable estimates $\mathbf{X}^{(t-1)}$, and communicates with neighbors. Line 14 applies correction and updates the variable estimates by a gradient-style update.

4 Theoretical guarantees

This section theoretically analyzes the privacy and convergence properties of PORTER under various assumptions. Section 4.1 lists all assumptions required for convergence analysis, Section 4.2 shows the privacy and convergence of PORTER-DP using a specific perturbation variance, and Section 4.3 shows the convergence of PORTER corresponding to clipped SGD without privacy.

4.1 Assumptions

We start with smoothness assumption in Assumption 1, which is standard and required for all of our analysis.

Assumption 1 (*L-smoothness*). For any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ and any datum \mathbf{z} in dataset \mathcal{Z} ,

$$\|\nabla \ell(\mathbf{x}; \mathbf{z}) - \nabla \ell(\mathbf{y}; \mathbf{z})\|_2 \leq L \|\mathbf{x} - \mathbf{y}\|_2.$$

Note that the gradient clipping operator $\text{Clip}_\tau(\cdot)$ is utilized to ensure gradients are bounded. In addition, the boundedness of the gradient ensures the application of differentially-private mechanisms. However, stochastic gradients at different agents lose correct scaling after clipping, which breaks the stationary point property at local minima and introduces bias. To simplify analysis, one assumption that has been adopted widely in theoretical analysis [LZLC22] is the following bounded gradient assumption.

Assumption 2 (*Bounded gradient*). For any $\mathbf{x} \in \mathbb{R}^d$ and any datum \mathbf{z} in dataset \mathcal{Z} , $\|\nabla \ell(\mathbf{x}; \mathbf{z})\|_2 \leq \tau$.

Under Assumption 2, PORTER can skip the clipping operator, and $\mathbf{g}_{\tau_i}^{(t)}$ becomes an unbiased estimator of local gradient $\nabla f_i(\mathbf{x}_i^{(t)})$, while still allowing privacy analysis. However, Assumption 2 is rather strong and seldomly met in practice. For example, the gradient of a quadratic loss function is not bounded. In addition, it may result in an overly pessimistic clipping operation when there are possibly adversarial gradients with large norms in the samples. Going beyond the strong bounded gradient assumption, we consider a much milder assumption that bounds the local variance as follows, which is more standard in the analysis of unclipped stochastic algorithms.

Assumption 3 (*Bounded local variance*). For any $\mathbf{x} \in \mathbb{R}^d$ and $i \in [n]$,

$$\mathbb{E}_{\mathbf{z} \sim \mathcal{Z}_i} \|\nabla \ell(\mathbf{x}; \mathbf{z}) - \nabla f_i(\mathbf{x})\|_2^2 \leq \sigma_g^2.$$

An additional challenge is associated with dealing with the decentralized environment, where the local objective functions can be rather distinct from the global one. Our analysis identifies the following assumption, called bounded gradient dissimilarity, which says that the difference between the local gradient and the global gradient should be small relative to the global one.

Assumption 4 (Bounded gradient dissimilarity). *For any $\mathbf{x} \in \mathbb{R}^d$ and $i \in [n]$,*

$$\|\nabla f(\mathbf{x}) - \nabla f_i(\mathbf{x})\|_2 \leq \frac{1}{12} \|\nabla f(\mathbf{x})\|_2.$$

4.2 Privacy and convergence guarantees of PORTER-DP

We start by analyzing the privacy and convergence guarantees of PORTER-DP, assuming the batch size $b = 1$.

Privacy guarantee. Theorem 1 proves that PORTER-DP is (ϵ, δ) -LDP when setting the variance of Gaussian perturbation properly.

Theorem 1 (Local differential privacy). *Let $\phi_m = \frac{\sqrt{d \log(1/\delta)}}{m\epsilon}$ and $b = 1$. For any $\epsilon \leq T/m^2$ and $\delta \in (0, 1)$, PORTER-DP is (ϵ, δ) -LDP after T iterations if we set*

$$\sigma_p^2 = \frac{T\tau^2 \log(1/\delta)}{m^2\epsilon^2} = T\tau^2\phi_m^2/d. \quad (5)$$

Theorem 1 shows that PORTER-DP can achieve (ϵ, δ) -LDP regardless of whether the bounded gradient assumption presents, because using the clipping operator $\text{Clip}_\tau(\cdot)$ can guarantee all the stochastic gradients' ℓ_2 norms are bounded by τ , so that the perturbation variance can be set accordingly.

Convergence with bounded gradient assumption. We start by analyzing the convergence when the gradients are bounded under Assumption 2, in which case PORTER-DP can omit the clipping operator. Theorem 2 presents the convergence result of PORTER-DP using general compression operators (Definition 3).

Theorem 2 (Convergence of PORTER-DP with bounded gradient assumption). *Assume Assumptions 1 and 2 hold, and use general compression operators (Definition 3). Let $\Delta = \mathbb{E}[f(\bar{\mathbf{x}}^{(0)})] - f^*$. Set $\gamma = O((1 - \alpha)\rho)$, $\eta = O(\gamma^{4/3}\rho^{4/3}\phi_m)$, $T = \phi_m^{-2}$, $b = 1$ and $\sigma_p^2 = T\tau^2\phi_m^2/d$. PORTER-DP converges in average squared ℓ_2 gradient norm as*

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2^2 \lesssim \frac{\phi_m}{\rho^{4/3}(1 - \alpha)^{8/3}} \cdot \max\{\tau^2, L\Delta\}. \quad (6)$$

Theorem 2 shows the convergence error of the squared ℓ_2 gradient norm with explicit dependency on the compression ratio ρ and mixing rate α . When we fix ρ and α , Theorem 2 reaches an $O(\phi_m)$ final average squared ℓ_2 gradient norm, which matches the result of SoteriaFL-SGD [LZLC22], the state-of-the-art stochastic algorithm with local differential privacy guarantees and *unbiased* communication compression in the server-client setting. However, due to extra complexities induced by the decentralized setting and *biased* compression, PORTER-DP takes $O(\phi_m^{-2})$ iterations to converge while SoteriaFL-SGD only takes $O(\phi_m^{-1})$ iterations; in addition, PORTER-DP has a slightly worse dependency on the compression ratio ρ .

Convergence with bounded gradient assumption. A more interesting and challenging scenario is when Assumption 2 does not hold, PORTER-DP applies gradient clipping to ensure gradients are bounded to suit the privacy constraints. Fortunately, Theorem 3 describes the convergence behavior of Algorithm 1 in this case, under the much weaker bounded local variance and bounded dissimilarity assumptions.

Theorem 3 (Convergence of PORTER-DP without bounded gradient assumption). *Assume Assumptions 1, 3 and 4 hold, and use general compression operators (Definition 3). Let $\Delta = \mathbb{E}[f(\bar{\mathbf{x}}^{(0)})] - f^*$. Set $\tau = \max\{365\rho^{-4/3}(1 - \alpha)^{-8/3}\phi_m^{1/2}, 24\sigma_g\}$, $\gamma = O((1 - \alpha)\rho)$, $\eta = O(L^{-1})$, $b = 1$, $T = \phi_m^{-2}$ and $\sigma_p^2 = T\tau^2\phi_m^2/d$. Algorithm 1 converges in minimum ℓ_2 gradient norm as*

$$\min_{t \in [T]} \mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 \lesssim \max\{\rho^{-4/3}(1 - \alpha)^{-8/3}(L\Delta\phi_m)^{1/2}, \sigma_g\}. \quad (7)$$

Theorem 3 shows PORTER-DP converges in minimum ℓ_2 gradient norm with explicit dependency on compression ratio ρ , mixing rate α and gradient variance σ_g , under much weaker assumptions. To compare with Theorem 2, we can take the square root of (6), which translates to minimum ℓ_2 gradient norm convergence on the order of $O(\rho^{-2/3}(1-\alpha)^{-4/3}\phi_m^{1/2} \cdot \max\{\tau, (L\Delta)^{1/2}\})$. In comparison, although Theorem 3 has worse dependency on compression ratio ρ and mixing rate α , it matches the dependency on the baseline privacy loss ϕ_m .

4.3 Convergence guarantees of PORTER-GC

Theorem 4 further establishes the convergence of PORTER-GC without the bounded gradient assumption, which applies the clipping operator to mini-batch stochastic gradients without privacy perturbation.

Theorem 4 (Convergence of PORTER-GC without bounded gradient assumption). *Assume Assumptions 1, 3 and 4 hold, and use general compression operators (Definition 3). Let $\Delta = \mathbb{E}[f(\bar{\mathbf{x}}^{(0)})] - f^*$. Set $\tau = O(\rho^{-\frac{2}{3}}(1-\alpha)^{-\frac{4}{3}}T^{-1/2})$, $\gamma = O((1-\alpha)\rho)$, $\eta = O(L^{-1})$ and $b = O(\sigma_g^2\nu^{-2})$. PORTER-GC converges in minimum ℓ_2 gradient norm as*

$$\min_{t \in [T]} \mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 \lesssim \frac{1}{\rho^{\frac{2}{3}}(1-\alpha)^{\frac{4}{3}}} \cdot \frac{1}{T^{1/2}}.$$

Theorem 4 suggests that by picking the clipping threshold τ and batch size b properly, PORTER-GC converges at an $O(1/T^{1/2})$ rate. In comparison, standard centralized SGD converges in average squared ℓ_2 gradient norm at an $O(1/T)$ rate, which also translates to a minimum ℓ_2 gradient norm convergence in the form of $\min_{t \in [T]} \mathbb{E} \|\nabla f(\mathbf{x}^{(t)})\|_2 \lesssim 1/T^{1/2}$. Therefore, using gradient clipping for decentralized SGD does not affect the convergence rate, providing proper hyper parameter choices.

When the gradients are bounded, we can omit the clipping operator in PORTER-GC, which become the same as BEER [ZLL⁺22]. Recall that BEER guarantees a minimum ℓ_2 gradient norm convergence at the rate $\frac{1}{\rho^{1/2}(1-\alpha)^{3/2}} \cdot \frac{1}{T^{1/2}}$. In comparison, Theorem 4 has a better dependency on the mixing rate α , but has a slightly worse dependency on the compression ratio ρ , which again emphasizes that gradient clipping does not harm convergence.

5 Numerical experiments

This section presents numerical experiments to examine the performance of PORTER-DP, with comparison to the state-of-the-art server/client private stochastic algorithm SoteriaFL-SGD, which also utilizes communication compression and guarantees local differential privacy. More specifically, we evaluate the convergence of utility and accuracy in terms of communication rounds and communication bits, to analyze the privacy-utility-communication trade-offs of different algorithms.

For all experiments, we split shuffled datasets evenly to 10 agents that are connected by an Erdős-Rényi random graph with connecting probability $p = 0.8$. We use the FDLA matrix [XB04] as the mixing matrix to perform weighted information aggregation to accelerate convergence. We use biased random sparsification (cf. Example 1) for all algorithms where $k = \lfloor \frac{d}{20} \rfloor$, i.e., the compressor randomly selects 5% elements from each vector. We also apply gradient clipping with $\tau = 1$ to all algorithms for simplicity. For each experiment, all algorithms are initialized to the same starting point, and use best-tuned learning rates, batch size 1 and $\sigma_p = \frac{\tau\sqrt{T\log(1/\delta)}}{m\epsilon}$.

5.1 Logistic regression with nonconvex regularization

We run experiments on logistic regression with nonconvex regularization on the a9a dataset [CL11]. Following [WJZ⁺19], the objective function is defined as

$$\ell(\mathbf{x}; \{\mathbf{f}, l\}) = \log(1 + l \exp(-\mathbf{x}^\top \mathbf{f})) + \lambda \sum_{i=1}^d \frac{x_i^2}{1 + x_i^2},$$

where $\{\mathbf{f}, l\}$ represents a training tuple, $\mathbf{f} \in \mathbb{R}^d$ is the feature vector and $l \in \{0, 1\}$ is the label, and λ is the regularization parameter which is set to $\lambda = 0.2$.




Figure 2: The train utility and test accuracy vs. communication rounds for logistic regression with nonconvex regularization on the a9a dataset when guaranteeing $(10^{-2}, 10^{-3})$ -LDP and $(10^{-1}, 10^{-3})$ -LDP, respectively. Both PORTER-DP and SoteriaFL-SGD employ random_{162} compression (cf. Example 1).

Figure 2 shows the convergence results of PORTER-DP and SoteriaFL-SGD for logistic regression with nonconvex regularization on the a9a dataset to reach $(10^{-2}, 10^{-3})$ -LDP and $(10^{-1}, 10^{-3})$ -LDP, respectively. Under $(10^{-2}, 10^{-3})$ -LDP, which is a stricter privacy setting, PORTER-DP converges faster than SoteriaFL-SGD in test accuracy, while converges slightly slower in train utility. Under $(10^{-1}, 10^{-3})$ -LDP, PORTER-DP performs slightly worse than SoteriaFL-SGD. Given that PORTER-DP operates under the decentralized topology with much weaker information exchange, the results highlight PORTER-DP’s communication efficiency by showing it can achieve similar performance as its server/client counterpart, i.e. SoteriaFL-SGD, especially under strict privacy constraints.

5.2 One-hidden-layer neural network training

We evaluate PORTER-DP by training a one-hidden layer neural network on the MNIST dataset [LJB⁺95]. The network uses 64 hidden neurons, sigmoid activation functions and cross-entropy loss, where the loss function over a training pair $\{\mathbf{f}, l\}$ is defined as

$$\ell(\mathbf{x}; (\mathbf{f}, l)) = \text{CrossEntropy}(\text{softmax}(\mathbf{W}_2 \text{sigmoid}(\mathbf{W}_1 \mathbf{f} + \mathbf{c}_1) + \mathbf{c}_2), l).$$

Here, the model parameter is defined by $\mathbf{x} = \text{vec}(\mathbf{W}_1, \mathbf{c}_1, \mathbf{W}_2, \mathbf{c}_2)$, where the dimensions of the network parameters $\mathbf{W}_1, \mathbf{c}_1, \mathbf{W}_2, \mathbf{c}_2$ are $64 \times 784, 64 \times 1, 10 \times 64,$ and 10×1 , respectively.




