
On the Linear Speedup Analysis of Communication Efficient Momentum SGD for Distributed Non-Convex Optimization

Hao Yu¹ Rong Jin¹ Sen Yang¹

Abstract

Recent developments on large-scale distributed machine learning applications, e.g., deep neural networks, benefit enormously from the advances in distributed non-convex optimization techniques, e.g., distributed Stochastic Gradient Descent (SGD). A series of recent works study the linear speedup property of distributed SGD variants with reduced communication. The linear speedup property enable us to scale out the computing capability by adding more computing nodes into our system. The reduced communication complexity is desirable since communication overhead is often the performance bottleneck in distributed systems. Recently, momentum methods are more and more widely adopted in training machine learning models and can often converge faster and generalize better. For example, many practitioners use distributed SGD with momentum to train deep neural networks with big data. However, it remains unclear whether any distributed momentum SGD possesses the same linear speedup property as distributed SGD and has reduced communication complexity. This paper fills the gap by considering a distributed communication efficient momentum SGD method and proving its linear speedup property.

1. Introduction

Consider distributed non-convex optimization scenarios where N workers jointly solve the following consensus optimization problem:

$$\min_{\mathbf{x} \in \mathbb{R}^m} f(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N \underbrace{\mathbb{E}_{\xi_i \sim \mathcal{D}_i} [F_i(\mathbf{x}; \xi_i)]}_{\triangleq f_i(\mathbf{x})} \quad (1)$$

¹Machine Intelligence Technology Lab, Alibaba Group (U.S.) Inc., Bellevue, WA. Correspondence to: Hao Yu <eeyuhao@gmail.com>.

where $f_i(\mathbf{x}) \triangleq \mathbb{E}_{\xi_i \sim \mathcal{D}_i} [F_i(\mathbf{x}; \xi_i)]$ are smooth non-convex functions with possibly different distributions \mathcal{D}_i . Problem (1) is particularly important in deep learning since it captures data parallelism for training deep neural networks. In deep learning with data parallelism, each $F_i(\cdot; \cdot) = F(\cdot; \cdot)$ represents the common deep neural network to be jointly trained and each \mathcal{D}_i represents the distribution of the local data set accessed by worker i . The scenario where each \mathcal{D}_i are different also captures the **federated learning** setting recently proposed in (McMahan et al., 2017) where mobile clients with private local training data and slow intermittent network connections cooperatively train a high-quality centralized model.

Stochastic Gradient Descent (SGD) (and its momentum variants) have been the dominating methodology for solving machine learning problem. For large-scale distributed machine learning problems, such as training deep neural networks, a parallel version of SGD, known as parallel mini-batch SGD, is widely adopted (Dean et al., 2012; Dekel et al., 2012; Li et al., 2014). With N parallel workers, parallel mini-batch SGD can solve problem (1) with $O(1/\sqrt{NT})$ convergence, which is N times faster¹ than the $O(1/\sqrt{T})$ convergence attained by SGD over a single node (Ghadimi & Lan, 2013; Lian et al., 2015). Obviously, such linear speedup with respect to (w.r.t.) number of workers is desired in distributed training as it implies perfect computation scalability by increasing the number of used workers. However, such linear speedup is difficult to harvest in practice because the classical parallel mini-batch SGD requires all workers to synchronize local gradients or models at **every** iteration such that inter-node communication cost becomes the performance bottleneck. To eliminate potential communication bottlenecks, such as high latency or low bandwidths, in computing infrastructures, many distributed SGD variants have been recently proposed. For example, (Lian et al., 2017; Jiang et al., 2017; Assran et al., 2018) consider decentralized parallel SGD where global gradient aggregations

¹If an algorithm has $O(1/\sqrt{T})$ convergence, then it takes $1/\epsilon^2$ iterations to attain an $O(\epsilon)$ accurate solution. Similarly, if another algorithm has $O(1/\sqrt{NT})$ convergence, then it takes $1/(N\epsilon^2)$ iterations, which is N times smaller than $1/\epsilon^2$, to attain an $O(\epsilon)$ accurate solution. In this sense, the second algorithm is N times faster than the first one.

used in the classical parallel mini-batch SGD are replaced with local aggregations between neighbors. To reduce the communication cost at each iteration, compression or sparsification based parallel SGD are considered in (Seide et al., 2014; Strom, 2015; Aji & Heafield, 2017; Wen et al., 2017; Alistarh et al., 2017). Recently, (Yu et al., 2018; Wang & Joshi, 2018a; Jiang & Agrawal, 2018) prove that certain parallel SGD variants that strategically skip communication rounds can achieve the fast $O(1/\sqrt{NT})$ convergence with significantly less communication rounds. See Table 1 for a summary on recent works studying distributed SGD with $O(1/\sqrt{NT})$ convergence and reduced communication complexity.

It is worth noting that existing convergence analyses on distributed methods for solving problem (1) focus on parallel SGD **without momentum**. In practice, momentum SGD is more commonly used to train deep neural networks since it often can converge faster and generalize better (Krizhevsky et al., 2012; Sutskever et al., 2013; Yan et al., 2018). For example, momentum SGD is suggested for training ResNet for image classification tasks to obtain the best test accuracy (He et al., 2016). See Figure 1 for the comparison of test accuracy between training with momentum and training without momentum.² In fact, while previous works (Lian et al., 2017; Stich, 2018; Yu et al., 2018; Jiang & Agrawal, 2018) prove that distributed vanilla SGD (without momentum) can train deep neural networks with $O(1/\sqrt{NT})$ convergence using significantly fewer communications rounds, most of their experiments use momentum SGD rather than vanilla SGD (to achieve the targeted test accuracy). In addition, previous empirical works (Chen & Huo, 2016; Lin et al., 2018) on distributed training for deep neural networks explicitly observe that momentum is necessary to improve the convergence and test accuracy. In this perspective, there is a huge gap between the current practices, i.e., using momentum SGD rather than vanilla SGD in distributed training for deep neural networks, and existing theoretical analyses such as (Yu et al., 2018; Wang & Joshi, 2018a; Jiang & Agrawal, 2018) studying the convergence rate and communication complexity of SGD without momentum. Momentum methods are originally proposed by Polyak in (Polyak, 1964) to minimize deterministic strongly convex quadratic functions. Its convergence rate for deterministic convex optimization, which is not necessarily strongly convex, is established in (Ghadimi et al., 2014). For non-convex optimization, the convergence for deterministic non-convex

²As observed in Figure 1, the final test accuracy of momentum SGD is roughly 2.5% better than that of vanilla SGD (without momentum) over CIFAR10. In practice, people usually use momentum SGD to train ResNet in both single GPU and multiple GPU scenarios. Some practitioners conjecture that it is possible to avoid the performance degradation of vanilla SGD if its hyper-parameters are well tuned. However, even if this conjecture is true, the hyper-parameter tuning can be extremely time-consuming.

optimization is proven in (Zavriev & Kostyuk, 1993) and the $O(1/\sqrt{T})$ convergence rate for stochastic non-convex optimization is recently shown in (Yan et al., 2018). However, to our best knowledge, it remains as an open question: “Whether any distributed momentum SGD can achieve the same $O(1/\sqrt{NT})$ convergence, i.e., linear speedup w.r.t. number of workers, with reduced communication complexity as SGD (without momentum) methods in (Yu et al., 2018; Wang & Joshi, 2018a; Jiang & Agrawal, 2018) ?”




Figure 1: Test accuracy performance of ResNet56 over CIFAR10 when trained with SGD with and without momentum on a single GPU using hyper-parameters suggested in (He et al., 2016).

Our Contributions: In this paper, we consider parallel restarted SGD with momentum (described in Algorithm 1), which can be viewed as the momentum extension of parallel restarted SGD, also referred to as local SGD, considered in (Stich, 2018; Yu et al., 2018). Such a method is also faithful to the common practice “model averaging with momentum” used by practitioners for training deep neural networks in (Chen & Huo, 2016; McMahan et al., 2017; Su et al., 2018; Lin et al., 2018). We show that parallel restarted SGD with momentum can solve problem (1) with $O(1/\sqrt{NT})$ convergence, i.e., achieving linear speedup w.r.t. number of workers. Moreover, to achieve the fast $O(1/\sqrt{NT})$ convergence, T iterations of our algorithm only requires $O(N^{3/2}T^{1/2})$ communication rounds when all workers can access identical data sets; or $O(N^{3/4}T^{3/4})$ communication rounds when workers access non-identical data sets. To our knowledge, this is the first time that a distributed momentum SGD method for non-convex stochastic optimization is proven to possess the same linear speedup property (with communication reduction) as distributed SGD (without momentum) in (Lian et al., 2017; Yu et al., 2018; Wang & Joshi, 2018a; Jiang & Agrawal, 2018).

Recall that momentum SGD degrades to vanilla SGD when its momentum coefficient is set 0. Algorithm 1 with $\beta = 0$ degrades to the parallel SGD methods (without momentum) considered in (Yu et al., 2018; Wang & Joshi, 2018a;

Table 1: Number of communication rounds involved in T iterations of existing distributed SGD **without momentum** with $O(1/\sqrt{NT})$ convergence for **non-convex** stochastic optimization (1).

REFERENCE	IDENTICAL $f_i(\mathbf{x})$, I.E., $\kappa = 0$	NON-IDENTICAL $f_i(\mathbf{x})$, I.E., $\kappa \neq 0$	EXTRA ASSUMPTIONS
(LIAN ET AL., 2017)	$O(T)$	$O(T)$	NO
(YU ET AL., 2018)	$O(N^{3/4}T^{3/4})$	$O(N^{3/4}T^{3/4})$	BOUNDED GRADIENTS
(JIANG & AGRAWAL, 2018)	$O(N^{5/2}T^{1/2})$	$O(N^{5/4}T^{3/4})$	NO
(WANG & JOSHI, 2018A)	$O(N^{3/2}T^{1/2})$	NOT APPLICABLE	NO
THIS PAPER	$O(N^{3/2}T^{1/2})$	$O(N^{3/4}T^{3/4})$	NO

Jiang & Agrawal, 2018). Our communication complexity results for parallel SGD without momentum cases also improve the state-of-the-art. As shown in Table 1, the number of required communication rounds shown in this paper is the fewest in both identical training data set case and non-identical data set case. In particular, this paper relaxes the bounded gradient moment assumption used in (Yu et al., 2018) to a milder bounded variance assumption and reduces the communication complexity for the case with identical training data. Our analysis applies to the distributed training scenarios where workers access non-identical training sets, e.g., training with sharded data or federated learning, that can not be handled in (Wang & Joshi, 2018a).

This paper further considers momentum SGD with decentralized communication and proves that it possess the same linear speedup property as its vanilla SGD (without momentum) counterpart considered in (Lian et al., 2017).

2. Parallel Restarted SGD with Momentum

2.1. Preliminary

Following the convention in (distributed) stochastic optimization, we assume each worker can independently evaluate unbiased stochastic gradients $\nabla F_i(\mathbf{x}_i; \xi_i)$, $\xi_i \sim \mathcal{D}_i$ using its own local variable \mathbf{x}_i . Throughout this paper, we have the following standard assumption:

Assumption 1. *Problem (1) satisfies the following:*

- Smoothness:** Each function $f_i(\mathbf{x})$ in problem (1) is smooth with modulus L .
- Bounded variances:** There exist $\sigma > 0$ and $\kappa > 0$ such that

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

$$\frac{1}{N} \sum_{i=1}^N \|\nabla f_i(\mathbf{x}) - \nabla f(\mathbf{x})\|^2 \leq \kappa^2, \forall \mathbf{x} \quad (3)$$

Note that σ^2 in Assumption 1 quantifies the variance of stochastic gradients at local worker, and κ^2 quantifies the deviations between each worker’s local objective function $f_i(\mathbf{x})$. For distributed training of neural networks, if all

workers can access the same data set, then $\kappa = 0$. Assumption 1 was previously used in (Lian et al., 2017; Wen et al., 2017; Alistarh et al., 2017; Jiang & Agrawal, 2018). The bounded variance assumption in Assumption 1 is milder than the bounded second moment assumption used in (Stich, 2018; Yu et al., 2018).

Recall that if function $f(\cdot)$ is smooth with modulus L , then we have the key property that $f(\mathbf{x}) \leq f(\mathbf{y}) + \langle \nabla f(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle + \frac{L}{2} \|\mathbf{y} - \mathbf{x}\|^2, \forall \mathbf{x}, \mathbf{y}$ where $\langle \cdot, \cdot \rangle$ represents the inner product of two vectors. Another two useful properties implied by Assumption 1 are summarized in the next two lemmas.

Lemma 1. *Consider problem (1) under Assumption 1. Let \mathbf{g}_i be mutually independent unbiased stochastic gradients sampled at points \mathbf{x}_i such that $\mathbb{E}[\mathbf{g}_i] = \nabla f_i(\mathbf{x}_i), \forall i \in \{1, 2, \dots, N\}$. Then, we have*

$$\mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i\|^2] \leq \frac{1}{N} \sigma^2 + \mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i)\|^2]$$

Proof. See Supplement 6.1 □

Since $\mathbb{E}[\mathbf{g}_i] = \nabla f_i(\mathbf{x}_i)$, we have $\mathbb{E}[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i] = \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i)$. The implication of Lemma 1 is: while \mathbf{g}_i are sampled at different points \mathbf{x}_i , the variance of $\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i$, which is equal to $\mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i - \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i)\|^2] = \mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i\|^2] - \mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i)\|^2]$, scales down by N times when N workers sample gradients independently (even though at different \mathbf{x}_i). Note that if each \mathbf{g}_i are independently sampled by N workers at the same point $\mathbf{x}_i \equiv \mathbf{x}, \forall i$ (as in the classical parallel mini-batch SGD), then Lemma 1 reduces to $\mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i - \nabla f(\mathbf{x})\|^2] \leq \frac{\sigma^2}{N}$, which is the fundamental reason why the classical parallel mini-batch SGD can converges N times faster with N workers. However, to reduce the communication overhead associated with synchronization between workers, we shall consider distributed algorithms with fewer synchronization rounds such that \mathbf{x}_i at different workers can possibly deviate from

each other. In this case, huge deviations across \mathbf{x}_i can cause that \mathbf{g}_i sampled at \mathbf{x}_i provides unreliable or even contradicting gradient knowledge and hence slow down the convergence. Thus, to let N workers jointly solve problem (1) with fast convergence, the distributed algorithm needs to enforce certain concentration of \mathbf{g}_i . The following useful lemma relates the concentration of $\nabla f_i(\mathbf{x})$ with deviations across all \mathbf{x}_i .

Lemma 2. *Consider problem (1) under Assumption 1. For any N points $\mathbf{x}_i, i \in \{1, 2, \dots, N\}$, if we define $\bar{\mathbf{x}} = \frac{1}{N} \sum_{i=1}^N \mathbf{x}_i$, then we have*

$$\begin{aligned} & \frac{1}{N} \sum_{i=1}^N \|\nabla f_i(\mathbf{x}_i) - \frac{1}{N} \sum_{j=1}^N \nabla f_j(\mathbf{x}_j)\|^2 \\ & \leq \frac{6L^2}{N} \sum_{i=1}^N \|\mathbf{x}_i - \bar{\mathbf{x}}\|^2 + \frac{3}{N} \sum_{i=1}^N \|\nabla f_i(\bar{\mathbf{x}}) - \nabla f(\bar{\mathbf{x}})\|^2 \end{aligned}$$

Proof. See Supplement 6.2 \square

2.2. Parallel Restarted SGD with Momentum

2.2.1. ALGORITHM 1 WITH POLYAK’S MOMENTUM

Consider applying the distributed algorithm described in Algorithm 1 to solve problem (1) with N workers. In the literature, there are two momentum methods for SGD, i.e., Polyak’s momentum method (also known as the heavy ball method) and Nesterov’s momentum method. Algorithm 1 can use either of them for local worker updates by providing two options: “Option I” is Polyak’s momentum; “Option II” is Nesterov’s momentum.³ We also note that the update steps described in (4) or (5) can be locally performed at each worker in parallel.

The synchronization step (6) can be interpreted as restarting momentum SGD with **new** initial point $\hat{\mathbf{x}}$ and momentum buffer $\hat{\mathbf{u}}$, i.e., resetting $\mathbf{u}_i^{(t)} = \hat{\mathbf{u}}$ and $\mathbf{x}_i^{(t)} = \hat{\mathbf{x}}$. In this perspective, we call Algorithm 1 “parallel restarted SGD with momentum”. Note that if we choose $\beta = 0$ in Algorithm 1, then it degrades to the “parallel restarted SGD”, also known as “local SGD” or “SGD with periodic averaging”, in (Stich, 2018; Yu et al., 2018; Lin et al., 2018; Wang & Joshi, 2018a; Zhou & Cong, 2018; Jiang & Agrawal, 2018).

Since inter-node communication is only needed by Algorithm 1 to calculate global averages $\hat{\mathbf{x}}$ and $\hat{\mathbf{u}}$ in (6) and happens only once every I iterations, the total number of inter-communication rounds involved in T iterations of Algorithm 1 is given by T/I . If we use **all-reduce** operations

³The current description in (4) and (5) is identical to the default implementations of momentum methods in PyTorch’s SGD optimizer. Polyak’s and Nesterov’s momentum have many other equivalent representations. These equivalent representations will be further discussed later.

Algorithm 1 Parallel Restarted SGD with momentum (PR-SGD-Momentum)

1: **Input:** Initialize local solutions $\mathbf{x}_i^{(0)} = \hat{\mathbf{x}} \in \mathbb{R}^m$ and momentum buffers $\mathbf{u}_i^{(0)} = \mathbf{0}, \forall i \in \{1, 2, \dots, N\}$. Set learning rate $\gamma > 0$, momentum coefficient $\beta \in [0, 1)$, node synchronization interval $I > 0$ and number of iterations T

2: **for** $t = 1, 2, \dots, T - 1$ **do**

3: Each worker i samples its stochastic gradient $\mathbf{g}_i^{(t-1)} = \nabla F_i(\mathbf{x}_i^{(t-1)}; \xi_i^{(t-1)})$ with $\xi_i^{(t-1)} \sim \mathcal{D}_i$.

4: Each worker i in parallel updates its local momentum buffer and solution via

$$\text{Option I: } \begin{cases} \mathbf{u}_i^{(t)} = \beta \mathbf{u}_i^{(t-1)} + \mathbf{g}_i^{(t-1)} \\ \mathbf{x}_i^{(t)} = \mathbf{x}_i^{(t-1)} - \gamma \mathbf{u}_i^{(t)} \end{cases} \quad \forall i. \quad (4)$$

Or

$$\text{Option II: } \begin{cases} \mathbf{u}_i^{(t)} = \beta \mathbf{u}_i^{(t-1)} + \mathbf{g}_i^{(t-1)} \\ \mathbf{v}_i^{(t)} = \beta \mathbf{u}_i^{(t)} + \mathbf{g}_i^{(t-1)} \\ \mathbf{x}_i^{(t)} = \mathbf{x}_i^{(t-1)} - \gamma \mathbf{v}_i^{(t)} \end{cases} \quad \forall i. \quad (5)$$

5: **if** t is a multiple of I , i.e., $t \bmod I = 0$, **then**

6: Each worker resets its momentum buffer and local solution as the node averages via

$$\begin{cases} \mathbf{u}_i^{(t)} = \hat{\mathbf{u}} \triangleq \frac{1}{N} \sum_{j=1}^N \mathbf{u}_j^{(t)} \\ \mathbf{x}_i^{(t)} = \hat{\mathbf{x}} \triangleq \frac{1}{N} \sum_{j=1}^N \mathbf{x}_j^{(t)} \end{cases} \quad \forall i \quad (6)$$

7: **end if**

8: **end for**

to compute the global averages, the per-round communication is relatively cheap and does not involve a centralized parameter server (Goyal et al., 2017).

Algorithm 1 uses $\{\mathbf{x}_i^{(t)}\}_{t \geq 0}$ to denote local solution sequences generated at each worker i . For each iteration t , we define

$$\bar{\mathbf{x}}^{(t)} \triangleq \frac{1}{N} \sum_{i=1}^N \mathbf{x}_i^{(t)} \quad (7)$$

as the averages of local solutions $\mathbf{x}_i^{(t)}$ across all N nodes. Our performance analysis will be performed regarding the aggregated version $\bar{\mathbf{x}}^{(t)}$ defined in (7) so that “consensus” is no longer our concern for problem (1). Performing convergence analysis for the node-average version $\bar{\mathbf{x}}^{(t)}$ has been an important technique used in previous works on distributed consensus optimization (Lian et al., 2017; Mania et al., 2017; Stich, 2018; Yu et al., 2018; Jiang & Agrawal, 2018).

2.3. Performance Analysis

This subsection analyzes the convergence rates of Algorithm 1 with Polyak's and Nesterov's momentum, respectively.