Figure 3: The train utility and test accuracy vs. communication rounds for training a one-hidden-layer neural network on the MNIST dataset when guaranteeing $(10^{-2}, 10^{-3})$ -LDP and $(10^{-1}, 10^{-3})$ -LDP, respectively. Both PORTER-DP and SoteriaFL-SGD employ random_{2583} compression (cf. Example 1).

Figure 3 shows the convergence results of PORTER-DP and SoteriaFL-SGD for training a one-hidden-layer neural network on the MNIST dataset to reach $(10^{-2}, 10^{-3})$ -LDP and $(10^{-1}, 10^{-3})$ -LDP, respectively. Under both privacy settings, PORTER-DP converges at a similar rate as SoteriaFL-SGD in train utility. However, in terms of convergence in test accuracy, PORTER-DP outperforms SoteriaFL-SGD under the stricter

$(10^{-2}, 10^{-3})$ -LDP, while the two algorithms have similar performance under the other setting. This experiment again emphasizes PORTER-DP’s communication efficiency in comparison to the server/client algorithm SoteriaFL-SGD.

6 Conclusions

In this paper, we propose an algorithmic framework called PORTER, which incorporates practically-relevant gradient clipping and communication compression simultaneously in the design of decentralized nonconvex optimization algorithms. We propose two variants: PORTER-DP and PORTER-GC. While they share a similar structure that makes use of gradient tracking, communication compression, and error feedback, their focuses are on different perspectives motivated by applications in privacy preserving and neural network training, respectively. PORTER-DP adds privacy perturbation to clipped gradients to protect the local differential privacy of each agent, with explicit utility and communication complexities. PORTER applies gradient clipping to mini-batch stochastic gradients, which converges in minimum ℓ_2 gradient norm at similar rate as centralized SGD without clipping under proper choices of hyperparameters. The development of PORTER offers a simple analysis framework to understand gradient clipping in decentralized nonconvex optimization without bounded gradient assumptions, highlighting the potential of achieving both communication efficiency and privacy preserving in the decentralized framework.

Acknowledgements

The work of B. Li and Y. Chi is supported in part by ONR under the grant N00014-19-1-2404, by AFRL under FA8750-20-2-0504, and by NSF under the grants CCF-1901199, CCF-2007911 and CNS-2148212. B. Li is also gratefully supported by the Wei Shen and Xuehong Zhang Presidential Fellowship at Carnegie Mellon University.

References

- [ACCÖ21] S. Asoodeh, W.-N. Chen, F. P. Calmon, and A. Özgür. Differentially private federated learning: An information-theoretic perspective. In *2021 IEEE International Symposium on Information Theory (ISIT)*, pages 344–349, July 2021.
- [ACG⁺16] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 308–318, October 2016.
- [AGL⁺17] D. Alistarh, D. Grubic, J. Li, R. Tomioka, and M. Vojnovic. QSGD: Communication-efficient sgd via gradient quantization and encoding. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [AHJ⁺18] D. Alistarh, T. Hoefler, M. Johansson, N. Konstantinov, S. Khirirat, and C. Renggli. The convergence of sparsified gradient methods. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- [ALBR19] M. Assran, N. Loizou, N. Ballas, and M. Rabbat. Stochastic gradient push for distributed deep learning. In *International Conference on Machine Learning*, pages 344–353. PMLR, 2019.
- [BBP13] Y. Bengio, N. Boulanger-Lewandowski, and R. Pascanu. Advances in optimizing recurrent networks. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 8624–8628, May 2013.
- [BJ13] P. Bianchi and J. Jakubowicz. Convergence of a multi-agent projected stochastic gradient algorithm for non-convex optimization. *IEEE Transactions on Automatic Control*, 58(2):391–405, February 2013.

- [CABP13] K. Chatzikokolakis, M. E. Andrés, N. E. Bordenabe, and C. Palamidessi. Broadening the scope of differential privacy using metrics. In E. De Cristofaro and M. Wright, editors, *Privacy Enhancing Technologies*, Lecture Notes in Computer Science, pages 82–102, Berlin, Heidelberg, 2013. Springer.
- [CL11] C.-C. Chang and C.-J. Lin. LIBSVM: a library for support vector machines. *ACM transactions on intelligent systems and technology (TIST)*, 2(3):1–27, May 2011.
- [CWH20] X. Chen, S. Z. Wu, and M. Hong. Understanding gradient clipping in private SGD: A geometric perspective. In *Advances in Neural Information Processing Systems*, volume 33, pages 13773–13782. Curran Associates, Inc., 2020.
- [DJW13] J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Local privacy and statistical minimax rates. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 429–438, October 2013.
- [DKX⁺22] R. Das, S. Kale, Z. Xu, T. Zhang, and S. Sanghavi. Beyond uniform Lipschitz condition in differentially private optimization. *arXiv preprint arXiv:2206.10713*, 2022.
- [DLC⁺22] J. Du, S. Li, X. Chen, S. Chen, and M. Hong. Dynamic differential-privacy preserving sgd. *arXiv preprint arXiv:2111.00173*, 2022.
- [DLS16] P. Di Lorenzo and G. Scutari. Next: In-network nonconvex optimization. *IEEE Transactions on Signal and Information Processing over Networks*, 2(2):120–136, 2016.
- [DMNS06] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In S. Halevi and T. Rabin, editors, *Theory of Cryptography*, Lecture Notes in Computer Science, pages 265–284, Berlin, Heidelberg, 2006. Springer.
- [DSZOR15] C. M. De Sa, C. Zhang, K. Olukotun, and C. Ré. Taming the wild: A unified analysis of Hogwild-style algorithms. *Advances in Neural Information Processing Systems*, 28, 2015.
- [Dwo08] C. Dwork. Differential privacy: A survey of results. In M. Agrawal, D. Du, Z. Duan, and A. Li, editors, *Theory and Applications of Models of Computation*, Lecture Notes in Computer Science, pages 1–19, Berlin, Heidelberg, 2008. Springer.
- [FKT20] V. Feldman, T. Koren, and K. Talwar. Private stochastic convex optimization: Optimal rates in linear time. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2020, pages 439–449, New York, NY, USA, June 2020. Association for Computing Machinery.
- [FLFL23] H. Fang, X. Li, C. Fan, and P. Li. Improved convergence of differential private SGD with gradient clipping. In *The Eleventh International Conference on Learning Representations*, 2023.
- [FSG⁺21] I. Fatkhullin, I. Sokolov, E. Gorbunov, Z. Li, and P. Richtárik. EF21 with bells & whistles: Practical algorithmic extensions of modern error feedback. *arXiv preprint arXiv:2110.03294*, 2021.
- [GBLR21] E. Gorbunov, K. P. Burlachenko, Z. Li, and P. Richtarik. MARINA: Faster non-convex distributed learning with compression. In *Proceedings of the 38th International Conference on Machine Learning*, pages 3788–3798. PMLR, July 2021.
- [HDJ⁺20] X. Huang, Y. Ding, Z. L. Jiang, S. Qi, X. Wang, and Q. Liao. DP-FL: A novel differentially private federated learning framework for the unbalanced data. *World Wide Web*, 23(4):2529–2545, July 2020.
- [HP23] K. Huang and S. Pu. CEDAS: A compressed decentralized stochastic gradient method with improved convergence. *arXiv preprint arXiv:2301.05872*, 2023.

- [HSZ⁺22] Y. Huang, Y. Sun, Z. Zhu, C. Yan, and J. Xu. Tackling data heterogeneity: A new unified framework for decentralized sgd with sample-induced topology. In *Proceedings of the 39th International Conference on Machine Learning*, pages 9310–9345. PMLR, June 2022.
- [INS⁺19] R. Iyengar, J. P. Near, D. Song, O. Thakkar, A. Thakurta, and L. Wang. Towards practical differentially private convex optimization. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 299–316, May 2019.
- [KDG03] D. Kempe, A. Dobra, and J. Gehrke. Gossip-based computation of aggregate information. In *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings.*, pages 482–491, October 2003.
- [KFI17] S. Kanai, Y. Fujiwara, and S. Iwamura. Preventing gradient explosions in gated recurrent units. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [KHS23] A. Koloskova, H. Hendrikx, and S. U. Stich. Revisiting gradient clipping: Stochastic bias and tight convergence guarantees. In *Proceedings of the 40th International Conference on Machine Learning*, Proceedings of Machine Learning Research. PMLR, 23–29 Jul 2023.
- [KLL16] J. Kim, J. K. Lee, and K. M. Lee. Accurate image super-resolution using very deep convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1646–1654, 2016.
- [KSJ19] A. Koloskova, S. Stich, and M. Jaggi. Decentralized stochastic optimization and gossip algorithms with compressed communication. In K. Chaudhuri and R. Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 3478–3487. PMLR, 09–15 Jun 2019.
- [LCCC20] B. Li, S. Cen, Y. Chen, and Y. Chi. Communication-efficient distributed optimization in networks with gradient tracking and variance reduction. *Journal of Machine Learning Research*, 21(180):1–51, 2020.
- [LJB⁺95] Y. LeCun, L. D. Jackel, L. Bottou, C. Cortes, J. S. Denker, H. Drucker, I. Guyon, U. A. Muller, E. Sackinger, and P. Simard. Learning algorithms for classification: A comparison on handwritten digit recognition. *Neural networks: the statistical mechanics perspective*, 261(276):2, 1995.
- [LKQR20] Z. Li, D. Kovalev, X. Qian, and P. Richtarik. Acceleration for compressed gradient descent in distributed and federated optimization. In *Proceedings of the 37th International Conference on Machine Learning*, pages 5895–5904. PMLR, November 2020.
- [LLC22] B. Li, Z. Li, and Y. Chi. DESTRESS: Computation-optimal and communication-efficient decentralized nonconvex finite-sum optimization. *SIAM Journal on Mathematics of Data Science*, 4(3):1031–1051, September 2022.
- [LLP23] Y. Liao, Z. Li, and S. Pu. A linearly convergent robust compressed Push-Pull method for decentralized optimization. *arXiv preprint arXiv:2303.07091*, 2023.
- [LR21] Z. Li and P. Richtarik. CANITA: Faster rates for distributed convex optimization with communication compression. In *Advances in Neural Information Processing Systems*, volume 34, pages 13770–13781. Curran Associates, Inc., 2021.
- [LY22] L. Luo and H. Ye. An optimal stochastic algorithm for decentralized nonconvex finite-sum optimization. *arXiv preprint arXiv:2210.13931*, 2022.
- [LZLC22] Z. Li, H. Zhao, B. Li, and Y. Chi. SoteriaFL: A unified framework for private federated learning with communication compression. In *Advances in Neural Information Processing Systems*, volume 35, pages 4285–4300. Curran Associates, Inc., 2022.

- [LZZ⁺17] X. Lian, C. Zhang, H. Zhang, C.-J. Hsieh, W. Zhang, and J. Liu. Can decentralized algorithms outperform centralized algorithms? a case study for decentralized parallel stochastic gradient descent. In *Advances in Neural Information Processing Systems*, pages 5330–5340, 2017.
- [MGTR19] K. Mishchenko, E. Gorbunov, M. Takáč, and P. Richtárik. Distributed learning with compressed gradient differences. *arXiv preprint arXiv:1901.09269*, 2019.
- [MS23] T. Murata and T. Suzuki. DIFF2: Differential private optimization via gradient differences for nonconvex distributed learning. *arXiv preprint arXiv:2302.03884*, 2023.
- [NBD22] M. Noble, A. Bellet, and A. Dieuleveut. Differentially private federated learning on heterogeneous data. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, pages 10110–10145. PMLR, May 2022.
- [NOR18] A. Nedić, A. Olshevsky, and M. G. Rabbat. Network topology and communication-computation tradeoffs in decentralized optimization. *Proceedings of the IEEE*, 106(5):953–976, 2018.
- [NOS17] A. Nedić, A. Olshevsky, and W. Shi. Achieving geometric convergence for distributed optimization over time-varying graphs. *SIAM Journal on Optimization*, 27(4):2597–2633, 2017.
- [NRB20] M. Nokleby, H. Raja, and W. U. Bajwa. Scaling-Up Distributed Processing of Data Streams for Machine Learning. *Proceedings of the IEEE*, 108(11):1984–2012, November 2020.
- [PMB13] R. Pascanu, T. Mikolov, and Y. Bengio. On the difficulty of training recurrent neural networks. In *Proceedings of the 30th International Conference on Machine Learning*, pages 1310–1318. PMLR, May 2013.
- [QL18] G. Qu and N. Li. Harnessing smoothness to accelerate distributed optimization. *IEEE Transactions on Control of Network Systems*, 5(3):1245–1260, 2018.
- [RLDJ23] A. Reiszadeh, H. Li, S. Das, and A. Jadbabaie. Variance-reduced clipping for non-convex optimization. *arXiv preprint arXiv:2303.00883*, 2023.
- [RSF21] P. Richtarik, I. Sokolov, and I. Fatkhullin. EF21: A new, simpler, theoretically better, and practically faster error feedback. In *Advances in Neural Information Processing Systems*, volume 34, pages 4384–4396. Curran Associates, Inc., 2021.
- [SCJ18] S. U. Stich, J.-B. Cordonnier, and M. Jaggi. Sparsified SGD with memory. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- [SDGD21] N. Singh, D. Data, J. George, and S. Diggavi. SQuARM-SGD: Communication-Efficient Momentum SGD for Decentralized Optimization. *IEEE Journal on Selected Areas in Information Theory*, 2(3):954–969, September 2021.
- [SFD⁺14] F. Seide, H. Fu, J. Droppo, G. Li, and D. Yu. 1-bit stochastic gradient descent and its application to data-parallel distributed training of speech DNNs. In *Interspeech 2014*, pages 1058–1062. ISCA, September 2014.
- [Sha07] D. Shah. Gossip algorithms. *Foundations and Trends[®] in Networking*, 3(1):1–125, 2007.
- [SLH20] H. Sun, S. Lu, and M. Hong. Improving the sample and communication complexity for decentralized non-convex optimization: Joint gradient estimation and tracking. In *International Conference on Machine Learning*, pages 9217–9228. PMLR, 2020.
- [TGZ⁺18] H. Tang, S. Gan, C. Zhang, T. Zhang, and J. Liu. Communication compression for decentralized training. *Advances in Neural Information Processing Systems*, 31:7652–7662, 2018.
- [TLQ⁺19] H. Tang, X. Lian, S. Qiu, L. Yuan, C. Zhang, T. Zhang, and J. Liu. Deepsqueeze: Decentralization meets error-compensated compression. *arXiv preprint arXiv:1907.07346*, 2019.