We first consider Algorithm 1 with Option I given by (4), i.e., Polyak's momentum. Polyak's momentum SGD, also known as the heavy ball method, is the classical momentum SGD used for training deep neural networks and often provides fast convergence and good generalization (Krizhevsky et al., 2012; Sutskever et al., 2013; Yan et al., 2018). If we let $\mathbf{x}_i^{(-1)} = \mathbf{x}_i^{(0)}$, then (4) in Algorithm 1 can be equivalently written as the following single variable version

$$\mathbf{x}_i^{(t)} = \mathbf{x}_i^{(t-1)} - \gamma \mathbf{g}_i^{(t-1)} + \beta [\mathbf{x}_i^{(t-1)} - \mathbf{x}_i^{(t-2)}] \quad (8)$$

where the last term $\beta [\mathbf{x}_i^{(t-1)} - \mathbf{x}_i^{(t-2)}]$ is often called Polyak's momentum term.

It's interesting to note that if we define an auxiliary sequence

$$\bar{\mathbf{u}}^{(t)} \triangleq \frac{1}{N} \sum_{i=1}^N \mathbf{u}_i^{(t)} \quad (9)$$

which is the node average sequence of local buffer variables $\mathbf{u}_i^{(t)}$ from Algorithm 1 with Option I, then we have

$$\begin{cases} \bar{\mathbf{u}}^{(t)} &= \beta \bar{\mathbf{u}}^{(t-1)} + \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t-1)} \\ \bar{\mathbf{x}}^{(t)} &= \bar{\mathbf{x}}^{(t-1)} - \gamma \bar{\mathbf{u}}^{(t)} \end{cases}, \forall t \geq 1 \quad (10)$$

where $\bar{\mathbf{x}}^{(t)}$ is defined in (7).

An important observation is that if $\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t-1)}$ in (10) is replaced with $\nabla f(\bar{\mathbf{x}}^{(t-1)})$, then (10) is exactly a standard **single-node** SGD momentum method with momentum buffer $\bar{\mathbf{u}}^{(t)}$ and solution $\bar{\mathbf{x}}^{(t)}$. That is, under Algorithm 1, N workers essentially jointly update $\bar{\mathbf{u}}^{(t)}$ and $\bar{\mathbf{x}}^{(t)}$ with momentum SGD using an **inaccurate** stochastic gradient $\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t-1)}$. However, since (6) periodically (every I iterations) synchronizes all local variables $\mathbf{u}_i^{(t)}$ and $\mathbf{x}_i^{(t)}$, our intuition is by synchronizing frequently enough the inaccuracy of the used gradient counterpart in (10) shall not damage the convergence too much. The next theorem summarizes the convergence rate of Algorithm 1 and characterizes the effect of synchronization interval I in its convergence.

Remark 1. *Our Algorithm 1 is different from a common heuristic model averaging strategy for momentum SGD suggested in (Wang & Joshi, 2018b) and in Microsoft's CNTK framework (Seide & Agarwal, 2016). The strategy in (Seide & Agarwal, 2016; Wang & Joshi, 2018b) let each worker run local momentum SGD in parallel, and periodically reset momentum buffer variables to zero and average local models. However, there is no convergence analysis for this strategy. In contrast, this paper shall provide rigorous convergence analysis for our Algorithm 1. Our experiment in*

Supplement 6.7.2 further shows that Algorithm 1 has faster convergence than the strategy in (Seide & Agarwal, 2016; Wang & Joshi, 2018b) and yields better test accuracy when used to train ResNet for CIFAR10.

Theorem 1. *Consider problem (1) under Assumption 1. If we choose $\gamma \leq \frac{(1-\beta)^2}{(1+\beta)L}$ and $I \leq \frac{1-\beta}{6L\gamma}$ in Algorithm 1 with **Option I**, then for all $T \geq 1$, we have*

$$\begin{aligned} & \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2] \\ & \leq \frac{2(1-\beta)}{\gamma T} (f(\bar{\mathbf{x}}^{(0)}) - f^*) + \frac{L\gamma}{(1-\beta)^2} \frac{\sigma^2}{N} + \frac{4L^2\gamma^2 I \sigma^2}{(1-\beta)^2} + \frac{9L^2\gamma^2 I^2 \kappa^2}{(1-\beta)^2} \\ & = O\left(\frac{1}{\gamma T}\right) + O\left(\frac{\gamma}{N}\sigma^2\right) + O(\gamma^2 I \sigma^2) + O(\gamma^2 I^2 \kappa^2) \end{aligned}$$

where $\{\bar{\mathbf{x}}^{(t)}\}_{t \geq 0}$ is the sequence defined in (7); and σ and κ are constants defined in Assumption 1; and f^* is the minimum value of problem (1).

Proof. See Supplement 6.3. \square

The next corollary summarizes that Algorithm 1 with Option I using N workers can solve problem (1) with the fast $O\left(\frac{1}{\sqrt{NT}}\right)$ convergence, i.e., achieving the linear speedup (w.r.t. number of workers). Note that $O\left(\frac{1}{\sqrt{NT}}\right)$ dominates $O\left(\frac{N}{T}\right)$ in Corollary 1 when T is sufficiently large.

Corollary 1 (Linear Speedup). *Consider problem (1) under Assumption 1. If we choose $\gamma = \frac{\sqrt{N}}{\sqrt{T}}$ and $I = 1$ in Algorithm 1 with **Option I**, then for any $T \geq \frac{36L^2N}{(1-\beta)^2}$, we have*

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2] = O\left(\frac{1}{\sqrt{NT}}\right) + O\left(\frac{N}{T}\right). \quad (11)$$

Proof. This corollary follows because the selection of γ and I satisfies the conditions in Theorem 1. \square

The next corollary summarizes that the desired linear speedup can be attained with communication skipping, i.e., using $I > 1$ in Algorithm 1. In particular, to achieve the linear speedup, T iterations in Algorithm 1 with Option I only require $O(N^{3/2}T^{1/2})$ communication rounds when $\kappa = 0$, i.e., workers access the common training data set; or only require $O(N^{3/4}T^{3/4})$ communication rounds when $\kappa \neq 0$, i.e., workers access non-identical training data sets.

Corollary 2 (Linear Speedup with Reduced Communication). *Consider problem (1) under Assumption 1.*

- **Case $\kappa = 0$:** *If we choose $\gamma = \frac{\sqrt{N}}{\sqrt{T}}$ and $I \leq \frac{1-\beta}{6L} \frac{\sqrt{T}}{N^{3/2}} = O\left(\frac{T^{1/2}}{N^{3/2}}\right)$ in Algorithm 1 with **Option I**, then for any $T \geq \frac{(1+\beta)^2 L^2}{(1-\beta)^4} N$, we have*

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2] = O\left(\frac{1}{\sqrt{NT}}\right). \quad (12)$$

By using $I = O(\frac{T^{1/2}}{N^{3/2}})$, the total number of inter-node communication rounds involved in Algorithm 1 is $O(N^{3/2}T^{1/2})$.

- **Case $\kappa \neq 0$:** If we choose $\gamma = \frac{\sqrt{N}}{\sqrt{T}}$ and $I \leq \frac{1-\beta}{6L} \frac{T^{1/4}}{N^{3/4}} = O(\frac{T^{1/4}}{N^{3/4}})$ in Algorithm 1 with **Option I**, then for any $T \geq \frac{(1+\beta)^2 L^2}{(1-\beta)^4} N$, we have

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2] = O\left(\frac{1}{\sqrt{NT}}\right). \quad (13)$$

By using $I = O(\frac{T^{1/4}}{N^{3/4}})$, the total number of inter-node communication rounds involved in Algorithm 1 is $O(N^{3/4}T^{3/4})$.

Proof. This corollary follows simply because the selection of γ and I satisfies the conditions in Theorem 1. \square

Remark 2. Following the convention in non-convex optimization, the convergence rate in Corollary 2 is measured by the expected squared gradient norm used in (Ghadimi & Lan, 2013; Lian et al., 2017; Yu et al., 2018; Jiang & Agrawal, 2018). To attain the average $\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2]$ in expectation, one neat strategy suggested in (Ghadimi & Lan, 2013) is to generate a random iteration count R uniformly from $\{0, 1, \dots, T-1\}$, terminate Algorithm 1 at iteration R and output $\bar{\mathbf{x}}^{(R)}$ as the solution.

Remark 3. Note that Theorem 1 and Corollary 2 hold for any $0 \leq \beta < 1$. Recall that Algorithm 1 with $\beta = 0$ degrades to parallel restarted SGD (without momentum). For parallel SGD (without momentum), the communication complexity implied by Corollary 2 improves the state-of-the-art as summarized in Table 1.

2.3.1. ALGORITHM 1 WITH NESTEROV'S MOMENTUM

Now consider using Nesterov's momentum described in Option II in Algorithm 1. Option II given by (5) introduces extra auxiliary variables $\bar{\mathbf{v}}_i^{(t)}$ and uses them in the update of local solutions $\mathbf{x}_i^{(t)}$. It is not difficult⁴ to show that (5) yields the same solution sequences $\{\mathbf{x}_i^{(t)}\}_{t \geq 0}$ as

$$\begin{cases} \mathbf{y}_i^{(t)} = \mathbf{x}_i^{(t-1)} - \gamma \mathbf{g}_i^{(t-1)} \\ \mathbf{x}_i^{(t)} = \mathbf{y}_i^{(t)} + \beta [\mathbf{y}_i^{(t)} - \mathbf{y}_i^{(t-1)}] \end{cases} \quad \forall i \quad (14)$$

with $\mathbf{y}_i^{(0)} \triangleq \mathbf{0}, \forall i$. The only difference between (5) and (14) is the adoption of different cache variables.

Equation (14) is more widely used to describe SGD with Nesterov's momentum in the literature (Nesterov, 2004). However, by writing Nesterov's momentum SGD as (5), an important observation is that momentum buffer variables

⁴The derivation on the equivalence is in Supplement 6.4.

$\mathbf{u}_i^{(t)}$ in Polyak's and Nesterov's momentum methods evolve according to the same dynamic (with stochastic gradients sampled at different points). By adapting the convergence rate analysis for Polyak's momentum, Theorem 2 (formally proven in Supplement 6.5) summarizes the convergence for Nesterov's momentum.

Theorem 2. Consider problem (1) under Assumption 1. If we choose $\gamma \leq \frac{(1-\beta)^2}{L(1+\beta^3)}$ and $I \leq \frac{1-\beta}{6L\gamma}$ in Algorithm 1 with **Option II**, then for all $T > 0$, we have

$$\begin{aligned} & \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2] \\ & \leq \frac{2(1-\beta)}{\gamma T} (f(\bar{\mathbf{x}}^{(0)}) - f^*) + \frac{L\gamma}{(1-\beta)^2} \frac{\sigma^2}{N} + \frac{4L^2\gamma^2 I \sigma^2}{(1-\beta)^2} + \frac{9L^2\gamma^2 I^2 \kappa^2}{(1-\beta)^2} \\ & = O\left(\frac{1}{\gamma T}\right) + O\left(\frac{\gamma}{N}\sigma^2\right) + O(\gamma^2 I \sigma^2) + O(\gamma^2 I^2 \kappa^2) \end{aligned}$$

where $\{\bar{\mathbf{x}}^{(t)}\}_{t \geq 0}$ is the sequence defined in (7); and σ and κ are constants defined in Assumption 1; and f^* is the minimum value of problem (1).

Remark 4. By comparing Theorem 1 and Theorem 2, we conclude that the order of convergence rate for both options in Algorithm 1 (with slightly different choices of learning rate γ) have the same dependence on γ, N and I . It is then immediate that Algorithm 1 with Option II can achieve the linear speedup with the same (reduced) communication complexity as summarized in Corollary 2 for Option I.

3. Extension: Momentum SGD with Decentralized Communication

Note that Algorithm 1 requires to compute global averages of all N nodes at each communication step (6). Such global averages are possible without a central parameter center by using distributive average protocols such as all-reduce primitives (Goyal et al., 2017). However, the feasibility and performance of exiting all-reduce protocols are restricted to the underlying network topology. For example, the performance of a ring based all-reduce protocol can be poor in a line network, which does not have a physical ring. In this subsection, we consider distributed momentum SGD with decentralized communication (described in Algorithm 2) where the communication pattern can be fully customized. Let $\mathbf{W} \in \mathbb{R}^{N \times N}$ be a symmetric doubly stochastic matrix. By definition, a symmetric doubly stochastic matrix satisfies (1) $W_{ij} \in [0, 1], \forall i, j \in \{1, 2, \dots, N\}$; (2) $\mathbf{W}^T = \mathbf{W}$; (3) $\sum_{j=1}^N W_{ij} = 1, \forall i \in \{1, 2, \dots, N\}$. For a computer network with N nodes, we can use \mathbf{W} to encode the communication between nodes, i.e., if node i and node j are disconnected, we let $W_{ij} = 0$. This section further assumes the mixing matrix \mathbf{W} that obeys the network connection topology is selected to satisfy the following assumption. The same assumption is used in (Lian et al., 2017; Jiang et al., 2017; Wang & Joshi, 2018a).

Algorithm 2 Distributed Momentum SGD with Decentralized Communication

- 1: **Input:** Initialize local solutions $\mathbf{x}_i^{(0)} = \hat{\mathbf{x}} \in \mathbb{R}^m$ and momentum buffers $\mathbf{u}_i^{(0)} = \mathbf{0}, \forall i \in \{1, 2, \dots, N\}$. Set mixing matrix \mathbf{W} , learning rate $\gamma > 0$, momentum coefficient $\beta \in [0, 1)$ and number of iterations T
- 2: **for** $t = 1, 2, \dots, T - 1$ **do**
- 3: Each worker i samples its stochastic gradient $\mathbf{g}_i^{(t-1)} = \nabla F_i(\mathbf{x}_i^{(t-1)}; \xi_i^{(t-1)})$ with $\xi_i^{(t-1)} \sim \mathcal{D}_i$.
- 4: Each worker i in parallel updates its local auxiliary momentum buffer $\tilde{\mathbf{u}}_i^{(t)}$ and solution $\tilde{\mathbf{x}}_i^{(t)}$ via

$$\text{Option I: } \begin{cases} \tilde{\mathbf{u}}_i^{(t)} = \beta \mathbf{u}_i^{(t-1)} + \mathbf{g}_i^{(t-1)} \\ \tilde{\mathbf{x}}_i^{(t)} = \mathbf{x}_i^{(t-1)} - \gamma \tilde{\mathbf{u}}_i^{(t)} \end{cases} \quad \forall i. \quad (16)$$

Or

$$\text{Option II: } \begin{cases} \tilde{\mathbf{u}}_i^{(t)} = \beta \mathbf{u}_i^{(t-1)} + \mathbf{g}_i^{(t-1)} \\ \tilde{\mathbf{v}}_i^{(t)} = \beta \tilde{\mathbf{u}}_i^{(t)} + \mathbf{g}_i^{(t-1)} \\ \tilde{\mathbf{x}}_i^{(t)} = \mathbf{x}_i^{(t-1)} - \gamma \tilde{\mathbf{v}}_i^{(t)} \end{cases} \quad \forall i. \quad (17)$$

- 5: Each worker i updates its local momentum buffer $\mathbf{u}_i^{(t)}$ and solution $\mathbf{x}_i^{(t)}$ based on the neighborhood weighted averages given by

$$\begin{cases} \mathbf{u}_i^{(t)} = \sum_{j=1}^N \tilde{\mathbf{u}}_j^{(t)} W_{ji} \\ \mathbf{x}_i^{(t)} = \sum_{j=1}^N \tilde{\mathbf{x}}_j^{(t)} W_{ji} \end{cases} \quad \forall i \quad (18)$$

6: **end for**

Assumption 2. The mixing matrix $\mathbf{W} \in \mathbb{R}^{N \times N}$ is a symmetric doubly stochastic matrix satisfying $\lambda_1(\mathbf{W}) = 1$ and

$$\max\{|\lambda_2(\mathbf{W})|, |\lambda_N(\mathbf{W})|\} \leq \sqrt{\rho} < 1 \quad (15)$$

where $\lambda_i(\mathbf{W})$ is the i -th largest eigenvalue of \mathbf{W} .

Compared with PR-SGD-Momentum described in Algorithm 1, Algorithm 2 uses a **local** aggregation step (18) at every iteration. While Algorithm 2 can not skip the aggregation steps as Algorithm 1 does and hence involves I times more communication rounds, the per-step communication is more flexible since each node only computes neighborhood averages rather than global averages. As a consequence, Algorithm 2 is more suitable to distributed machine learning with mobiles where the network topology is heterogeneous and diverse. By extending our analysis for Algorithm 1, the next theorem proves the linear speedup of Algorithm 2.

Theorem 3. Consider problem (1) under Assumptions 1 and 2. If we choose $\gamma \leq \min\{\frac{(1-\beta)^2(1-\sqrt{\rho})^2}{6L}, \frac{(1-\beta)(1-\sqrt{\rho})}{4L}\}$ in Algorithm 2,

then for all $T \geq 1$, we have

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2] = O\left(\frac{1}{\gamma T}\right) + O\left(\frac{\gamma}{N}\sigma^2\right) + O(\gamma^2\sigma^2) + O(\gamma^2\kappa^2)$$

where $\{\bar{\mathbf{x}}^{(t)}\}_{t \geq 0}$ defined in (7) are the averages of local solutions across N workers; and σ, κ and ρ are constants defined in Assumptions 1-2.

Proof. See Supplement 6.6. \square

This next corollary further summarizes that using $\gamma = O(\frac{\sqrt{N}}{\sqrt{T}})$ in Algorithm 2 can achieve $O(\frac{1}{\sqrt{NT}})$ convergence with the linear speedup.

Corollary 3. Consider problem (1) under Assumptions 1 and 2. If we choose $\gamma = \frac{\sqrt{N}}{\sqrt{T}}$ in Algorithm 2, then for all $T \geq \max\{\frac{36NL^2}{(1-\beta)^4(1-\sqrt{\rho})^4}, \frac{32NL^2}{(1-\beta)^2(1-\sqrt{\rho})^2}\}$, we have

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2] = O\left(\frac{1}{\sqrt{NT}}\right) + O\left(\frac{N}{T}\right)$$


where $\{\bar{\mathbf{x}}^{(t)}\}_{t \geq 0}$ defined in (7) are the averages of local solutions across N workers.

Note that Algorithm 2 with $\beta = 0$ degrades to the decentralized SGD without momentum considered in (Lian et al., 2017), where the linear speedup of decentralized SGD is first shown. Recently, many variants of decentralized SGD without momentum with linear speedup have been developed for distributed deep learning (Tang et al., 2018; Wang & Joshi, 2018a; Assran et al., 2018). To our knowledge, this is the first time that decentralized momentum SGD is shown to possess the same linear speedup.

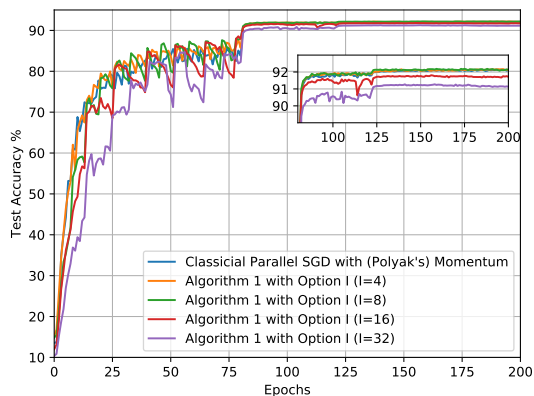
4. Experiments

In this section, we validate our theory with experiments on training ResNet (He et al., 2016) for the image classification tasks over CIFAR-10 (Krizhevsky & Hinton, 2009) and ImageNet (Deng et al., 2009). Our experiments are conducted on a machine with 8 NVIDIA P100 GPUs. The local batch size at each GPU is 64. The learning rate is initialized to 0.3 and is divided by 10 when all GPUs jointly access 80 and 120 epochs.⁵ The momentum coefficient, i.e., β in Algorithm 1, is set to 0.9 for both Polyak's and Nesterov's momentum. All algorithms are implemented using PyTorch 0.4. To study how communication skipping affects the convergence of Algorithm 1, we run Algorithm

⁵Such decaying learning rates are used to achieve good test accuracy by practitioners (He et al., 2016). This deviates from the constant learning rates used in the theory to establish the linear speedup. On one hand, it is possible follow the analysis techniques in (Yu et al., 2018) to extend this paper's theory to establish a similar linear speedup. On the other hand, our supplement 6.7.1 reports extra experiments to verify the linear speedup of Algorithm 1 using constant learning rates faithful to the theory.




(a) Training loss v.s. epochs.




(b) Test accuracy v.s. epochs.

Figure 2: Algorithm 1 with Option I: convergence v.s. epochs for ResNet56 over CIFAR10.



(a) Training loss v.s. wall clock time.



(b) Test accuracy v.s. wall clock time.

Figure 3: Algorithm 1 with Option I: convergence v.s. wall clock time for ResNet56 over CIFAR10.