- [TMHP20] H. Taheri, A. Mokhtari, H. Hassani, and R. Pedarsani. Quantized decentralized stochastic learning over directed graphs. In *Proceedings of the 37th International Conference on Machine Learning*, pages 9324–9333. PMLR, November 2020.
- [WJ21] J. Wang and G. Joshi. Cooperative SGD: A unified framework for the design and analysis of local-update SGD algorithms. *The Journal of Machine Learning Research*, 22(1):213:9709–213:9758, January 2021.
- [WJEG19] L. Wang, B. Jayaraman, D. Evans, and Q. Gu. Efficient privacy-preserving stochastic nonconvex optimization. *arXiv preprint arXiv:1910.13659*, 2019.
- [WJZ⁺19] Z. Wang, K. Ji, Y. Zhou, Y. Liang, and V. Tarokh. SpiderBoost and momentum: Faster variance reduction algorithms. In *Advances in Neural Information Processing Systems*, pages 2406–2416, 2019.
- [WXDX20] D. Wang, H. Xiao, S. Devadas, and J. Xu. On differentially private stochastic convex optimization with heavy-tailed data. In *Proceedings of the 37th International Conference on Machine Learning*, pages 10081–10091. PMLR, November 2020.
- [WYX17] D. Wang, M. Ye, and J. Xu. Differentially private empirical risk minimization revisited: Faster and more general. *Advances in Neural Information Processing Systems*, 30, 2017.
- [XB04] L. Xiao and S. Boyd. Fast linear iterations for distributed averaging. *Systems & Control Letters*, 53(1):65–78, September 2004.
- [XKK22a] R. Xin, U. A. Khan, and S. Kar. Fast decentralized nonconvex finite-sum optimization with recursive variance reduction. *SIAM Journal on Optimization*, 32(1):1–28, March 2022.
- [XKK22b] R. Xin, U. A. Khan, and S. Kar. A Fast Randomized Incremental Gradient Method for Decentralized Nonconvex Optimization. *IEEE Transactions on Automatic Control*, 67(10):5150–5165, October 2022.
- [XPNK20] R. Xin, S. Pu, A. Nedić, and U. A. Khan. A general framework for decentralized optimization with first-order methods. *Proceedings of the IEEE*, 108(11):1869–1889, 2020.
- [XYD19] H. Xiao, Y. Ye, and S. Devadas. Local differential privacy in decentralized optimization. *arXiv preprint arXiv:1902.06101*, 2019.
- [YCC⁺23] Y. Yan, J. Chen, P.-Y. Chen, X. Cui, S. Lu, and Y. Xu. Compressed decentralized proximal stochastic gradient method for nonconvex composite problems with heterogeneous data. *arXiv preprint arXiv:2302.14252*, 2023.
- [YGG17] Y. You, I. Gitman, and B. Ginsburg. Scaling SGD batch size to 32k for ImageNet training. Technical Report UCB/EECS-2017-156, EECS Department, University of California, Berkeley, Sep 2017.
- [YZCL22] X. Yang, H. Zhang, W. Chen, and T.-Y. Liu. Normalized/clipped SGD with perturbation for differentially private non-convex optimization. *arXiv preprint arXiv:2206.13033*, 2022.
- [ZBLR21] H. Zhao, K. Burlachenko, Z. Li, and P. Richtárik. Faster rates for compressed federated learning with client-variance reduction. *arXiv preprint arXiv:2112.13097*, 2021.
- [ZCH⁺22] X. Zhang, X. Chen, M. Hong, S. Wu, and J. Yi. Understanding clipping for federated learning: Convergence and client-level differential privacy. In *Proceedings of the 39th International Conference on Machine Learning*, pages 26048–26067. PMLR, June 2022.
- [ZHSJ20a] J. Zhang, T. He, S. Sra, and A. Jadbabaie. Why gradient clipping accelerates training: A theoretical justification for adaptivity. In *International Conference on Learning Representations*, 2020.

- [ZHSJ20b] J. Zhang, T. He, S. Sra, and A. Jadbabaie. Why gradient clipping accelerates training: A theoretical justification for adaptivity. In *International Conference on Learning Representations*, March 2020.
- [ZJFW20] B. Zhang, J. Jin, C. Fang, and L. Wang. Improved analysis of clipping algorithms for non-convex optimization. *Advances in Neural Information Processing Systems*, 33:15511–15521, 2020.
- [ZKV⁺20] J. Zhang, S. P. Karimireddy, A. Veit, S. Kim, S. J. Reddi, S. Kumar, and S. Sra. Why are adaptive methods good for attention models? *arXiv preprint arXiv:1912.03194*, October 2020.
- [ZLL⁺22] H. Zhao, B. Li, Z. Li, P. Richtárik, and Y. Chi. Beer: Fast $O(1/T)$ rate for decentralized nonconvex optimization with communication compression. *Advances in Neural Information Processing Systems*, 35, 2022.
- [ZM10] M. Zhu and S. Martínez. Discrete-time dynamic average consensus. *Automatica*, 46(2):322–329, 2010.

A Proof of Theorem 1

This section proves Theorem 1 in the following steps: 1) define privacy loss and moment generating function, 2) define mechanisms and sub-mechanisms, and 3) bound the overall moment generating function and show the choice of perturbation variance satisfies all conditions.

Moment generating function. Let o and aux denote an outcome and an auxiliary input, respectively. Then, we can define the privacy loss of an outcome o on neighboring dataset \mathbf{Z} and $\mathbf{Z}^{(i)}$ as

$$c(o; \mathcal{M}, \text{aux}, \mathbf{Z}, \mathbf{Z}^{(i)}) = \log \frac{\mathbb{P}(\mathcal{M}(\text{aux}, \mathbf{Z}) = o)}{\mathbb{P}(\mathcal{M}(\text{aux}, \mathbf{Z}^{(i)}) = o)},$$

and its log moment generating functions as

$$\alpha_i^{\mathcal{M}}(\lambda; \text{aux}, \mathbf{Z}, \mathbf{Z}^{(i)}) = \log \mathbb{E}_{o \sim \mathcal{M}(\text{aux}, \mathbf{Z})} [\exp(\lambda c(o; \mathcal{M}, \text{aux}, \mathbf{Z}, \mathbf{Z}^{(i)}))].$$

Taking maximum over conditions, the unconditioned log moment generating function is

$$\hat{\alpha}_i^{\mathcal{M}}(\lambda) = \max_{\text{aux}, \mathbf{Z}, \mathbf{Z}^{(i)}} \alpha_i^{\mathcal{M}}(\lambda; \text{aux}, \mathbf{Z}, \mathbf{Z}^{(i)}).$$

Sub-mechanisms. Definition 4 defines the LDP mechanism, but it is not enough to model decentralized algorithms. To model the perturbation operation happens on agent i at time t , we define a sub-mechanism as $\mathcal{M}_i^{(t)} : \mathcal{D} \rightarrow \mathcal{R}$, where $i \in [n], t \in [T]$, which can be understood as the perturbation added on agent i at time t . In addition, we define another mechanism $\mathcal{C} : \mathcal{R} \rightarrow \mathcal{R}$ to model the compression operator and $\mathcal{C} \circ \mathcal{M}_i^{(t)}$ to represent the full update at an agent, and use \mathcal{M} to represent the full algorithm.

Proof of LDP. The overall log moment generating function for agent i can be bounded using [LZLC22, Lemma 2] as

$$\hat{\alpha}_i^{\mathcal{M}}(\lambda) \leq \sum_{t=1}^T \hat{\alpha}_i^{\mathcal{C} \circ \mathcal{M}_i^{(t)}}(\lambda) \leq \sum_{t=1}^T \hat{\alpha}_i^{\mathcal{M}_i^{(t)}}(\lambda).$$

Let $q = \frac{1}{m}$ denote the probability each data sample is chosen. For agent i and $\lambda > 0$, assume $q \leq \frac{\tau}{16\sigma_p}$ and $\lambda \leq \frac{\sigma_p^2}{\tau^2} \log \frac{\tau}{q\sigma_p}$. We can apply [ACG⁺16, Lemma 3] to bound each $\hat{\alpha}_i^{\mathcal{M}_i^{(t)}}(\lambda)$ as

$$\hat{\alpha}_i^{\mathcal{M}_i^{(t)}}(\lambda) \leq \frac{q^2 \lambda (\lambda + 1) \tau^2}{(1 - q) \sigma_p^2} + O\left(\frac{q^3 \lambda^3 \tau^3}{\sigma_p^3}\right) = O\left(\frac{q^2 \lambda^2 \tau^2}{\sigma_p^2}\right).$$

To conclude the proof, we can verify there exists some λ that satisfies the following inequalities when choosing $\sigma_p = \frac{\tau q \sqrt{T \log(1/\delta)}}{\epsilon}$ and $q = \frac{1}{m}$, namely,

$$\begin{aligned} \left(\frac{Tq\tau\lambda}{\sigma_p}\right)^2 &\leq \frac{\lambda\epsilon}{2}, \\ \exp(-\lambda\epsilon/2) &\leq \delta, \\ \lambda &\leq \frac{\sigma_p^2}{\tau^2} \log \frac{\tau}{q\sigma_p}. \end{aligned}$$

B Proof of Theorem 2

This section proves Theorem 2 in the following 4 subsections: Appendix B.1 derives the descent inequality, Appendices B.2 and B.3 create two linear systems to bound the sum of consensus errors in the descent inequality, and finally Appendix B.4 specifies hyper parameters to obtain convergence rate.

With Assumption 2, we can skip the compression operator, in PORTER-DP. To reuse this section's results in later analysis, we assume Assumption 3 in deriving descent lemma and linear systems, and lift this assumption when computing convergence rate in Appendix B.4 using $\sigma_g \leq 2\tau$.

Denote

$$\mathbf{G}^{(t)} = \frac{1}{b} \sum_{\mathbf{Z} \in \mathcal{Z}^{(t)}} \nabla \ell(\mathbf{X}^{(t-1)}; \mathbf{Z}).$$

B.1 Function value descent

Using Taylor expansion, and taking expectation conditioned on time t ,

$$\begin{aligned} \mathbb{E}_t[f(\bar{\mathbf{x}}^{(t+1)}) - f(\bar{\mathbf{x}}^{(t)})] &\leq \mathbb{E}_t \langle \nabla f(\bar{\mathbf{x}}^{(t)}), -\eta \bar{\mathbf{v}}^{(t+1)} \rangle + \frac{L}{2} \mathbb{E}_t \|\eta \bar{\mathbf{v}}^{(t+1)}\|_2^2 \\ &= -\eta \langle \nabla f(\bar{\mathbf{x}}^{(t)}), \mathbb{E}_t[\bar{\mathbf{v}}^{(t+1)}] \rangle + \frac{\eta^2 L}{2} \mathbb{E}_t \|\bar{\mathbf{v}}^{(t+1)}\|_2^2 \\ &= -\eta \langle \nabla f(\bar{\mathbf{x}}^{(t)}), \mathbb{E}_t[\bar{\mathbf{g}}_\tau^{(t+1)} + \bar{\mathbf{e}}^{(t+1)}] \rangle + \frac{\eta^2 L}{2} \mathbb{E}_t \|\bar{\mathbf{v}}^{(t+1)}\|_2^2, \end{aligned}$$

where the last equality is due to $\bar{\mathbf{v}}^{(t)} = \bar{\mathbf{g}}_p^{(t)}$ that can be proved by induction.

Because $\mathbb{E}_t[\mathbf{e}_i^{(t)}] = \mathbf{0}_d$ and stochastic gradients are unbiased,

$$\begin{aligned} &\mathbb{E}_t[f(\bar{\mathbf{x}}^{(t+1)}) - f(\bar{\mathbf{x}}^{(t)})] \\ &= -\eta \langle \nabla f(\bar{\mathbf{x}}^{(t)}), \nabla F(\mathbf{X}^{(t)}) (\frac{1}{n} \mathbf{1}_n) \rangle + \frac{\eta^2 L}{2} \mathbb{E}_t \|\bar{\mathbf{v}}^{(t+1)}\|_2^2 \\ &= \frac{\eta}{2} \left(\|\nabla f(\bar{\mathbf{x}}^{(t)}) - \nabla F(\mathbf{X}^{(t)}) (\frac{1}{n} \mathbf{1}_n)\|_2^2 - \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2^2 - \|\nabla F(\mathbf{X}^{(t)}) (\frac{1}{n} \mathbf{1}_n)\|_2^2 \right) + \frac{\eta^2 L}{2} \mathbb{E}_t \|\bar{\mathbf{v}}^{(t+1)}\|_2^2 \\ &\leq -\frac{\eta}{2} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2^2 + \frac{\eta L^2}{2n} \|\mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top\|_F^2 + \frac{\eta^2 L}{2} \mathbb{E}_t \|\bar{\mathbf{v}}^{(t+1)}\|_2^2 - \frac{\eta}{2} \|\nabla F(\mathbf{X}^{(t)}) (\frac{1}{n} \mathbf{1}_n)\|_2^2, \end{aligned} \quad (8)$$

where the last inequality is due to Assumption 1.