1 with $I \in \{4, 8, 16, 32\}$ and compare the test accuracy convergence between Algorithm 1 and the classical parallel mini-batch SGD with momentum, which uses an inter-node communication step at every iteration and can be interpreted as Algorithm 1 with $I = 1$. Figure 2 plots the convergence of training loss and test accuracy in terms of the number of epochs that are jointly accessed by all used GPUs. That is, if the x -axis value is 8, then each GPU access 1 epoch of training data. The same convention is followed by other figures for multiple GPU training in this paper. By plotting the convergence in terms of the number of epochs that are jointly accessed by all GPUs, we can verify the $O(\frac{1}{\sqrt{NT}})$ convergence for Algorithm 1 proven in this paper. To verify the benefit of skipping communication in our Algorithm 1, Figure 3 plots the convergence of training loss and test accuracy in terms of the wall clock time. Since Algorithm

1 with $I > 0$ skip communication steps, it is much faster than the classical parallel momentum SGD when measured by the wall clock time. More numerical experiments on training ResNet over ImageNet, comparisons with a model averaging strategy suggested in (Seide & Agarwal, 2016; Wang & Joshi, 2018b), and Algorithm 1 with Nesterov’s momentum are available in Supplement 6.7.

5. Conclusion

This paper considers parallel restarted SGD with momentum and prove that it can achieve $O(1/\sqrt{NT})$ convergence with $O(N^{3/2}T^{1/2})$ or $O(N^{3/4}T^{3/4})$ communication rounds depending whether each node accesses identical objective functions $f_i(x)$ or not. We further show that distributed momentum SGD with decentralized communication can achieve $O(1/\sqrt{NT})$ convergence.

References

- Aji, A. F. and Heafield, K. Sparse communication for distributed gradient descent. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2017.
- Alistarh, D., Grubic, D., Li, J., Tomioka, R., and Vojnovic, M. QSGD: Communication-efficient SGD via gradient quantization and encoding. In *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- Assran, M., Loizou, N., Ballas, N., and Rabbat, M. Stochastic gradient push for distributed deep learning. *arXiv:1811.10792*, 2018.
- Chen, K. and Huo, Q. Scalable training of deep learning machines by incremental block training with intra-block parallel optimization and blockwise model-update filtering. In *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016.
- Dean, J., Corrado, G., Monga, R., Chen, K., Devin, M., Mao, M., Senior, A., Tucker, P., Yang, K., Le, Q. V., et al. Large scale distributed deep networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2012.
- Dekel, O., Gilad-Bachrach, R., Shamir, O., and Xiao, L. Optimal distributed online prediction using mini-batches. *Journal of Machine Learning Research*, 13(165–202), 2012.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. ImageNet: A large-scale hierarchical image database. In *IEEE conference on computer vision and pattern recognition (CVPR)*, 2009.
- Ghadimi, E., Feyzmahdavian, H. R., and Johansson, M. Global convergence of the heavy-ball method for convex optimization. *arXiv:1412.7457*, 2014.
- Ghadimi, S. and Lan, G. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- Goyal, P., Dollár, P., Girshick, R., Noordhuis, P., Wesolowski, L., Kyrola, A., Tulloch, A., Jia, Y., and He, K. Accurate, large minibatch SGD: training imagenet in 1 hour. *arXiv:1706.02677*, 2017.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *IEEE conference on computer vision and pattern recognition (CVPR)*, 2016.
- Horn, R. A. and Johnson, C. R. *Matrix Analysis*. Cambridge University Press, 1985.
- Jiang, P. and Agrawal, G. A linear speedup analysis of distributed deep learning with sparse and quantized communication. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- Jiang, Z., Balu, A., Hegde, C., and Sarkar, S. Collaborative deep learning in fixed topology networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- Krizhevsky, A. and Hinton, G. Learning multiple layers of features from tiny images. *Technical report, University of Toronto*, 2009.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2012.
- Li, M., Andersen, D. G., Smola, A. J., and Yu, K. Communication efficient distributed machine learning with the parameter server. In *Advances in Neural Information Processing Systems (NIPS)*, 2014.
- Lian, X., Huang, Y., Li, Y., and Liu, J. Asynchronous parallel stochastic gradient for nonconvex optimization. In *Advances in Neural Information Processing Systems (NIPS)*, 2015.
- Lian, X., Zhang, C., Zhang, H., Hsieh, C.-J., Zhang, W., and Liu, J. Can decentralized algorithms outperform centralized algorithms? A case study for decentralized parallel stochastic gradient descent. In *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- Lin, T., Stich, S. U., and Jaggi, M. Don’t use large mini-batches, use local SGD. *arXiv:1808.07217*, 2018.
- Mania, H., Pan, X., Papailiopoulos, D., Recht, B., Ramchandran, K., and Jordan, M. I. Perturbed iterate analysis for asynchronous stochastic optimization. *SIAM Journal on Optimization*, 27(4):2202–2229, 2017.
- McMahan, H. B., Moore, E., Ramage, D., Hampson, S., et al. Communication-efficient learning of deep networks from decentralized data. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2017.
- Nesterov, Y. *Introductory Lectures on Convex Optimization: A Basic Course*. Springer Science & Business Media, 2004.
- Polyak, B. T. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5):1–17, 1964.
- Seide, F. and Agarwal, A. CNTK: Microsoft’s open-source deep-learning toolkit. In *Proceedings of ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2016.

- Seide, F., Fu, H., Droppo, J., Li, G., and Yu, D. 1-bit stochastic gradient descent and its application to data-parallel distributed training of speech DNNs. In *Annual Conference of the International Speech Communication Association (INTERSPEECH)*, 2014.
- Stich, S. U. Local SGD converges fast and communicates little. *arXiv:1805.09767*, 2018.
- Strom, N. Scalable distributed DNN training using commodity GPU cloud computing. In *Annual Conference of the International Speech Communication Association (INTERSPEECH)*, 2015.
- Su, H., Chen, H., and Xu, H. Experiments on parallel training of deep neural network using model averaging. *arXiv:1507.01239v3*, 2018.
- Sutskever, I., Martens, J., Dahl, G., and Hinton, G. On the importance of initialization and momentum in deep learning. In *International conference on machine learning (ICML)*, 2013.
- Tang, H., Lian, X., Yan, M., Zhang, C., and Liu, J. D^2 : Decentralized training over decentralized data. *arXiv:1803.07068*, 2018.
- Wang, J. and Joshi, G. Cooperative SGD: A unified framework for the design and analysis of communication-efficient SGD algorithms. *arXiv:1808.07576*, 2018a.
- Wang, J. and Joshi, G. Adaptive communication strategies to achieve the best error-runtime trade-off in local-update SGD. *arXiv:1810.08313*, 2018b.
- Wen, W., Xu, C., Yan, F., Wu, C., Wang, Y., Chen, Y., and Li, H. Terngrad: Ternary gradients to reduce communication in distributed deep learning. In *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- Yan, Y., Yang, T., Li, Z., Lin, Q., and Yang, Y. A unified analysis of stochastic momentum methods for deep learning. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2018.
- Yu, H., Yang, S., and Zhu, S. Parallel restarted SGD with faster convergence and less communication: Demystifying why model averaging works for deep learning. *arXiv:1807.06629*, 2018.
- Zavriev, S. and Kostyuk, F. Heavy-ball method in nonconvex optimization problems. *Computational Mathematics and Modeling*, 4(4):336–341, 1993.
- Zhou, F. and Cong, G. On the convergence properties of a K -step averaging stochastic gradient descent algorithm for nonconvex optimization. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2018.

6. Supplement

6.1. Proof of Lemma 1

The lemma follows simply because unbiased stochastic gradients \mathbf{g}_i are independently sampled. The formal proof is as follows:

$$\begin{aligned}
 & \mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i\|^2] \\
 & \stackrel{(a)}{=} \mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N (\mathbf{g}_i - \nabla f_i(\mathbf{x}_i))\|^2] + \mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i)\|^2] \\
 & \stackrel{(b)}{=} \frac{1}{N^2} \sum_{i=1}^N \mathbb{E}[\|\mathbf{g}_i - \nabla f_i(\mathbf{x}_i)\|^2] + \mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i)\|^2] \\
 & \stackrel{(c)}{\leq} \frac{1}{N} \sigma^2 + \mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i)\|^2]
 \end{aligned}$$

where (a) follows from the facts that $\mathbb{E}[\mathbf{g}_i] = \nabla f_i(\mathbf{x}_i), \forall i$ and identity $\mathbb{E}[\|\mathbf{Z}\|^2] = \mathbb{E}[\|\mathbf{Z} - \mathbb{E}[\mathbf{Z}]\|^2] + \|\mathbb{E}[\mathbf{Z}]\|^2$ holds for any random vector \mathbf{Z} ; (b) holds from the facts that $\mathbf{g}_i - \nabla f_i(\mathbf{x}_i)$ are independent random vectors with $\mathbf{0}$ means and identity $\mathbb{E}[\|\sum_{i=1}^N \mathbf{Z}_i\|^2] = \sum_{i=1}^N \mathbb{E}[\|\mathbf{Z}_i\|^2]$ holds when \mathbf{Z}_i are independent with $\mathbf{0}$ means; and (c) follows from Assumption 1.

6.2. Proof of Lemma 2

This lemma follows from simple algebraic manipulations as follows:

$$\begin{aligned}
 & \frac{1}{N} \sum_{i=1}^N \|\nabla f_i(\mathbf{x}_i) - \frac{1}{N} \sum_{j=1}^N \nabla f_j(\mathbf{x}_j)\|^2 \\
 & = \frac{1}{N} \sum_{i=1}^N \|\nabla f_i(\mathbf{x}_i) - \nabla f_i(\bar{\mathbf{x}}) + \nabla f_i(\bar{\mathbf{x}}) - \nabla f(\bar{\mathbf{x}}) + \nabla f(\bar{\mathbf{x}}) - \frac{1}{N} \sum_{j=1}^N \nabla f_j(\mathbf{x}_j)\|^2 \\
 & \stackrel{(a)}{\leq} \frac{1}{N} \sum_{i=1}^N \left[3\|\nabla f_i(\mathbf{x}_i) - \nabla f_i(\bar{\mathbf{x}})\|^2 + 3\|\nabla f_i(\bar{\mathbf{x}}) - \nabla f(\bar{\mathbf{x}})\|^2 + 3\|\nabla f(\bar{\mathbf{x}}) - \frac{1}{N} \sum_{j=1}^N \nabla f_j(\mathbf{x}_j)\|^2 \right] \\
 & \stackrel{(b)}{\leq} \frac{1}{N} \sum_{i=1}^N \left[3L^2\|\mathbf{x}_i - \bar{\mathbf{x}}\|^2 + 3\|\nabla f_i(\bar{\mathbf{x}}) - \nabla f(\bar{\mathbf{x}})\|^2 + 3\frac{1}{N} \sum_{j=1}^N L^2\|\bar{\mathbf{x}} - \mathbf{x}_j\|^2 \right] \\
 & = \frac{6L^2}{N} \sum_{i=1}^N \|\mathbf{x}_i - \bar{\mathbf{x}}\|^2 + \frac{3}{N} \sum_{i=1}^N \|\nabla f_i(\bar{\mathbf{x}}) - \nabla f(\bar{\mathbf{x}})\|^2
 \end{aligned}$$