Let $\Delta = \mathbb{E}[f(\bar{\mathbf{x}}^{(0)})] - f^*$. Take full expectation and average (8) over $t = 1, \dots, T$, the expected utility can be bounded by

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2^2 &\leq \frac{2\Delta}{\eta T} + \frac{1}{T} \cdot \frac{L^2}{n} \sum_{t=1}^T \mathbb{E} \|\mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top\|_F^2 \\ &\quad + \frac{1}{T} \cdot \eta L \sum_{t=1}^T \mathbb{E} \|\bar{\mathbf{v}}^{(t+1)}\|_2^2 - \frac{1}{T} \sum_{t=1}^T \mathbb{E} \|\nabla F(\mathbf{X}^{(t)}) (\frac{1}{n} \mathbf{1}_n)\|_2^2. \end{aligned} \quad (9)$$

B.2 Sum of variable consensus errors

This subsection creates a linear system to bound $\sum_{t=1}^T \mathbb{E} \|\mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top\|_F^2$ by $\sum_{t=1}^T \mathbb{E} \|\mathbf{V}^{(t)} - \bar{\mathbf{v}}^{(t)} \mathbf{1}_n^\top\|_F^2$ and $\sum_{t=1}^T \mathbb{E} \|\bar{\mathbf{v}}^{(t)}\|_2^2$. To simplify notations, let $\widehat{\mathbf{W}} = \mathbf{I}_n + \gamma(\mathbf{W} - \mathbf{I}_n)$, and denote the mixing rate of $\widehat{\mathbf{W}}$ by $\hat{\alpha} = \|\widehat{\mathbf{W}} - (\frac{1}{n} \mathbf{1}_n \mathbf{1}_n^\top)\|_{\text{op}}$. Lemma 1 analyzes the mixing rate of the regularized mixing matrix.

Lemma 1 (Mixing rate of regularized mixing matrix). *Assuming $0 < \gamma \leq 1$. The mixing rate of $\widehat{\mathbf{W}}$ can be bounded as*

$$\hat{\alpha} \leq 1 + \gamma(\alpha - 1). \quad (10)$$

Proof of Lemma 1. Let $\lambda_1 = 1 > \lambda_2 \geq \dots \geq \lambda_n > -1$ denote the eigenvalues of \mathbf{W} . Corresponding eigenvalues of $\widehat{\mathbf{W}}$ are $1 + \gamma(\lambda_i - 1)$, $i = 1, \dots, n$.

The mixing rate of $\widehat{\mathbf{W}}$ is

$$\begin{aligned} \hat{\alpha} &= \max \{ |1 + \gamma(\lambda_2 - 1)|, |1 + \gamma(\lambda_n - 1)| \} \\ &\leq \max \{ |1 - \gamma| + \gamma|\lambda_2|, |1 - \gamma| + \gamma|\lambda_n| \} \\ &= 1 + \gamma(\alpha - 1). \end{aligned}$$

□

B.2.1 Variable consensus error

Take expectation conditioned on time t , and use Young's inequality, the variable consensus error can be bounded as

$$\begin{aligned} & \mathbb{E}_t \|\mathbf{X}^{(t+1)} - \bar{\mathbf{x}}^{(t+1)} \mathbf{1}_n^\top\|_F^2 \\ &= \mathbb{E}_t \left\| \left(\mathbf{X}^{(t)} + \gamma \mathbf{Q}_x^{(t+1)} (\mathbf{W} - \mathbf{I}_n) - \eta \mathbf{V}^{(t+1)} \right) \left(\mathbf{I}_n - \left(\frac{1}{n} \mathbf{1}_n \mathbf{1}_n^\top \right) \right) \right\|_F^2 \\ &= \mathbb{E}_t \left\| \left(\mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top \right) \left(\widehat{\mathbf{W}} - \left(\frac{1}{n} \mathbf{1}_n \mathbf{1}_n^\top \right) \right) + \gamma \left(\mathbf{Q}_x^{(t+1)} - \mathbf{X}^{(t)} \right) (\mathbf{W} - \mathbf{I}_n) - \eta \mathbf{V}^{(t+1)} \left(\mathbf{I}_n - \left(\frac{1}{n} \mathbf{1}_n \mathbf{1}_n^\top \right) \right) \right\|_F^2 \\ &\leq \frac{2}{1 + \hat{\alpha}^2} \left\| \left(\mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top \right) \left(\widehat{\mathbf{W}} - \left(\frac{1}{n} \mathbf{1}_n \mathbf{1}_n^\top \right) \right) \right\|_F^2 \\ &\quad + \frac{2}{1 - \hat{\alpha}^2} \mathbb{E}_t \left\| \gamma \left(\mathbf{Q}_x^{(t+1)} - \mathbf{X}^{(t)} \right) (\mathbf{W} - \mathbf{I}_n) - \eta \mathbf{V}^{(t+1)} \left(\mathbf{I}_n - \left(\frac{1}{n} \mathbf{1}_n \mathbf{1}_n^\top \right) \right) \right\|_F^2 \\ &\leq \frac{2}{1 + \hat{\alpha}^2} \left\| \left(\mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top \right) \left(\widehat{\mathbf{W}} - \left(\frac{1}{n} \mathbf{1}_n \mathbf{1}_n^\top \right) \right) \right\|_F^2 \\ &\quad + \frac{4}{1 - \hat{\alpha}^2} \mathbb{E}_t \left\| \gamma \left(\mathbf{Q}_x^{(t+1)} - \mathbf{X}^{(t)} \right) (\mathbf{W} - \mathbf{I}_n) \right\|_F^2 + \frac{4}{1 - \hat{\alpha}^2} \mathbb{E}_t \left\| \eta \mathbf{V}^{(t+1)} \left(\mathbf{I}_n - \left(\frac{1}{n} \mathbf{1}_n \mathbf{1}_n^\top \right) \right) \right\|_F^2 \\ &\stackrel{(i)}{\leq} \frac{2\hat{\alpha}^2}{1 + \hat{\alpha}^2} \left\| \mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top \right\|_F^2 + \frac{16\gamma^2}{1 - \hat{\alpha}^2} \mathbb{E}_t \left\| \mathbf{Q}_x^{(t+1)} - \mathbf{X}^{(t)} \right\|_F^2 + \frac{4\eta^2}{1 - \hat{\alpha}^2} \mathbb{E}_t \left\| \mathbf{V}^{(t+1)} - \bar{\mathbf{v}}^{(t+1)} \mathbf{1}_n^\top \right\|_F^2 \\ &\stackrel{(ii)}{\leq} \hat{\alpha} \left\| \mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top \right\|_F^2 + \frac{16(1 - \rho)\gamma^2}{1 - \hat{\alpha}} \left\| \mathbf{Q}_x^{(t)} - \mathbf{X}^{(t)} \right\|_F^2 + \frac{4\eta^2}{1 - \hat{\alpha}} \mathbb{E}_t \left\| \mathbf{V}^{(t+1)} - \bar{\mathbf{v}}^{(t+1)} \mathbf{1}_n^\top \right\|_F^2, \end{aligned} \quad (11)$$

where (i) is obtained by $\|\mathbf{W} - \mathbf{I}_n\|_{\text{op}} \leq 2$, (ii) uses $2\hat{\alpha} \leq 1 + \hat{\alpha}^2$, $1 - \hat{\alpha} \leq 1 - \hat{\alpha}^2$ and Definition 3.

B.2.2 Variable quantization error

Assume γ satisfies the following inequality (which will be verified in Appendix B.4)

$$\gamma^2 \leq \frac{\rho^2}{96(1 - \rho)}. \quad (12)$$

Taking expectation conditioned on time t , the variable quantization error can be decomposed and bounded as

$$\mathbb{E}_t \left\| \mathbf{Q}_x^{(t+1)} - \mathbf{X}^{(t+1)} \right\|_F^2$$

$$\begin{aligned}
&= \mathbb{E}_t \left\| \mathbf{Q}_x^{(t)} + \mathcal{C}(\mathbf{X}^{(t)} - \mathbf{Q}_x^{(t)}) - \mathbf{X}^{(t+1)} \right\|_F^2 \\
&= \mathbb{E}_t \left\| \mathcal{C}(\mathbf{X}^{(t)} - \mathbf{Q}_x^{(t)}) - (\mathbf{X}^{(t)} - \mathbf{Q}_x^{(t)}) - (\mathbf{X}^{(t+1)} - \mathbf{X}^{(t)}) \right\|_F^2 \\
&\stackrel{(i)}{\leq} \frac{2}{1+(1-\rho)} \mathbb{E}_t \left\| \mathcal{C}(\mathbf{X}^{(t)} - \mathbf{Q}_x^{(t)}) - (\mathbf{X}^{(t)} - \mathbf{Q}_x^{(t)}) \right\|_F^2 + \frac{2}{1-(1-\rho)} \mathbb{E}_t \left\| \mathbf{X}^{(t+1)} - \mathbf{X}^{(t)} \right\|_F^2 \\
&\stackrel{(ii)}{\leq} \frac{2(1-\rho)}{1+(1-\rho)} \left\| \mathbf{X}^{(t)} - \mathbf{Q}_x^{(t)} \right\|_F^2 + \frac{2}{\rho} \mathbb{E}_t \left\| \mathbf{X}^{(t+1)} - \mathbf{X}^{(t)} \right\|_F^2 \\
&= \frac{2(1-\rho)}{1+(1-\rho)} \left\| \mathbf{X}^{(t)} - \mathbf{Q}_x^{(t)} \right\|_F^2 \\
&\quad + \frac{2}{\rho} \mathbb{E}_t \left\| \gamma(\mathbf{Q}_x^{(t+1)} - \mathbf{X}^{(t)})(\mathbf{W} - \mathbf{I}_n) + \gamma(\mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top)(\mathbf{W} - \mathbf{I}_n) - \eta \mathbf{V}^{(t+1)} \right\|_F^2 \\
&\stackrel{(iii)}{\leq} \left(1 - \frac{\rho}{2}\right) \left\| \mathbf{X}^{(t)} - \mathbf{Q}_x^{(t)} \right\|_F^2 + \frac{24\gamma^2}{\rho} \mathbb{E}_t \left\| \mathbf{Q}_x^{(t+1)} - \mathbf{X}^{(t)} \right\|_F^2 + \frac{24\gamma^2}{\rho} \left\| \mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top \right\|_F^2 + \frac{6\eta^2}{\rho} \mathbb{E}_t \left\| \mathbf{V}^{(t+1)} \right\|_F^2 \\
&\leq \left(1 - \frac{\rho}{2} + \frac{24(1-\rho)\gamma^2}{\rho}\right) \left\| \mathbf{Q}_x^{(t)} - \mathbf{X}^{(t)} \right\|_F^2 + \frac{24\gamma^2}{\rho} \left\| \mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top \right\|_F^2 + \frac{6\eta^2}{\rho} \mathbb{E}_t \left\| \mathbf{V}^{(t+1)} \right\|_F^2 \\
&\stackrel{(iv)}{\leq} \left(1 - \frac{\rho}{4}\right) \left\| \mathbf{Q}_x^{(t)} - \mathbf{X}^{(t)} \right\|_F^2 + \frac{24\gamma^2}{\rho} \left\| \mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top \right\|_F^2 + \frac{6\eta^2}{\rho} \mathbb{E}_t \left\| \mathbf{V}^{(t+1)} \right\|_F^2, \tag{13}
\end{aligned}$$

where (i) is obtained by applying Young's inequality, (ii) uses Definition 3, (iii) uses the fact $\|\mathbf{W} - \mathbf{I}_n\|_{\text{op}} \leq 2$, and (iv) uses (12).

B.2.3 Linear system

Let $\mathbf{e}_1^{(t)} = \begin{bmatrix} \left\| \mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top \right\|_F^2 \\ \left\| \mathbf{Q}_x^{(t)} - \mathbf{X}^{(t)} \right\|_F^2 \end{bmatrix}$, we can take full expectation and rewrite (11) and (13) in matrix form as

$$\begin{aligned}
\mathbb{E}[\mathbf{e}_1^{(t+1)}] &\leq \begin{bmatrix} \hat{\alpha} & \frac{16(1-\rho)\gamma^2}{1-\hat{\alpha}} \\ \frac{24\gamma^2}{\rho} & 1 - \frac{\rho}{4} \end{bmatrix} \mathbb{E}[\mathbf{e}_1^{(t)}] + \begin{bmatrix} \frac{4\eta^2}{1-\hat{\alpha}} \mathbb{E} \left\| \mathbf{V}^{(t+1)} - \bar{\mathbf{v}}^{(t+1)} \mathbf{1}_n^\top \right\|_F^2 \\ \frac{6\eta^2}{\rho} \mathbb{E} \left\| \mathbf{V}^{(t+1)} \right\|_F^2 \end{bmatrix} \\
&:= \mathbf{G}_1 \mathbb{E}[\mathbf{e}_1^{(t)}] + \mathbf{b}_1^{(t)}. \tag{14}
\end{aligned}$$

We can compute $(\mathbf{I}_n - \mathbf{G}_1)^{-1}$ and verify all its entries are positive:

$$\begin{aligned}
(\mathbf{I}_n - \mathbf{G}_1)^{-1} &= \frac{1}{(1-\hat{\alpha}) \cdot \frac{\rho}{4} - \frac{16(1-\rho)\gamma^2}{1-\hat{\alpha}} \cdot \frac{24\gamma^2}{\rho}} \begin{bmatrix} \frac{\rho}{4} & \frac{16(1-\rho)\gamma^2}{1-\hat{\alpha}} \\ \frac{24\gamma^2}{\rho} & 1 - \hat{\alpha} \end{bmatrix} \\
&\leq \frac{1}{\frac{1}{8}(1-\hat{\alpha})\rho} \begin{bmatrix} \frac{\rho}{4} & \frac{16(1-\rho)\gamma^2}{1-\hat{\alpha}} \\ \frac{24\gamma^2}{\rho} & 1 - \hat{\alpha} \end{bmatrix}, \tag{15}
\end{aligned}$$

where we assume the following inequality to prove (15), which will be validated in Appendix B.4:

$$(1-\hat{\alpha}) \cdot \frac{\rho}{4} - \frac{16(1-\rho)\gamma^2}{1-\hat{\alpha}} \cdot \frac{24\gamma^2}{\rho} \geq \frac{1}{8}(1-\hat{\alpha})\rho. \tag{16}$$

Sum expected error vectors $\mathbb{E}[\mathbf{e}_1^{(t)}]$ over $t = 1, \dots, T$,

$$\begin{aligned}
\sum_{t=1}^T \mathbb{E}[\mathbf{e}_1^{(t)}] &\leq \sum_{t=1}^T (\mathbf{G}_1 \mathbb{E}[\mathbf{e}_1^{(t-1)}] + \mathbf{b}_1^{(t-1)}) \\
&\leq \mathbf{G}_1 \sum_{t=1}^T \mathbb{E}[\mathbf{e}_1^{(t)}] + \mathbf{G}_1 \mathbb{E}[\mathbf{e}_1^{(0)}] + \sum_{t=1}^T \mathbf{b}_1^{(t-1)}.
\end{aligned}$$

Reorganize terms, multiply $(\mathbf{I}_n - \mathbf{G}_1)^{-1}$ on both sides and use $\mathbf{e}_1^{(0)} = \mathbf{0}_2$, the sum of error vectors can be bounded as

$$\sum_{t=1}^T \mathbb{E}[\mathbf{e}_1^{(t)}] \leq (\mathbf{I}_n - \mathbf{G}_1)^{-1} \sum_{t=0}^{T-1} \mathbf{b}_1^{(t)}. \quad (17)$$

The sum of consensus error can be computed as

$$\begin{aligned} & \sum_{t=1}^T \mathbb{E} \|\mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top\|_{\mathbb{F}}^2 \\ & \leq \begin{bmatrix} 1 & 0 \end{bmatrix} (\mathbf{I}_n - \mathbf{G})^{-1} \sum_{t=0}^{T-1} \mathbf{b}_1^{(t)} \\ & = \frac{1}{\frac{1}{8}(1-\hat{\alpha})\rho} \begin{bmatrix} \frac{\rho}{4} & \frac{16(1-\rho)\gamma^2}{1-\hat{\alpha}} \end{bmatrix} \begin{bmatrix} \frac{4\eta^2}{1-\hat{\alpha}} \sum_{t=1}^T \mathbb{E} \|\mathbf{V}^{(t)} - \bar{\mathbf{v}}^{(t)} \mathbf{1}_n^\top\|_{\mathbb{F}}^2 \\ \frac{6\eta^2}{\rho} \sum_{t=1}^T \mathbb{E} \|\mathbf{V}^{(t)}\|_{\mathbb{F}}^2 \end{bmatrix} \\ & = \frac{\eta^2}{\frac{1}{8}(1-\hat{\alpha})\rho} \left(\frac{\rho}{1-\hat{\alpha}} \sum_{t=1}^T \mathbb{E} \|\mathbf{V}^{(t)} - \bar{\mathbf{v}}^{(t)} \mathbf{1}_n^\top\|_{\mathbb{F}}^2 + \frac{16(1-\rho)\gamma^2}{1-\hat{\alpha}} \cdot \frac{6}{\rho} \sum_{t=1}^T \mathbb{E} \|\mathbf{V}^{(t)}\|_{\mathbb{F}}^2 \right) \\ & \stackrel{(i)}{=} \frac{8\eta^2}{(1-\hat{\alpha})\rho} \left(\frac{\rho}{1-\hat{\alpha}} + \frac{96(1-\rho)\gamma^2}{(1-\hat{\alpha})\rho} \right) \sum_{t=1}^T \mathbb{E} \|\mathbf{V}^{(t)} - \bar{\mathbf{v}}^{(t)} \mathbf{1}_n^\top\|_{\mathbb{F}}^2 + \frac{768(1-\rho)\gamma^2\eta^2}{(1-\hat{\alpha})^2\rho^2} \sum_{t=1}^T n \mathbb{E} \|\bar{\mathbf{v}}^{(t)}\|_{\mathbb{F}}^2 \\ & \stackrel{(ii)}{\leq} \frac{16\eta^2}{(1-\hat{\alpha})^2} \sum_{t=1}^T \mathbb{E} \|\mathbf{V}^{(t)} - \bar{\mathbf{v}}^{(t)} \mathbf{1}_n^\top\|_{\mathbb{F}}^2 + \frac{768(1-\rho)\gamma^2\eta^2}{(1-\hat{\alpha})^2\rho^2} \sum_{t=1}^T n \mathbb{E} \|\bar{\mathbf{v}}^{(t)}\|_{\mathbb{F}}^2, \end{aligned} \quad (18)$$

where we use the equality $\|\mathbf{V}^{(t)}\|_{\mathbb{F}}^2 = \|\mathbf{V}^{(t)} - \bar{\mathbf{v}}^{(t)} \mathbf{1}_n^\top\|_{\mathbb{F}}^2 + n \|\bar{\mathbf{v}}^{(t)}\|_2^2$ for (i) and use (12) for (ii).

B.3 Sum of gradient consensus errors

This section creates a linear system to bound the sum of gradient consensus error $\sum_{t=1}^T \mathbb{E} \|\mathbf{V}^{(t)} - \bar{\mathbf{v}}^{(t)} \mathbf{1}_n^\top\|_{\mathbb{F}}^2$ by $\sum_{t=1}^T \mathbb{E} \|\bar{\mathbf{v}}^{(t)}\|_2^2$ and constant terms.

B.3.1 Gradient consensus error

Take expectation conditioned on time t and reorganize terms, the gradient consensus error can be expanded as

$$\begin{aligned} & \mathbb{E}_t \|\mathbf{V}^{(t+1)} - \bar{\mathbf{v}}^{(t+1)} \mathbf{1}_n^\top\|_{\mathbb{F}}^2 \\ & = \mathbb{E}_t \left\| \left(\mathbf{V}^{(t)} + \gamma \mathbf{Q}_v^{(t+1)} (\mathbf{W} - \mathbf{I}_n) + \mathbf{G}_p^{(t+1)} - \mathbf{G}_p^{(t)} \right) (\mathbf{I}_n - (\frac{1}{n} \mathbf{1}_n \mathbf{1}_n^\top)) \right\|_{\mathbb{F}}^2 \\ & = \mathbb{E}_t \left\| \left(\mathbf{V}^{(t)} - \bar{\mathbf{v}}^{(t)} \mathbf{1}_n^\top \right) (\widehat{\mathbf{W}} - (\frac{1}{n} \mathbf{1}_n \mathbf{1}_n^\top)) + \gamma (\mathbf{Q}_v^{(t+1)} - \mathbf{V}^{(t)}) (\mathbf{W} - \mathbf{I}_n) + (\mathbf{G}_p^{(t+1)} - \mathbf{G}_p^{(t)}) (\mathbf{I}_n - (\frac{1}{n} \mathbf{1}_n \mathbf{1}_n^\top)) \right\|_{\mathbb{F}}^2. \end{aligned}$$

Then, take full expectation, use the update formula and Young's inequality similarly to (11),

$$\begin{aligned} & \mathbb{E} \|\mathbf{V}^{(t+1)} - \bar{\mathbf{v}}^{(t+1)} \mathbf{1}_n^\top\|_{\mathbb{F}}^2 \\ & \leq \frac{2\hat{\alpha}^2}{1+\hat{\alpha}^2} \|\mathbf{V}^{(t)} - \bar{\mathbf{v}}^{(t)} \mathbf{1}_n^\top\|_{\mathbb{F}}^2 \\ & \quad + \frac{2}{1-\hat{\alpha}^2} \mathbb{E} \left\| \gamma (\mathbf{Q}_v^{(t+1)} - \mathbf{V}^{(t)}) (\mathbf{W} - \mathbf{I}_n) + (\mathbf{G}_p^{(t+1)} - \mathbf{G}_p^{(t)}) (\mathbf{I}_n - (\frac{1}{n} \mathbf{1}_n \mathbf{1}_n^\top)) \right\|_{\mathbb{F}}^2 \\ & \leq \hat{\alpha} \|\mathbf{V}^{(t)} - \bar{\mathbf{v}}^{(t)} \mathbf{1}_n^\top\|_{\mathbb{F}}^2 + \frac{4}{1-\hat{\alpha}} \mathbb{E} \|\gamma (\mathbf{Q}_v^{(t+1)} - \mathbf{V}^{(t)}) (\mathbf{W} - \mathbf{I}_n)\|_{\mathbb{F}}^2 \\ & \quad + \frac{4}{1-\hat{\alpha}} \mathbb{E} \left\| (\mathbf{G}_p^{(t+1)} - \mathbf{G}_p^{(t)}) (\mathbf{I}_n - (\frac{1}{n} \mathbf{1}_n \mathbf{1}_n^\top)) \right\|_{\mathbb{F}}^2 \end{aligned}$$

$$\begin{aligned}
&\stackrel{(i)}{\leq} \hat{\alpha} \|\mathbf{V}^{(t)} - \bar{\mathbf{v}}^{(t)} \mathbf{1}_n^\top\|_F^2 + \frac{16\gamma^2}{1-\hat{\alpha}} \mathbb{E} \|\mathbf{Q}_v^{(t+1)} - \mathbf{V}^{(t)}\|_F^2 + \frac{4}{1-\hat{\alpha}} \mathbb{E} \|\mathbf{G}_p^{(t+1)} - \mathbf{G}_p^{(t)}\|_F^2 \\
&\stackrel{(ii)}{\leq} \hat{\alpha} \|\mathbf{V}^{(t)} - \bar{\mathbf{v}}^{(t)} \mathbf{1}_n^\top\|_F^2 + \frac{16(1-\rho)\gamma^2}{1-\hat{\alpha}} \mathbb{E} \|\mathbf{Q}_v^{(t)} - \mathbf{V}^{(t)}\|_F^2 + \frac{16n(\tau^2 + \sigma_p^2 d)}{1-\hat{\alpha}}, \tag{19}
\end{aligned}$$

where (i) is proved using the facts $\|\mathbf{W} - \mathbf{I}_n\|_{\text{op}} \leq 2$ and $\|\mathbf{I}_n - (\frac{1}{n} \mathbf{1}_n \mathbf{1}_n^\top)\|_{\text{op}} \leq 1$, (ii) is due to Definition 3 and

$$\begin{aligned}
\mathbb{E} \|\mathbf{G}_p^{(t+1)} - \mathbf{G}_p^{(t)}\|_F^2 &\leq 2\mathbb{E} \|\mathbf{G}_p^{(t+1)}\|_F^2 + 2\mathbb{E} \|\mathbf{G}_p^{(t)}\|_F^2 \\
&= 2(\mathbb{E} \|\mathbf{G}_\tau^{(t+1)}\|_F^2 + n\sigma_p^2 d) + 2(\mathbb{E} \|\mathbf{G}_\tau^{(t)}\|_F^2 + n\sigma_p^2 d) \\
&\leq 4n(\tau^2 + \sigma_p^2 d). \tag{20}
\end{aligned}$$

B.3.2 Gradient quantization error

$$\begin{aligned}
\mathbb{E}_t \|\mathbf{Q}_v^{(t+1)} - \mathbf{V}^{(t+1)}\|_F^2 &= \mathbb{E}_t \|(\mathbf{Q}_v^{(t+1)} - \mathbf{V}^{(t)}) - (\mathbf{V}^{(t+1)} - \mathbf{V}^{(t)})\|_F^2 \\
&\leq \frac{2}{1+(1-\rho)} \mathbb{E}_t \|\mathbf{Q}_v^{(t+1)} - \mathbf{V}^{(t)}\|_F^2 + \frac{2}{1-(1-\rho)} \mathbb{E}_t \|\mathbf{V}^{(t+1)} - \mathbf{V}^{(t)}\|_F^2 \\
&\leq \frac{2(1-\rho)}{2-\rho} \|\mathbf{Q}_v^{(t)} - \mathbf{V}^{(t)}\|_F^2 + \frac{2}{\rho} \mathbb{E}_t \|\gamma \mathbf{Q}_v^{(t+1)} (\mathbf{W} - \mathbf{I}_n) + \mathbf{G}_p^{(t+1)} - \mathbf{G}_p^{(t)}\|_F^2 \\
&\leq \frac{2(1-\rho)}{2-\rho} \|\mathbf{Q}_v^{(t)} - \mathbf{V}^{(t)}\|_F^2 + \frac{6\gamma^2}{\rho} \mathbb{E}_t \|(\mathbf{Q}_v^{(t+1)} - \mathbf{V}^{(t)}) (\mathbf{W} - \mathbf{I}_n)\|_F^2 \\
&\quad + \frac{6\gamma^2}{\rho} \|\mathbf{V}^{(t)} (\mathbf{W} - \mathbf{I}_n)\|_F^2 + \frac{6}{\rho} \mathbb{E}_t \|\mathbf{G}_p^{(t+1)} - \mathbf{G}_p^{(t)}\|_F^2 \\
&\stackrel{(i)}{\leq} \left(1 - \frac{\rho}{2} + \frac{24\gamma^2(1-\rho)}{\rho}\right) \|\mathbf{Q}_v^{(t)} - \mathbf{V}^{(t)}\|_F^2 + \frac{24\gamma^2}{\rho} \|\mathbf{V}^{(t)} - \bar{\mathbf{v}}^{(t)} \mathbf{1}_n^\top\|_F^2 + \frac{24n(\tau^2 + \sigma_p^2 d)}{\rho} \\
&\stackrel{(ii)}{\leq} \left(1 - \frac{\rho}{4}\right) \|\mathbf{Q}_v^{(t)} - \mathbf{V}^{(t)}\|_F^2 + \frac{24\gamma^2}{\rho} \|\mathbf{V}^{(t)} - \bar{\mathbf{v}}^{(t)} \mathbf{1}_n^\top\|_F^2 + \frac{24n(\tau^2 + \sigma_p^2 d)}{\rho}, \tag{21}
\end{aligned}$$

where we use (20) and the fact $\frac{2(1-\rho)}{2-\rho} = 1 - \frac{\rho}{2-\rho} \geq 1 - \frac{\rho}{2}$ when $\rho \geq 0$ to reach (i) and use (12) to reach (ii).