where (a) follows from the basic inequality $\|\mathbf{a}_1 + \mathbf{a}_2 + \mathbf{a}_3\|^2 \leq 3\|\mathbf{a}_1\|^2 + 3\|\mathbf{a}_2\|^2 + 3\|\mathbf{a}_3\|^2$ for any vectors $\mathbf{a}_1, \mathbf{a}_2$ and \mathbf{a}_3 ; (b) follows because $\|\nabla f_i(\mathbf{x}_i) - \nabla f_i(\bar{\mathbf{x}})\| \leq L\|\mathbf{x}_i - \bar{\mathbf{x}}\|$ by the smoothness of $f_i(\cdot)$ in Assumption 1 and $\|\nabla f(\bar{\mathbf{x}}) - \frac{1}{N} \sum_{j=1}^N \nabla f_j(\mathbf{x}_j)\|^2 = \|\frac{1}{N} \sum_{j=1}^N \nabla f_j(\bar{\mathbf{x}}) - \frac{1}{N} \sum_{j=1}^N \nabla f_j(\mathbf{x}_j)\|^2 \leq \frac{1}{N} \sum_{j=1}^N \|\nabla f_j(\bar{\mathbf{x}}) - \nabla f_j(\mathbf{x}_j)\|^2 \leq \frac{1}{N} \sum_{j=1}^N L^2\|\bar{\mathbf{x}} - \mathbf{x}_j\|^2$, where the first inequality follows from the convexity of $\|\cdot\|^2$ and Jensen's inequality and the second inequality follows from the smoothness of each $f_j(\cdot)$.

6.3. Proof of Theorem 1

To prove this theorem, we first introduce an auxiliary sequence $\{\bar{\mathbf{z}}^{(t)}\}_{t \geq 0}$ given by

$$\bar{\mathbf{z}}^{(t)} \triangleq \begin{cases} \bar{\mathbf{x}}^{(t)}, & t = 0 \\ \frac{1}{1-\beta} \bar{\mathbf{x}}^{(t)} - \frac{\beta}{1-\beta} \bar{\mathbf{x}}^{(t-1)}, & t \geq 1 \end{cases} \quad (19)$$

where $\bar{\mathbf{x}}^{(t)}$ defined in (7) are the node averages of local solutions from Algorithm 1 with Option I. A similar auxiliary sequence $\{\bar{\mathbf{z}}^{(t)}\}_{t \geq 0}$ has been used for the convergence analysis of standard momentum methods (single node without restarting) in (Ghadimi et al., 2014) and (Yan et al., 2018).

6.3.1. USEFUL LEMMAS

Before the main proof of Theorem 1, we further introduce several useful lemmas.

Lemma 3. Consider the sequence $\{\bar{\mathbf{z}}^{(t)}\}_{t \geq 0}$ defined in (19). Algorithm 1 with Option I ensures that for all $t \geq 0$, we have

$$\bar{\mathbf{z}}^{(t+1)} - \bar{\mathbf{z}}^{(t)} = -\frac{\gamma}{1-\beta} \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)}. \quad (20)$$

Proof. To prove $\bar{\mathbf{z}}^{(t+1)} - \bar{\mathbf{z}}^{(t)} = -\frac{\gamma}{1-\beta} \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)}$, we consider $t = 0$ and $t \geq 1$ separately.

- **Case $t = 0$:** We have

$$\begin{aligned} \bar{\mathbf{z}}^{(t+1)} - \bar{\mathbf{z}}^{(t)} &= \bar{\mathbf{z}}^{(1)} - \bar{\mathbf{z}}^{(0)} \\ &\stackrel{(a)}{=} \frac{1}{1-\beta} \bar{\mathbf{x}}^{(1)} - \frac{\beta}{1-\beta} \bar{\mathbf{x}}^{(0)} - \bar{\mathbf{x}}^{(0)} \\ &= \frac{1}{1-\beta} [\bar{\mathbf{x}}^{(1)} - \bar{\mathbf{x}}^{(0)}] \\ &\stackrel{(b)}{=} -\frac{\gamma}{1-\beta} \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(0)} \end{aligned}$$

where (a) follows from (19); (b) follows from (10) by noting that $\bar{\mathbf{u}}^{(0)} = \mathbf{0}$.

- **Case $t \geq 1$:** We have

$$\begin{aligned} \bar{\mathbf{z}}^{(t+1)} - \bar{\mathbf{z}}^{(t)} &\stackrel{(a)}{=} \frac{1}{1-\beta} [\bar{\mathbf{x}}^{(t+1)} - \bar{\mathbf{x}}^{(t)}] - \frac{\beta}{1-\beta} [\bar{\mathbf{x}}^{(t)} - \bar{\mathbf{x}}^{(t-1)}] \\ &\stackrel{(b)}{=} -\frac{\gamma}{1-\beta} \bar{\mathbf{u}}^{(t+1)} + \frac{\gamma\beta}{1-\beta} \bar{\mathbf{u}}^{(t)} \\ &\stackrel{(c)}{=} -\frac{\gamma}{1-\beta} [\beta \bar{\mathbf{u}}^{(t)} + \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)}] + \frac{\gamma\beta}{1-\beta} \bar{\mathbf{u}}^{(t)} \\ &= -\frac{\gamma}{1-\beta} \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)} \end{aligned}$$

where (a) follows from (19); and (b) and (c) follow from (10).

□

Lemma 4. Let $\{\bar{\mathbf{x}}^{(t)}\}_{t \geq 0}$ defined in (7) be the node averages of local solutions from Algorithm 1 with Option I. Let $\{\bar{\mathbf{z}}^{(t)}\}_{t \geq 0}$ be defined in (19). For all $T \geq 1$, Algorithm 1 with Option I ensures that

$$\sum_{t=0}^{T-1} \|\bar{\mathbf{z}}^{(t)} - \bar{\mathbf{x}}^{(t)}\|^2 \leq \frac{\gamma^2 \beta^2}{(1-\beta)^4} \sum_{t=0}^{T-1} \left\| \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)} \right] \right\|^2 \quad (21)$$

Proof. Recall that $\bar{\mathbf{u}}^{(0)} = \mathbf{0}$. Recursively applying the first equation in (10) for t times yields:

$$\bar{\mathbf{u}}^{(t)} = \sum_{\tau=0}^{t-1} \beta^{t-1-\tau} \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(\tau)} \right], \quad \forall t \geq 1 \quad (22)$$

Note that for all $t \geq 1$, we have

$$\bar{\mathbf{z}}^{(t)} - \bar{\mathbf{x}}^{(t)} \stackrel{(a)}{=} \frac{\beta}{1-\beta} [\bar{\mathbf{x}}^{(t)} - \bar{\mathbf{x}}^{(t-1)}] \stackrel{(b)}{=} -\frac{\beta\gamma}{1-\beta} \bar{\mathbf{u}}^{(t)} \quad (23)$$

where (a) follows from (19) and (b) follows from the second equation in (10).

Combining (22) and (23) yields

$$\bar{\mathbf{z}}^{(t)} - \bar{\mathbf{x}}^{(t)} = -\frac{\beta\gamma}{1-\beta} \sum_{\tau=0}^{t-1} \beta^{t-1-\tau} \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(\tau)} \right], \quad \forall t \geq 1 \quad (24)$$

Define $s_t \triangleq \sum_{\tau=0}^{t-1} \beta^{t-1-\tau} = \frac{1-\beta^t}{1-\beta}$. For all $t \geq 1$, we have

$$\begin{aligned} \|\bar{\mathbf{z}}^{(t)} - \bar{\mathbf{x}}^{(t)}\|^2 &= \frac{\gamma^2 \beta^2}{(1-\beta)^2} s_t^2 \left\| \sum_{\tau=0}^{t-1} \frac{\beta^{t-1-\tau}}{s_t} \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(\tau)} \right] \right\|^2 \\ &\stackrel{(a)}{\leq} \frac{\gamma^2 \beta^2}{(1-\beta)^2} s_t^2 \sum_{\tau=0}^{t-1} \frac{\beta^{t-1-\tau}}{s_t} \left\| \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(\tau)} \right] \right\|^2 \\ &= \frac{\gamma^2 \beta^2 (1-\beta^t)}{(1-\beta)^3} \sum_{\tau=0}^{t-1} \beta^{t-1-\tau} \left\| \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(\tau)} \right] \right\|^2 \\ &\leq \frac{\gamma^2 \beta^2}{(1-\beta)^3} \sum_{\tau=0}^{t-1} \beta^{t-1-\tau} \left\| \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(\tau)} \right] \right\|^2 \end{aligned} \quad (25)$$

where (a) follows from the convexity of $\|\cdot\|^2$ and Jensen's inequality.

Fix $T \geq 1$. Note that $\bar{\mathbf{z}}^{(0)} - \bar{\mathbf{x}}^{(0)} = \mathbf{0}$. Summing (25) over $t \in \{1, 2, \dots, T-1\}$ yields

$$\begin{aligned} \sum_{t=0}^{T-1} \|\bar{\mathbf{z}}^{(t)} - \bar{\mathbf{x}}^{(t)}\|^2 &\leq \frac{\gamma^2 \beta^2}{(1-\beta)^3} \sum_{t=1}^{T-1} \sum_{\tau=0}^{t-1} \beta^{t-1-\tau} \left\| \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(\tau)} \right] \right\|^2 \\ &= \frac{\gamma^2 \beta^2}{(1-\beta)^3} \sum_{\tau=0}^{T-2} \left(\left\| \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(\tau)} \right] \right\|^2 \sum_{l=\tau+1}^{T-1} \beta^{l-1-\tau} \right) \\ &\stackrel{(a)}{\leq} \frac{\gamma^2 \beta^2}{(1-\beta)^4} \sum_{\tau=0}^{T-2} \left\| \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(\tau)} \right] \right\|^2 \\ &\leq \frac{\gamma^2 \beta^2}{(1-\beta)^4} \sum_{\tau=0}^{T-1} \left\| \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(\tau)} \right] \right\|^2 \end{aligned} \quad (26)$$

where (a) follows by noting that $\sum_{l=\tau+1}^{T-1} \beta^{l-1-\tau} = \frac{1-\beta^{T-\tau-1}}{1-\beta} \leq \frac{1}{1-\beta}$. □