B.3.3 Linear system

Let $\mathbf{e}_2^{(t)} = \begin{bmatrix} \|\mathbf{V}^{(t)} - \bar{\mathbf{v}}^{(t)} \mathbf{1}_n^\top\|_F^2 \\ \|\mathbf{Q}_v^{(t)} - \mathbf{V}^{(t)}\|_F^2 \end{bmatrix}$. We can write (19) and (21) in matrix form as

$$\begin{aligned}
\mathbb{E}[\mathbf{e}_2^{(t+1)}] &\leq \begin{bmatrix} \hat{\alpha} & \frac{16(1-\rho)\gamma^2}{1-\hat{\alpha}} \\ \frac{24\gamma^2}{\rho} & 1 - \frac{\rho}{4} \end{bmatrix} \mathbb{E}[\mathbf{e}_2^{(t)}] + \begin{bmatrix} \frac{16n(\tau^2 + \sigma_p^2 d)}{1-\hat{\alpha}} \\ \frac{24n(\tau^2 + \sigma_p^2 d)}{\rho} \end{bmatrix} \\
&:= \mathbf{G}_2 \mathbb{E}[\mathbf{e}_2^{(t)}] + \mathbf{b}_2^{(t)}.
\end{aligned}$$

Because $\mathbf{G}_2 = \mathbf{G}_1$, we can use the same argument as in Appendix B.2.3, and use (15) to prove

$$\begin{aligned}
\sum_{t=1}^T \mathbb{E} \|\mathbf{V}^{(t)} - \bar{\mathbf{v}}^{(t)} \mathbf{1}_n^\top\|_F^2 &\leq [1 \quad 0] (\mathbf{I}_n - \mathbf{G}_2)^{-1} \left(\mathbb{E}[\mathbf{e}_2^{(0)}] + \sum_{t=0}^{T-1} \mathbf{b}_2^{(t)} \right) \\
&\leq \frac{1}{\frac{1}{8}(1-\hat{\alpha})\rho} \begin{bmatrix} \frac{\rho}{4} & \frac{16(1-\rho)\gamma^2}{1-\hat{\alpha}} \\ \frac{16Tn(\tau^2 + \sigma_p^2 d)}{1-\hat{\alpha}} & \frac{24Tn(\tau^2 + \sigma_p^2 d)}{\rho} \end{bmatrix} \\
&= \frac{Tn(\tau^2 + \sigma_p^2 d)}{\frac{1}{8}(1-\hat{\alpha})\rho} \cdot \left(\frac{\rho}{4} \cdot \frac{16}{1-\hat{\alpha}} + \frac{16(1-\rho)\gamma^2}{1-\hat{\alpha}} \cdot \frac{24}{\rho} \right)
\end{aligned}$$

$$\begin{aligned}
&\leq \frac{Tn(\tau^2 + \sigma_p^2 d)}{\frac{1}{8}(1 - \hat{\alpha})\rho} \cdot \left(\frac{4\rho}{1 - \hat{\alpha}} + \frac{4\rho}{1 - \hat{\alpha}} \right) \\
&= \frac{64}{(1 - \hat{\alpha})^2} Tn(\tau^2 + \sigma_p^2 d),
\end{aligned} \tag{22}$$

where we use (12) to prove the last inequality.

With (22), we can bound (18) by

$$\begin{aligned}
&\sum_{t=1}^T \mathbb{E} \|\mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top\|_F^2 \\
&\leq \frac{16\eta^2}{(1 - \hat{\alpha})^2} \sum_{t=1}^T \mathbb{E} \|\mathbf{V}^{(t)} - \bar{\mathbf{v}}^{(t)} \mathbf{1}_n^\top\|_F^2 + \frac{768(1 - \rho)\gamma^2\eta^2}{(1 - \hat{\alpha})^2\rho^2} \sum_{t=1}^T n\mathbb{E} \|\bar{\mathbf{v}}^{(t)}\|_F^2 \\
&\leq \frac{16\eta^2}{(1 - \hat{\alpha})^2} \cdot \frac{64}{(1 - \hat{\alpha})^2} Tn(\tau^2 + \sigma_p^2 d) + \frac{768(1 - \rho)\gamma^2\eta^2}{(1 - \hat{\alpha})^2\rho^2} \sum_{t=1}^T n\mathbb{E} \|\bar{\mathbf{v}}^{(t)}\|_F^2 \\
&\leq \frac{1024\eta^2}{(1 - \hat{\alpha})^4} Tn(\tau^2 + \sigma_p^2 d) + \frac{8\eta^2}{(1 - \hat{\alpha})^2} \sum_{t=1}^T n\mathbb{E} \|\bar{\mathbf{v}}^{(t)}\|_F^2,
\end{aligned} \tag{23}$$

where we use (12) again to prove the last inequality.

B.4 Convergence rate

Note bounded gradient assumption can imply Assumption 3 for some $\sigma_g \leq 2\tau$, we can bound the expected norm of average gradient estimate as

$$\begin{aligned}
\mathbb{E} \|\bar{\mathbf{v}}^{(t)}\|_2^2 &= \mathbb{E} \|\bar{\mathbf{g}}_p^{(t)}\|_2^2 \\
&= \mathbb{E} \|\bar{\mathbf{g}}_\tau^{(t)}\|_2^2 + \frac{\sigma_p^2 d}{n} \\
&\leq \mathbb{E} \|\nabla F(\mathbf{X}^{(t)}) (\frac{1}{n} \mathbf{1}_n)\|_2^2 + \frac{\sigma_g^2}{b} + \frac{\sigma_p^2 d}{n} \\
&\leq \mathbb{E} \|\nabla F(\mathbf{X}^{(t)}) (\frac{1}{n} \mathbf{1}_n)\|_2^2 + \frac{4\tau^2}{b} + \frac{\sigma_p^2 d}{n}.
\end{aligned} \tag{24}$$

We assume

$$\eta L \leq \frac{1}{8}(1 - \hat{\alpha})^{\frac{4}{3}}. \tag{25}$$

Using (23) (24), expected utility (9) can be bounded by

$$\begin{aligned}
\frac{1}{T} \sum_{t=1}^T \mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2^2 &\leq \frac{2\Delta}{\eta T} + \frac{1}{T} \cdot \frac{L^2}{n} \sum_{t=1}^T \mathbb{E} \|\mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top\|_F^2 \\
&\quad + \frac{1}{T} \sum_{t=1}^T \eta L \mathbb{E} \|\bar{\mathbf{v}}^{(t)}\|_2^2 - \frac{1}{T} \sum_{t=1}^T \mathbb{E} \|\nabla F(\mathbf{X}^{(t)}) (\frac{1}{n} \mathbf{1}_n)\|_2^2 \\
&\leq \frac{2\Delta}{\eta T} + \frac{1}{T} \cdot \frac{L^2}{n} \left(\frac{1024\eta^2}{(1 - \hat{\alpha})^4} Tn(\tau^2 + \sigma_p^2 d) + \frac{8\eta^2}{(1 - \hat{\alpha})^2} \sum_{t=1}^T n\mathbb{E} \|\bar{\mathbf{v}}^{(t)}\|_2^2 \right) \\
&\quad + \frac{1}{T} \sum_{t=1}^T \eta L \mathbb{E} \|\bar{\mathbf{v}}^{(t)}\|_2^2 - \frac{1}{T} \sum_{t=1}^T \mathbb{E} \|\nabla F(\mathbf{X}^{(t)}) (\frac{1}{n} \mathbf{1}_n)\|_2^2 \\
&\stackrel{(i)}{\leq} \frac{2\Delta}{\eta T} + \frac{1024\eta^2 L^2}{(1 - \hat{\alpha})^4} (\tau^2 + \sigma_p^2 d) + \frac{2\eta L}{(1 - \hat{\alpha})^{\frac{4}{3}} T} \sum_{t=1}^T \mathbb{E} \|\bar{\mathbf{v}}^{(t)}\|_2^2 - \frac{1}{T} \sum_{t=1}^T \mathbb{E} \|\nabla F(\mathbf{X}^{(t)}) (\frac{1}{n} \mathbf{1}_n)\|_2^2
\end{aligned}$$

$$\begin{aligned}
&\leq \frac{2\Delta}{\eta T} + \frac{1024\eta^2 L^2}{(1-\hat{\alpha})^4}(\tau^2 + \sigma_p^2 d) + \frac{2\eta L}{(1-\hat{\alpha})^{\frac{4}{3}}} \left(\frac{4\tau^2}{b} + \frac{\sigma_p^2 d}{n} \right) \\
&\stackrel{\text{(ii)}}{=} \frac{2\Delta}{\eta T} + \frac{1024\eta^2 L^2 \tau^2}{(1-\hat{\alpha})^4} (1 + T\phi_m^2) + \frac{8\eta L \tau^2}{(1-\hat{\alpha})^{\frac{4}{3}}} (1 + T\phi_m^2) \\
&\stackrel{\text{(iii)}}{=} \frac{2\Delta}{\eta T} + \frac{2048\eta^2 L^2 \tau^2}{(1-\hat{\alpha})^4} + \frac{16\eta L \tau^2}{(1-\hat{\alpha})^{\frac{4}{3}}}
\end{aligned} \tag{26}$$

where we use (25) for (i), substitute $b = 1$ and $\sigma_p^2 d = \left(\frac{\tau \sqrt{T \log(1/\delta)}}{m\epsilon} \right)^2 = T\tau^2 \phi_m^2$ for (ii), and substitute $T = \phi_m^{-2}$ for (iii).

We set the step size as

$$\eta = \frac{\gamma^{\frac{4}{3}}(1-\alpha)^{\frac{4}{3}}}{32} \cdot \frac{\phi_m}{L},$$

(26) can be further bounded as

$$\begin{aligned}
\frac{1}{T} \sum_{t=1}^T \mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2^2 &\leq \frac{64L\Delta\phi_m}{\gamma^{\frac{4}{3}}(1-\alpha)^{\frac{4}{3}}} + \frac{2\tau^2\phi_m^2}{(1-\hat{\alpha})^{\frac{4}{3}}} + \frac{\tau^2\phi_m}{2} \\
&\leq \frac{64L\Delta\phi_m}{\gamma^{\frac{4}{3}}(1-\alpha)^{\frac{4}{3}}} + \frac{3\tau^2\phi_m}{(1-\hat{\alpha})^{\frac{4}{3}}} \\
&\leq \frac{67\phi_m}{\gamma^{\frac{4}{3}}(1-\alpha)^{\frac{4}{3}}} \max\{\tau^2, L\Delta\},
\end{aligned}$$

where we use Lemma 1 to reach the last inequality.

Lastly, set the hyper parameter γ as

$$\gamma = \frac{1}{100}(1-\alpha)\rho.$$

We can now verify conditions (12), (16) and the condition on η are all met to conclude the proof:

$$\gamma^2 \leq \frac{\rho^2}{10000} \quad \Rightarrow \tag{12}$$

$$\gamma^4 = \gamma^2 \cdot \frac{(1-\alpha)^2 \rho^2}{10000} \leq \frac{(1-\hat{\alpha})^2 \rho^2}{10000} \quad \Rightarrow \tag{16}$$

$$\eta L \leq \frac{(1-\hat{\alpha})^{\frac{4}{3}}}{32} \quad \Rightarrow \tag{25}$$

C Proof of Theorem 3

This section proves Theorem 3 in 2 subsections. Appendix C.1 derives the descent inequality using results from Appendices B.2 and B.3. Appendix C.2 first assumes all expected gradient norm $\mathbb{E}\|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2$ are greater than a threshold ν (i.e. $\mathbb{E}\|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 \geq \nu$ for all $t = 1, \dots, T$), then specifies parameters and proves the average of expected gradient norm is smaller than that threshold $\frac{1}{T} \sum_{t=1}^T \mathbb{E}\|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 \leq \nu$, which contradicts the assumption hence proves the algorithm reaches $\mathbb{E}\|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 \leq \nu$ within T steps.