Lemma 5. Consider problem (1) under Assumption 1. Let $\{\bar{\mathbf{x}}^{(t)}\}_{t \geq 0}$ defined in (7) be the node averages of local solutions from Algorithm 1 with Option I. If γ and I are chosen to satisfy $\frac{12L^2\gamma^2 I^2}{(1-\beta)^2} < 1$, then for all $T \geq 1$, we have

$$\sum_{t=0}^{T-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] \leq \frac{1}{1 - \frac{12L^2\gamma^2 I^2}{(1-\beta)^2}} \frac{2\gamma^2 I \sigma^2}{(1-\beta)^2} T + \frac{1}{1 - \frac{12L^2\gamma^2 I^2}{(1-\beta)^2}} \frac{6\gamma^2 I^2 \kappa^2}{(1-\beta)^2} T$$

Proof. Since $\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)} = \mathbf{0}$ when t is a multiple of I , our focus is to develop upper bounds for $\frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2]$ when t is not a multiple of I . Consider $t \geq 1$ that is not a multiple of I . Note that Algorithm 1 restarts ‘‘SGD with momentum’’ every I iterations (by resetting $\mathbf{u}_i^{(t)} = \hat{\mathbf{u}} \triangleq \frac{1}{N} \sum_{j=1}^N \mathbf{u}_j^{(t)}$ and $\mathbf{x}_i^{(t)} = \hat{\mathbf{x}} \triangleq \frac{1}{N} \sum_{j=1}^N \mathbf{x}_j^{(t)}$). Let $t_0 < t$ be the

largest iteration index that is a multiple of I , i.e., $t_0 \bmod I = 0$. Note that we must have $t - t_0 < I$ and $\mathbf{u}_i^{(t_0)} = \hat{\mathbf{u}}$, $\mathbf{x}_i^{(t_0)} = \hat{\mathbf{x}}$. For any $\tau \in \{t_0 + 1, \dots, t\}$, by recursively applying the first equation in (4), we have

$$\mathbf{u}_i^{(\tau)} = \beta^{\tau-t_0} \hat{\mathbf{u}} + \sum_{k=t_0}^{\tau-1} \beta^{\tau-1-k} \mathbf{g}_i^{(k)} \quad (27)$$

For any $\tau \in \{t_0 + 1, \dots, t\}$, by the second equation in (4), we have

$$\mathbf{x}_i^{(\tau)} - \mathbf{x}_i^{(\tau-1)} = -\gamma \mathbf{u}_i^{(\tau)} \quad (28)$$

Summing over $\tau \in \{t_0 + 1, \dots, t\}$ and noting that $\mathbf{x}_i^{(t_0)} = \hat{\mathbf{x}}$ yields

$$\begin{aligned} \mathbf{x}_i^{(t)} &= \hat{\mathbf{x}} - \gamma \sum_{\tau=t_0+1}^t \mathbf{u}_i^{(\tau)} \\ &\stackrel{(a)}{=} \hat{\mathbf{x}} - \gamma \left(\sum_{\tau=t_0+1}^t \beta^{\tau-t_0} \right) \hat{\mathbf{u}} - \gamma \sum_{\tau=t_0+1}^t \sum_{k=t_0}^{\tau-1} \beta^{\tau-1-k} \mathbf{g}_i^{(k)} \\ &= \hat{\mathbf{x}} - \gamma \left(\sum_{j=1}^{t-t_0} \beta^j \right) \hat{\mathbf{u}} - \gamma \sum_{k=t_0}^{t-1} \sum_{j=0}^{t-1-k} \beta^j \mathbf{g}_i^{(k)} \\ &= \hat{\mathbf{x}} - \gamma \left(\sum_{j=1}^{t-t_0} \beta^j \right) \hat{\mathbf{u}} - \gamma \sum_{\tau=t_0}^{t-1} \frac{1 - \beta^{t-\tau}}{1 - \beta} \mathbf{g}_i^{(\tau)} \end{aligned} \quad (29)$$

where (a) follows by substituting (27).

Using an argument similar to the above, we can further show

$$\bar{\mathbf{x}}^{(t)} = \hat{\mathbf{x}} - \gamma \left(\sum_{j=1}^{t-t_0} \beta^j \right) \hat{\mathbf{u}} - \gamma \sum_{\tau=t_0}^{t-1} \frac{1 - \beta^{t-\tau}}{1 - \beta} \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(\tau)} \quad (30)$$

Combining (29) and (30) yields

$$\begin{aligned} &\frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] \\ &= \gamma^2 \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\left\| \sum_{\tau=t_0}^{t-1} [\mathbf{g}_i^{(\tau)} - \frac{1}{N} \sum_{j=1}^N \mathbf{g}_j^{(\tau)}] \frac{1 - \beta^{t-\tau}}{1 - \beta} \right\|^2 \right] \\ &\stackrel{(a)}{\leq} 2\gamma^2 \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\left\| \sum_{\tau=t_0}^{t-1} [[\mathbf{g}_i^{(\tau)} - \nabla f_i(\mathbf{x}_i^{(\tau)})] - \frac{1}{N} \sum_{j=1}^N [\mathbf{g}_j^{(\tau)} - \nabla f_j(\mathbf{x}_j^{(\tau)})]] \frac{1 - \beta^{t-\tau}}{1 - \beta} \right\|^2 \right] \\ &\quad + 2\gamma^2 \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\left\| \sum_{\tau=t_0}^{t-1} [\nabla f_i(\mathbf{x}_i^{(\tau)}) - \frac{1}{N} \sum_{j=1}^N \nabla f_j(\mathbf{x}_j^{(\tau)})] \frac{1 - \beta^{t-\tau}}{1 - \beta} \right\|^2 \right] \end{aligned} \quad (31)$$

where (a) follows from the basic inequality $\|\mathbf{a}_1 + \mathbf{a}_2\|^2 \leq 2\|\mathbf{a}_1\|^2 + 2\|\mathbf{a}_2\|^2$.

Now we develop the respective upper bounds of the two terms on the right side of (31). We note that

$$\begin{aligned}
 & \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\left\| \sum_{\tau=t_0}^{t-1} [\mathbf{g}_i^{(\tau)} - \nabla f_i(\mathbf{x}_i^{(\tau)})] - \frac{1}{N} \sum_{j=1}^N [\mathbf{g}_j^{(\tau)} - \nabla f_j(\mathbf{x}_j^{(\tau)})] \right\| \frac{1-\beta^{t-\tau}}{1-\beta} \right]^2 \\
 & \stackrel{(a)}{\leq} \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\left\| \sum_{\tau=t_0}^{t-1} [\mathbf{g}_i^{(\tau)} - \nabla f_i(\mathbf{x}_i^{(\tau)})] \right\| \frac{1-\beta^{t-\tau}}{1-\beta} \right]^2 \\
 & \stackrel{(b)}{=} \frac{1}{N} \sum_{i=1}^N \sum_{\tau=t_0}^{t-1} \mathbb{E} \left[\left\| [\mathbf{g}_i^{(\tau)} - \nabla f_i(\mathbf{x}_i^{(\tau)})] \right\| \frac{1-\beta^{t-\tau}}{1-\beta} \right]^2 \\
 & \stackrel{(c)}{\leq} \frac{I\sigma^2}{(1-\beta)^2}
 \end{aligned} \tag{32}$$

where (a) follows from the inequality $\frac{1}{N} \sum_{i=1}^N \|\mathbf{a}_i - [\frac{1}{N} \sum_{j=1}^N \mathbf{a}_j]\|^2 = \frac{1}{N} \sum_{i=1}^N \|\mathbf{a}_i\|^2 - \|\frac{1}{N} \sum_{i=1}^N \mathbf{a}_i\|^2 \leq \frac{1}{N} \sum_{i=1}^N \|\mathbf{a}_i\|^2$ with $\mathbf{a}_i = \sum_{\tau=t_0}^{t-1} [\mathbf{g}_i^{(\tau)} - \nabla f_i(\mathbf{x}_i^{(\tau)})] \frac{1-\beta^{t-\tau}}{1-\beta}$; (b) follows because $\mathbb{E}[\mathbf{g}_i^{(\tau_2)} - \nabla f_i(\mathbf{x}_i^{(\tau_2)}) | \mathbf{g}_i^{(\tau_1)} - \nabla f_i(\mathbf{x}_i^{(\tau_1)})] = \mathbf{0}$ for any $\tau_2 > \tau_1$; (c) follows because $|\frac{1-\beta^{t-\tau}}{1-\beta}| \leq \frac{1}{1-\beta}$, $t-t_0 < I$ and $\mathbb{E}[\|\mathbf{g}_i^{(\tau)} - \nabla f_i(\mathbf{x}_i^{(\tau)})\|^2] \leq \sigma^2$ (by Assumption 1).

We further note that

$$\begin{aligned}
 & \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\left\| \sum_{\tau=t_0}^{t-1} [\nabla f_i(\mathbf{x}_i^{(\tau)}) - \frac{1}{N} \sum_{j=1}^N \nabla f_j(\mathbf{x}_j^{(\tau)})] \right\| \frac{1-\beta^{t-\tau}}{1-\beta} \right]^2 \\
 & \stackrel{(a)}{\leq} \frac{1}{N} \sum_{i=1}^N (t-t_0) \frac{1}{(1-\beta)^2} \sum_{\tau=t_0}^{t-1} \mathbb{E} \left[\left\| \nabla f_i(\mathbf{x}_i^{(\tau)}) - \frac{1}{N} \sum_{j=1}^N \nabla f_j(\mathbf{x}_j^{(\tau)}) \right\|^2 \right] \\
 & \stackrel{(b)}{\leq} \frac{I}{(1-\beta)^2} \sum_{\tau=t_0}^{t-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\left\| \nabla f_i(\mathbf{x}_i^{(\tau)}) - \frac{1}{N} \sum_{j=1}^N \nabla f_j(\mathbf{x}_j^{(\tau)}) \right\|^2 \right] \\
 & \stackrel{(c)}{\leq} \frac{6L^2I}{(1-\beta)^2} \sum_{\tau=t_0}^{t-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\mathbf{x}_i^{(\tau)} - \bar{\mathbf{x}}^{(\tau)}\|^2] + \frac{3I}{(1-\beta)^2} \sum_{\tau=t_0}^{t-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\nabla f_i(\bar{\mathbf{x}}^{(\tau)}) - \nabla f(\bar{\mathbf{x}}^{(\tau)})\|^2] \\
 & \stackrel{(d)}{\leq} \frac{6L^2I}{(1-\beta)^2} \sum_{\tau=t_0}^{t-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\mathbf{x}_i^{(\tau)} - \bar{\mathbf{x}}^{(\tau)}\|^2] + \frac{3I^2\kappa^2}{(1-\beta)^2}
 \end{aligned} \tag{33}$$

where (a) follows by applying the basic inequality $\|\sum_{i=1}^n \mathbf{a}_i\|^2 \leq n \sum_{i=1}^n \|\mathbf{a}_i\|^2$ for any vectors \mathbf{a}_i and any integer n and noting that $|\frac{1-\beta^{t-\tau}}{1-\beta}| \leq \frac{1}{1-\beta}$; (b) follows because $0 < t-t_0 < I$; (c) follows from Lemma 2; and (d) follows because $\frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\nabla f_i(\bar{\mathbf{x}}^{(\tau)}) - \nabla f(\bar{\mathbf{x}}^{(\tau)})\|^2] \leq \kappa^2$ by Assumption 1 and $0 < t-t_0 < I$.