C.1 Function value descent

Let $\delta_i^{(t)} = \frac{\tau}{\tau + \|\mathbf{g}_i^{(t)}\|_2}$ and $\delta^{(t)} = \frac{\tau}{\tau + \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2}$. Similar to Appendix B.1, use Taylor expansion and take expectation conditioned on t , we can expand the function value descent as

$$\mathbb{E}_t [f(\bar{\mathbf{x}}^{(t+1)}) - f(\bar{\mathbf{x}}^{(t)})]$$

$$\begin{aligned}
&\leq \mathbb{E}_t \langle \nabla f(\bar{\mathbf{x}}^{(t)}), -\eta \bar{\mathbf{v}}^{(t+1)} \rangle + \frac{L}{2} \mathbb{E}_t \|\eta \bar{\mathbf{v}}^{(t+1)}\|_2^2 \\
&= -\eta \mathbb{E}_t \langle \nabla f(\bar{\mathbf{x}}^{(t)}), \bar{\mathbf{g}}_p^{(t+1)} \rangle + \frac{\eta^2 L}{2} \mathbb{E}_t \|\bar{\mathbf{v}}^{(t+1)}\|_2^2 \\
&= -\eta \mathbb{E}_t \langle \nabla f(\bar{\mathbf{x}}^{(t)}), \bar{\mathbf{g}}_\tau^{(t+1)} \rangle + \frac{\eta^2 L}{2} \mathbb{E}_t \|\bar{\mathbf{v}}^{(t+1)}\|_2^2 \\
&= -\eta \mathbb{E}_t \langle \nabla f(\bar{\mathbf{x}}^{(t)}), \text{Clip}_\tau(\nabla f(\bar{\mathbf{x}}^{(t)})) \rangle + \eta \mathbb{E}_t \langle \nabla f(\bar{\mathbf{x}}^{(t)}), \text{Clip}_\tau(\nabla f(\bar{\mathbf{x}}^{(t)})) - \bar{\mathbf{g}}_\tau^{(t+1)} \rangle + \frac{\eta^2 L}{2} \mathbb{E}_t \|\bar{\mathbf{v}}^{(t+1)}\|_2^2 \\
&= -\eta \delta^{(t)} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2^2 + \eta \mathbb{E}_t \langle \nabla f(\bar{\mathbf{x}}^{(t)}), \text{Clip}_\tau(\nabla f(\bar{\mathbf{x}}^{(t)})) - \bar{\mathbf{g}}_\tau^{(t+1)} \rangle + \frac{\eta^2 L}{2} \mathbb{E}_t \|\bar{\mathbf{v}}^{(t+1)}\|_2^2 \\
&\leq -\eta \delta^{(t)} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2^2 + \eta \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 \mathbb{E}_t \|\text{Clip}_\tau(\nabla f(\bar{\mathbf{x}}^{(t)})) - \bar{\mathbf{g}}_\tau^{(t+1)}\|_2 + \frac{\eta^2 L}{2} \mathbb{E}_t \|\bar{\mathbf{v}}^{(t+1)}\|_2^2. \tag{27}
\end{aligned}$$

The $\mathbb{E}_t \|\text{Clip}_\tau(\nabla f(\bar{\mathbf{x}}^{(t)})) - \bar{\mathbf{g}}_\tau^{(t+1)}\|_2$ term in (27) is the error introduced by gradient clipping, which can be analyzed by splitting it to 4 terms as following

$$\begin{aligned}
&\mathbb{E}_t \|\text{Clip}_\tau(\nabla f(\bar{\mathbf{x}}^{(t)})) - \bar{\mathbf{g}}_\tau^{(t+1)}\|_2 \\
&= \mathbb{E}_t \left\| \frac{1}{n} \sum_{i=1}^n \frac{\tau}{\tau + \|\mathbf{g}_i^{(t)}\|_2} \mathbf{g}_i^{(t)} - \frac{\tau}{\tau + \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2} \nabla f(\bar{\mathbf{x}}^{(t)}) \right\|_2 \\
&= \mathbb{E}_t \left\| \frac{1}{n} \sum_{i=1}^n \left(\frac{\tau}{\tau + \|\mathbf{g}_i^{(t)}\|_2} \mathbf{g}_i^{(t)} - \frac{\tau}{\tau + \|\nabla f_i(\mathbf{x}_i^{(t)})\|_2} \mathbf{g}_i^{(t)} \right) \right. \\
&\quad + \frac{1}{n} \sum_{i=1}^n \left(\frac{\tau}{\tau + \|\nabla f_i(\mathbf{x}_i^{(t)})\|_2} \mathbf{g}_i^{(t)} - \frac{\tau}{\tau + \|\nabla f_i(\mathbf{x}_i^{(t)})\|_2} \nabla f_i(\mathbf{x}_i^{(t)}) \right) \\
&\quad + \frac{1}{n} \sum_{i=1}^n \left(\frac{\tau}{\tau + \|\nabla f_i(\mathbf{x}_i^{(t)})\|_2} \nabla f_i(\mathbf{x}_i^{(t)}) - \frac{\tau}{\tau + \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2} \nabla f_i(\mathbf{x}_i^{(t)}) \right) \\
&\quad \left. + \frac{1}{n} \sum_{i=1}^n \left(\frac{\tau}{\tau + \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2} \nabla f_i(\mathbf{x}_i^{(t)}) - \frac{\tau}{\tau + \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2} \nabla f(\bar{\mathbf{x}}^{(t)}) \right) \right\|_2 \\
&\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}_t \left\| \left(\frac{\tau}{\tau + \|\mathbf{g}_i^{(t)}\|_2} - \frac{\tau}{\tau + \|\nabla f_i(\mathbf{x}_i^{(t)})\|_2} \right) \mathbf{g}_i^{(t)} \right\|_2 \tag{28}
\end{aligned}$$

$$+ \frac{1}{n} \sum_{i=1}^n \mathbb{E}_t \left\| \frac{\tau}{\tau + \|\nabla f_i(\mathbf{x}_i^{(t)})\|_2} (\mathbf{g}_i^{(t)} - \nabla f_i(\mathbf{x}_i^{(t)})) \right\|_2 \tag{29}$$

$$+ \frac{1}{n} \sum_{i=1}^n \mathbb{E}_t \left\| \left(\frac{\tau}{\tau + \|\nabla f_i(\mathbf{x}_i^{(t)})\|_2} - \frac{\tau}{\tau + \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2} \right) \nabla f_i(\mathbf{x}_i^{(t)}) \right\|_2 \tag{30}$$

$$+ \left\| \frac{\tau}{\tau + \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2} \left(\frac{1}{n} \sum_{i=1}^n \nabla f_i(\mathbf{x}_i^{(t)}) - \nabla f(\bar{\mathbf{x}}^{(t)}) \right) \right\|_2. \tag{31}$$

Next, we bound each term separately using triangle inequality, Assumptions 3 and 4. Bound the first term (28) as

$$\begin{aligned}
&\frac{1}{n} \sum_{i=1}^n \mathbb{E}_t \left\| \left(\frac{\tau}{\tau + \|\mathbf{g}_i^{(t)}\|_2} - \frac{\tau}{\tau + \|\nabla f_i(\mathbf{x}_i^{(t)})\|_2} \right) \mathbf{g}_i^{(t)} \right\|_2 \\
&= \frac{1}{n} \sum_{i=1}^n \mathbb{E}_t \left\| \frac{\tau (\|\mathbf{g}_i^{(t)}\|_2 - \|\nabla f_i(\mathbf{x}_i^{(t)})\|_2)}{(\tau + \|\mathbf{g}_i^{(t)}\|_2)(\tau + \|\nabla f_i(\mathbf{x}_i^{(t)})\|_2)} \mathbf{g}_i^{(t)} \right\|_2 \\
&= \frac{1}{n} \sum_{i=1}^n \mathbb{E}_t \left(\left| \|\mathbf{g}_i^{(t)}\|_2 - \|\nabla f_i(\mathbf{x}_i^{(t)})\|_2 \right| \cdot \frac{\tau}{\tau + \|\nabla f_i(\mathbf{x}_i^{(t)})\|_2} \cdot \frac{\|\mathbf{g}_i^{(t)}\|_2}{\tau + \|\mathbf{g}_i^{(t)}\|_2} \right)
\end{aligned}$$

$$\begin{aligned}
&\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}_t \left| \|\mathbf{g}_i^{(t)}\|_2 - \|\nabla f_i(\mathbf{x}_i^{(t)})\|_2 \right| \\
&\leq \frac{1}{n} \sum_{i=1}^n \sqrt{\mathbb{E}_t (\|\mathbf{g}_i^{(t)}\|_2 - \|\nabla f_i(\mathbf{x}_i^{(t)})\|_2)^2} \\
&= \frac{1}{n} \sum_{i=1}^n \sqrt{\mathbb{E}_t (\|\mathbf{g}_i^{(t)}\|_2^2 + \|\nabla f_i(\mathbf{x}_i^{(t)})\|_2^2 - 2\|\mathbf{g}_i^{(t)}\|_2 \|\nabla f_i(\mathbf{x}_i^{(t)})\|_2)} \\
&\leq \frac{1}{n} \sum_{i=1}^n \sqrt{\mathbb{E}_t (\|\mathbf{g}_i^{(t)}\|_2^2 + \|\nabla f_i(\mathbf{x}_i^{(t)})\|_2^2 - 2\langle \mathbf{g}_i^{(t)}, \nabla f_i(\mathbf{x}_i^{(t)}) \rangle)} \\
&= \frac{1}{n} \sum_{i=1}^n \sqrt{\mathbb{E}_t \|\mathbf{g}_i^{(t)} - \nabla f_i(\mathbf{x}_i^{(t)})\|_2^2} \\
&\leq \frac{\sigma_g}{\sqrt{b}}. \tag{32}
\end{aligned}$$

Bound the second term (29) as

$$\frac{1}{n} \sum_{i=1}^n \mathbb{E}_t \left\| \frac{\tau}{\tau + \|\nabla f_i(\mathbf{x}_i^{(t)})\|_2} (\mathbf{g}_i^{(t)} - \nabla f_i(\mathbf{x}_i^{(t)})) \right\|_2 \leq \frac{1}{n} \sum_{i=1}^n \frac{\tau \sigma_g / \sqrt{b}}{\tau + \|\nabla f_i(\mathbf{x}_i^{(t)})\|_2} \leq \frac{\sigma_g}{\sqrt{b}}. \tag{33}$$

Bound the third term (30) as

$$\begin{aligned}
&\frac{1}{n} \sum_{i=1}^n \left\| \left(\frac{\tau}{\tau + \|\nabla f_i(\mathbf{x}_i^{(t)})\|_2} - \frac{\tau}{\tau + \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2} \right) \nabla f_i(\mathbf{x}_i^{(t)}) \right\|_2 \\
&= \frac{1}{n} \sum_{i=1}^n \left\| \frac{\tau (\|\nabla f_i(\mathbf{x}_i^{(t)})\|_2 - \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2)}{(\tau + \|\nabla f_i(\mathbf{x}_i^{(t)})\|_2)(\tau + \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2)} \nabla f_i(\mathbf{x}_i^{(t)}) \right\|_2 \\
&\leq \frac{1}{n} \sum_{i=1}^n \frac{\tau \left| \|\nabla f_i(\mathbf{x}_i^{(t)})\|_2 - \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 \right|}{\tau + \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2} \\
&\leq \frac{1}{n} \sum_{i=1}^n \delta^{(t)} \left| \|\nabla f_i(\mathbf{x}_i^{(t)})\|_2 - \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 \right| + \frac{1}{n} \sum_{i=1}^n \delta^{(t)} \left| \|\nabla f_i(\bar{\mathbf{x}}^{(t)})\|_2 - \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 \right| \\
&\leq \frac{1}{n} \sum_{i=1}^n \delta^{(t)} L \|\mathbf{x}_i^{(t)} - \bar{\mathbf{x}}^{(t)}\|_2 + \frac{1}{n} \sum_{i=1}^n \delta^{(t)} \cdot \frac{1}{12} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 \\
&\leq \frac{\delta^{(t)} L}{\sqrt{n}} \|\mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top\|_F + \frac{1}{12} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2, \tag{34}
\end{aligned}$$

where we use $\delta^{(t)} \leq 1$ to reach the last inequality.

Bound (31) as

$$\begin{aligned}
&\left\| \frac{\tau}{\tau + \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2} \left(\frac{1}{n} \sum_{i=1}^n \nabla f_i(\mathbf{x}_i^{(t)}) - \nabla f(\bar{\mathbf{x}}^{(t)}) \right) \right\|_2 \\
&= \frac{\tau}{\tau + \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2} \left\| \nabla F(\mathbf{X}^{(t)}) \left(\frac{1}{n} \mathbf{1}_n \right) - \nabla f(\bar{\mathbf{x}}^{(t)}) \right\|_2 \\
&\leq \frac{\delta^{(t)} L}{\sqrt{n}} \|\mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top\|_F. \tag{35}
\end{aligned}$$

Using (32), (33), (34) and (35), the function value descent inequality (27) becomes

$$\mathbb{E}_t [f(\bar{\mathbf{x}}^{(t+1)}) - f(\bar{\mathbf{x}}^{(t)})] \leq -\eta \delta^{(t)} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2^2 + \frac{\eta^2 L}{2} \mathbb{E}_t \|\bar{\mathbf{v}}^{(t+1)}\|_2^2$$

$$\begin{aligned}
& + \eta \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 \mathbb{E}_t \|\text{Clip}_\tau(\nabla f(\bar{\mathbf{x}}^{(t)})) - \bar{\mathbf{g}}_\tau^{(t+1)}\|_2 \\
\leq & -\eta \delta^{(t)} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2^2 + \frac{\eta^2 L}{2} \mathbb{E}_t \|\bar{\mathbf{v}}^{(t+1)}\|_2^2 \\
& + \eta \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 \left(\frac{2\sigma_g}{\sqrt{b}} + \frac{1}{12} \delta^{(t)} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 + \frac{2\delta^{(t)} L}{\sqrt{n}} \|\mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top\|_{\mathbb{F}} \right) \\
= & -\frac{11}{12} \eta \delta^{(t)} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2^2 + \frac{\eta^2 L}{2} \mathbb{E}_t \|\bar{\mathbf{v}}^{(t+1)}\|_2^2 \\
& + \eta \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 \left(\frac{2\sigma_g}{\sqrt{b}} + \frac{2\delta^{(t)} L}{\sqrt{n}} \|\mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top\|_{\mathbb{F}} \right) \\
\leq & -\frac{5}{12} \eta \delta^{(t)} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2^2 + \frac{\eta^2 L}{2} \mathbb{E}_t \|\bar{\mathbf{v}}^{(t+1)}\|_2^2 \\
& + \frac{2\eta\sigma_g}{\sqrt{b}} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 + \frac{2\delta^{(t)} \eta L^2}{n} \|\mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top\|_{\mathbb{F}}^2, \tag{36}
\end{aligned}$$

where the last inequality is due to

$$\begin{aligned}
& \eta \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 \cdot \frac{2\delta^{(t)} L}{\sqrt{n}} \|\mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top\|_{\mathbb{F}} \\
\leq & \eta \delta^{(t)} \cdot 2 \cdot \sqrt{\frac{1}{2} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2^2} \cdot \sqrt{\frac{2L^2}{n} \|\mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top\|_{\mathbb{F}}^2} \\
\leq & \frac{1}{2} \eta \delta^{(t)} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2^2 + \frac{2\delta^{(t)} \eta L^2}{n} \|\mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top\|_{\mathbb{F}}^2.
\end{aligned}$$

C.2 Convergence rate

Different from (24), with the use of gradient clipping operator, we can only bound the expected norm of average gradient estimate as

$$\begin{aligned}
\mathbb{E} \|\bar{\mathbf{v}}^{(t)}\|_2^2 &= \mathbb{E} \|\bar{\mathbf{g}}_p^{(t)}\|_2^2 \\
&= \mathbb{E} \|\bar{\mathbf{g}}_\tau^{(t)} + \bar{\mathbf{e}}^{(t)}\|_2^2 \\
&= \mathbb{E} \|\bar{\mathbf{g}}_\tau^{(t)}\|_2^2 + \mathbb{E} \|\bar{\mathbf{e}}^{(t)}\|_2^2 \\
&\leq \tau^2 + \frac{\sigma_p^2 d}{n}. \tag{37}
\end{aligned}$$

Let $\Delta = \mathbb{E}[f(\bar{\mathbf{x}}^{(0)})] - f^*$. The techniques used is similar to that used in Appendix B, so that we can reuse results from Appendices B.2 and B.3, namely (23), in the following proof. Take full expectation and use (37), sum (36) over $t = 1, \dots, T$,

$$\begin{aligned}
-\Delta &\leq -\frac{5\eta}{12} \sum_{t=1}^T \mathbb{E} \left(\delta^{(t)} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2^2 \right) + \frac{2\eta\sigma}{\sqrt{b}} \sum_{t=1}^T \mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 \\
&\quad + \frac{2\eta L^2}{n} \sum_{t=1}^T \mathbb{E} \|\mathbf{X}^{(t)} - \bar{\mathbf{x}}^{(t)} \mathbf{1}_n^\top\|_{\mathbb{F}}^2 + \frac{\eta^2 L}{2} \sum_{t=1}^T \mathbb{E} \|\bar{\mathbf{v}}^{(t+1)}\|_2^2 \\
&\leq -\frac{5\eta}{12} \sum_{t=1}^T \mathbb{E} \left(\delta^{(t)} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2^2 \right) + \frac{2\eta\sigma}{\sqrt{b}} \sum_{t=1}^T \mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 \\
&\quad + \frac{2\eta L^2}{n} \left(\frac{8\eta^2}{(1-\hat{\alpha})^2} \sum_{t=1}^T n \mathbb{E} \|\bar{\mathbf{v}}^{(t)}\|_2^2 + \frac{1024\eta^2}{(1-\hat{\alpha})^4} T n (\tau^2 + \sigma_p^2 d) \right) + \frac{\eta^2 L}{2} \sum_{t=1}^T \mathbb{E} \|\bar{\mathbf{v}}^{(t+1)}\|_2^2 \\
&= -\frac{5\eta}{12} \sum_{t=1}^T \mathbb{E} \left(\delta^{(t)} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2^2 \right) + \frac{2\eta\sigma}{\sqrt{b}} \sum_{t=1}^T \mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2
\end{aligned}$$

$$+ \frac{16\eta^3 L^2}{(1-\hat{\alpha})^2} \left(\tau^2 + \frac{\sigma_p^2 d}{n} \right) + \frac{2048\eta^3 L^2}{(1-\hat{\alpha})^4} T(\tau^2 + \sigma_p^2 d) + \frac{\eta^2 L T}{2} \left(\tau^2 + \frac{\sigma_p^2 d}{n} \right). \quad (38)$$

To be able to analyze the expected clipped gradient norm $\mathbb{E}(\delta^{(t)} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2^2)$, we need to use convexity and monotonicity from Lemma 2.

Lemma 2. Let $g(x) = \frac{x}{c+x}$ and $h(x) = xg(x) = \frac{x^2}{c+x}$. When $x \geq 0$, $g(x)$ and $h(x)$ increase monotonically, while $g(x)$ is concave and $h(x)$ is convex.

Proof of Lemma 2. It is sufficient to prove Lemma 2 by evaluating the first-order and second-order derivatives of $g(x)$ and $h(x)$.

Because $g'(x) = \frac{(c+x)-x}{(c+x)^2} = \frac{c}{(c+x)^2} > 0$ and $h'(x) = g(x) + xg'(x) \geq 0$, $g(x)$ and $h(x)$ increase monotonically.

$g(x)$ is concave because $g''(x) - \frac{2c(c+x)}{(c+x)^4} = -\frac{2c}{(c+x)^3} < 0$.

$h(x)$ is convex because $h''(x) = 2g'(x) + xg''(x) = \frac{2c}{(c+x)^2} - \frac{2cx}{(c+x)^3} = \frac{2c^2}{(c+x)^3} > 0$. \square

Next, we substitute $\tau = \nu$ (cf. Theorem 2), and assume the following inequality

$$\mathbb{E}\|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 \geq \nu. \quad (39)$$

By Lemma 2, the expectation of clipped gradients can be bounded as

$$\begin{aligned} \mathbb{E}\left(\delta^{(t)} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2^2\right) &= \mathbb{E}\left(\frac{\tau \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2^2}{\tau + \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2}\right) \\ &\geq \frac{\tau (\mathbb{E}\|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2)^2}{\tau + \mathbb{E}\|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2} \\ &\geq \frac{\tau \nu}{\tau + \nu} \mathbb{E}\|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 \\ &= \frac{\nu}{2} \mathbb{E}\|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2. \end{aligned} \quad (40)$$

Using (37), (40) Assumptions 3 and 4, and set $b = 1$, we can further bound the RHS of (38) as

$$\begin{aligned} -\Delta &\leq -\frac{5\eta\nu}{24} \sum_{t=1}^T \mathbb{E}\|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 + \frac{2\eta\sigma_g}{\sqrt{b}} \sum_{t=1}^T \mathbb{E}\|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 \\ &\quad + \frac{16\eta^3 L^2}{(1-\hat{\alpha})^2} \left(\tau^2 + \frac{\sigma_p^2 d}{n} \right) + \frac{2048\eta^3 L^2}{(1-\hat{\alpha})^4} T(\tau^2 + \sigma_p^2 d) + \frac{\eta^2 L T}{2} \left(\tau^2 + \frac{\sigma_p^2 d}{n} \right) \end{aligned} \quad (41)$$

$$\begin{aligned} &\stackrel{(i)}{\leq} -\frac{\eta\nu}{8} \sum_{t=1}^T \mathbb{E}\|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 + \frac{2048T\eta^3 L^2}{(1-\hat{\alpha})^4} (\tau^2 + \sigma_p^2 d) + \frac{3T\eta^2 L}{1-\hat{\alpha}} \left(\tau^2 + \frac{\sigma_p^2 d}{n} \right) \\ &= -\frac{\eta\nu}{8} \sum_{t=1}^T \mathbb{E}\|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 + \frac{2048T\eta^3 L^2 \tau^2}{(1-\hat{\alpha})^4} (1 + T\phi_m^2) + \frac{3T\eta^2 L \tau^2}{1-\hat{\alpha}} (1 + T\phi_m^2), \end{aligned} \quad (42)$$

where we use the condition $\nu \geq 24\sigma_g$ to prove (i) and substitute $\sigma_p^2 = T\tau^2\phi_m^2/d$ to prove (42).

Reorganize terms, (42) can be further bounded as

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbb{E}\|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 &\leq \frac{8\Delta}{\eta\nu T} + \frac{16384\eta^2 L^2 \nu}{(1-\hat{\alpha})^4} (1 + T\phi_m^2) + \frac{24\eta L \nu}{1-\hat{\alpha}} (1 + T\phi_m^2) \\ &\stackrel{(i)}{\leq} \frac{8\Delta\phi_m}{\eta\nu} + \frac{32768\eta^2 L^2 \nu}{(1-\hat{\alpha})^4} + \frac{48\eta L \nu}{1-\hat{\alpha}} \\ &\stackrel{(ii)}{\leq} \frac{8\Delta\phi_m}{\eta\nu} + \frac{4096\eta L \nu}{(1-\hat{\alpha})^{\frac{8}{3}}} + \frac{48\eta L \nu}{1-\hat{\alpha}}, \end{aligned}$$

$$\leq \frac{8\Delta\phi_m}{\eta\nu} + \frac{4144\eta L\nu}{(1-\hat{\alpha})^{\frac{8}{3}}}, \quad (43)$$

where we substitute $T = \phi_m^{-2}$ to prove (i), and use (25) for (ii).

Set $\eta = \frac{\gamma^{\frac{8}{3}}(1-\alpha)^{\frac{8}{3}}}{8288L}$ and $\gamma = \frac{1}{100}(1-\alpha)\rho$, (43) can be further bounded as

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 &< \frac{8288}{\gamma^{\frac{8}{3}}(1-\alpha)^{\frac{8}{3}}} \cdot \frac{192L\Delta\phi_m}{\nu} + \frac{\gamma^{\frac{8}{3}}(1-\alpha)^{\frac{8}{3}}}{8288} \cdot \frac{4144\nu}{(1-\hat{\alpha})^{\frac{8}{3}}} \\ &= \frac{8288}{\gamma^{\frac{8}{3}}(1-\alpha)^{\frac{8}{3}}} \cdot \frac{192L\Delta\phi_m}{\nu} + \frac{\nu}{2} \\ &< \frac{1784\sqrt{L\Delta\phi_m}}{\gamma^{\frac{4}{3}}(1-\alpha)^{\frac{4}{3}}}. \end{aligned} \quad (44)$$

Choosing $\nu = \frac{1784\sqrt{L\Delta\phi_m}}{\gamma^{\frac{4}{3}}(1-\alpha)^{\frac{4}{3}}}$, (44) simplifies to $\frac{1}{T} \sum_{t=1}^T \mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 < \nu$, which further implies that there exists some $t \in [T]$ such that $\mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 < \nu$. However, this contradicts the assumption (39), which leads to the convergence results in the theorem.

Lastly, we can verify conditions (12), (16) and (25) are all met, which concludes the proof.

D Proof of Theorem 4

This section proves Theorem 4 based on results from Appendix C. We first assume all expected gradient norm $\mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2$ are greater than a threshold ν (i.e. $\mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 \geq \nu$ for all $t = 1, \dots, T$). Set $b = \left(\frac{24\sigma_d}{\nu}\right)^2$ and $\sigma_p = 0$ in (41), we can reach the following descent inequality

$$-\Delta \leq -\frac{\eta\nu}{8} \sum_{t=1}^T \mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 + \frac{2048T\eta^3 L^2 \tau^2}{(1-\hat{\alpha})^4} + \frac{3T\eta^2 L \tau^2}{1-\hat{\alpha}}. \quad (45)$$

Reorganize terms,

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 &\leq \frac{8\Delta}{\eta\nu T} + \frac{16384\eta^2 L^2 \nu}{(1-\hat{\alpha})^4} + \frac{24\eta L\nu}{1-\hat{\alpha}} \\ &< \frac{8\Delta}{\eta\nu T} + \frac{2048\eta L\nu}{(1-\hat{\alpha})^{\frac{8}{3}}} + \frac{24\eta L\nu}{1-\hat{\alpha}}, \\ &\leq \frac{8\Delta}{\eta\nu T} + \frac{2072\eta L\nu}{(1-\hat{\alpha})^{\frac{8}{3}}}, \end{aligned}$$

where we use (25) for the second inequality.

Set $\eta = \frac{1}{L} \cdot \sqrt{\frac{8}{2072}}$ and $\nu = \sqrt{\frac{L\Delta}{(1-\hat{\alpha})^{\frac{8}{3}}T}}$, the average ℓ_2 norm of gradients can be bounded as

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 < \sqrt{\frac{L\Delta}{(1-\hat{\alpha})^{\frac{8}{3}}T}} = \nu,$$

which implies that there exists some $t \in [T]$, such that $\mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 < \nu$ which contradicts the assumption, and proves that PORTER-GC reaches $\mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(t)})\|_2 \leq \nu$ within T iterations