Substituting (32)-(33) into (31) yields

$$\frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] \leq \frac{2\gamma^2 I \sigma^2}{(1-\beta)^2} + \frac{12L^2\gamma^2 I}{(1-\beta)^2} \sum_{\tau=t_0}^{t-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\mathbf{x}_i^{(\tau)} - \bar{\mathbf{x}}^{(\tau)}\|^2] + \frac{6\gamma^2 I^2 \kappa^2}{(1-\beta)^2} \tag{34}$$

Recall that for each t , the index t_0 in the above equation is the largest integer such that $t_0 < t$, $t_0 \bmod I = 0$ and $t-t_0 < I$. Summing (34) over $t \in \{0, 1, \dots, T\}$ that are not a multiple of I and noting that $\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)} = \mathbf{0}$ if t is a multiple of I yields

$$\sum_{t=0}^{T-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] \leq \frac{2\gamma^2 I \sigma^2}{(1-\beta)^2} T + \frac{12L^2\gamma^2 I^2}{(1-\beta)^2} \sum_{t=0}^{T-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] + \frac{6\gamma^2 I^2 \kappa^2}{(1-\beta)^2} T$$

Collecting common terms and dividing both sides by $1 - \frac{12L^2\gamma^2 I^2}{(1-\beta)^2}$ yields

$$\sum_{t=0}^{T-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] \leq \frac{1}{1 - \frac{12L^2\gamma^2 I^2}{(1-\beta)^2}} \frac{2\gamma^2 I \sigma^2}{(1-\beta)^2} T + \frac{1}{1 - \frac{12L^2\gamma^2 I^2}{(1-\beta)^2}} \frac{6\gamma^2 I^2 \kappa^2}{(1-\beta)^2} T \tag{35}$$

□

6.3.2. MAIN PROOF OF THEOREM 1

With the above lemmas, we are ready to present the main proof of Theorem 1.

Fix $t \geq 0$. By the smoothness of function $f(\cdot)$ (in Assumption 1), we have

$$\mathbb{E}[f(\bar{\mathbf{z}}^{(t+1)})] \leq \mathbb{E}[f(\bar{\mathbf{z}}^{(t)})] + \mathbb{E}[\langle \nabla f(\bar{\mathbf{z}}^{(t)}), \bar{\mathbf{z}}^{(t+1)} - \bar{\mathbf{z}}^{(t)} \rangle] + \frac{L}{2} \mathbb{E}[\|\bar{\mathbf{z}}^{(t+1)} - \bar{\mathbf{z}}^{(t)}\|^2] \quad (36)$$

By Lemma 3, we have

$$\begin{aligned} & \mathbb{E}[\langle \nabla f(\bar{\mathbf{z}}^{(t)}), \bar{\mathbf{z}}^{(t+1)} - \bar{\mathbf{z}}^{(t)} \rangle] \\ &= -\frac{\gamma}{1-\beta} \mathbb{E}[\langle \nabla f(\bar{\mathbf{z}}^{(t)}), \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)} \rangle] \\ &\stackrel{(a)}{=} -\frac{\gamma}{1-\beta} \mathbb{E}[\langle \nabla f(\bar{\mathbf{z}}^{(t)}), \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \rangle] \\ &= -\frac{\gamma}{1-\beta} \mathbb{E}[\langle \nabla f(\bar{\mathbf{z}}^{(t)}) - \nabla f(\bar{\mathbf{x}}^{(t)}), \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \rangle] - \frac{\gamma}{1-\beta} \mathbb{E}[\langle \nabla f(\bar{\mathbf{x}}^{(t)}), \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \rangle] \end{aligned} \quad (37)$$

where (a) follows because $\bar{\mathbf{z}}^{(t)}$ and $\mathbf{x}_i^{(t)}$ are determined by $\boldsymbol{\xi}^{[t-1]} = [\boldsymbol{\xi}^{(0)}, \dots, \boldsymbol{\xi}^{(t-1)}]$, which is independent of $\boldsymbol{\xi}^{(t)}$, and $\mathbb{E}[\mathbf{g}_i^{(t)} | \boldsymbol{\xi}^{[t-1]}] = \mathbb{E}[\mathbf{g}_i^{(t)}] = \nabla f_i(\mathbf{x}_i^{(t)})$.

We note that

$$\begin{aligned} & -\frac{\gamma}{1-\beta} \langle \nabla f(\bar{\mathbf{z}}^{(t)}) - \nabla f(\bar{\mathbf{x}}^{(t)}), \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \rangle \\ & \stackrel{(a)}{\leq} \frac{1-\beta}{2\beta L} \|\nabla f(\bar{\mathbf{z}}^{(t)}) - \nabla f(\bar{\mathbf{x}}^{(t)})\|^2 + \frac{\beta L \gamma^2}{2(1-\beta)^3} \left\| \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \right\|^2 \\ & \stackrel{(b)}{\leq} \frac{(1-\beta)L}{2\beta} \|\bar{\mathbf{z}}^{(t)} - \bar{\mathbf{x}}^{(t)}\|^2 + \frac{\beta L \gamma^2}{2(1-\beta)^3} \left\| \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \right\|^2 \end{aligned} \quad (38)$$

where (a) follows by applying the basic inequality $\langle \mathbf{a}, \mathbf{b} \rangle \leq \frac{1}{2} \|\mathbf{a}\|^2 + \frac{1}{2} \|\mathbf{b}\|^2$ with $\mathbf{a} = -\frac{\sqrt{1-\beta}}{\sqrt{\beta L}} [\nabla f(\bar{\mathbf{z}}^{(t)}) - \nabla f(\bar{\mathbf{x}}^{(t)})]$ and $\mathbf{b} = \frac{\gamma \sqrt{\beta L}}{(1-\beta)^{3/2}} \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})$; and (b) follows from the smoothness of function $f(\cdot)$.

Applying the basic identity $\langle \mathbf{a}, \mathbf{b} \rangle = \frac{1}{2} [\|\mathbf{a}\|^2 + \|\mathbf{b}\|^2 - \|\mathbf{a} - \mathbf{b}\|^2]$ with $\mathbf{a} = \nabla f(\bar{\mathbf{x}}^{(t)})$ and $\mathbf{b} = \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})$ yields

$$\begin{aligned} & \langle \nabla f(\bar{\mathbf{x}}^{(t)}), \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \rangle \\ &= \frac{1}{2} \left(\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2 + \left\| \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \right\|^2 - \left\| \nabla f(\bar{\mathbf{x}}^{(t)}) - \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \right\|^2 \right) \\ & \stackrel{(a)}{\geq} \frac{1}{2} \left(\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2 + \left\| \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \right\|^2 - L^2 \frac{1}{N} \sum_{i=1}^N \|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2 \right) \end{aligned} \quad (39)$$

where (a) follows because $\left\| \nabla f(\bar{\mathbf{x}}^{(t)}) - \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \right\|^2 = \left\| \frac{1}{N} \sum_{i=1}^N \nabla f_i(\bar{\mathbf{x}}^{(t)}) - \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \right\|^2 \leq \frac{1}{N} \sum_{i=1}^N \|\nabla f_i(\bar{\mathbf{x}}^{(t)}) - \nabla f_i(\mathbf{x}_i^{(t)})\|^2 \leq \frac{1}{N} \sum_{i=1}^N L^2 \|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2$, where the first inequality follows from the convexity of $\|\cdot\|^2$ and Jensen's inequality and the second inequality follows from the smoothness of each $f_i(\cdot)$.

Substituting (38)-(39) into (37) yields

$$\begin{aligned}
 & \mathbb{E}[\langle \nabla f(\bar{\mathbf{z}}^{(t)}), \bar{\mathbf{z}}^{(t+1)} - \bar{\mathbf{z}}^{(t)} \rangle] \\
 & \leq \frac{(1-\beta)L}{2\beta} \mathbb{E}[\|\bar{\mathbf{z}}^{(t)} - \bar{\mathbf{x}}^{(t)}\|^2] + \left(\frac{\beta L \gamma^2}{2(1-\beta)^3} - \frac{\gamma}{2(1-\beta)} \right) \mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})\|^2] - \frac{\gamma}{2(1-\beta)} \mathbb{E}[\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2] \\
 & \quad + \frac{\gamma L^2}{2(1-\beta)} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] \tag{40}
 \end{aligned}$$

By Lemma 3, we have

$$\mathbb{E}[\|\bar{\mathbf{z}}^{(t+1)} - \bar{\mathbf{z}}^{(t)}\|^2] = \frac{\gamma^2}{(1-\beta)^2} \mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)}\|^2] \tag{41}$$

Substituting (40)-(41) into (36) yields

$$\begin{aligned}
 \mathbb{E}[f(\bar{\mathbf{z}}^{(t+1)})] & \leq \mathbb{E}[f(\bar{\mathbf{z}}^{(t)})] + \frac{(1-\beta)L}{2\beta} \mathbb{E}[\|\bar{\mathbf{z}}^{(t)} - \bar{\mathbf{x}}^{(t)}\|^2] + \left(\frac{\beta L \gamma^2}{2(1-\beta)^3} - \frac{\gamma}{2(1-\beta)} \right) \mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})\|^2] \\
 & \quad - \frac{\gamma}{2(1-\beta)} \mathbb{E}[\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2] + \frac{\gamma L^2}{2(1-\beta)} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] + \frac{\gamma^2 L}{2(1-\beta)^2} \mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)}\|^2] \tag{42}
 \end{aligned}$$

Dividing both sides by $\frac{\gamma}{2(1-\beta)}$ and rearranging terms yields

$$\begin{aligned}
 \mathbb{E}[\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2] & \leq \frac{2(1-\beta)}{\gamma} (\mathbb{E}[f(\bar{\mathbf{z}}^{(t)})] - \mathbb{E}[f(\bar{\mathbf{z}}^{(t+1)})]) - \left(1 - \frac{L\gamma\beta}{(1-\beta)^2}\right) \mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})\|^2] \\
 & \quad + \frac{(1-\beta)^2 L}{\beta\gamma} \mathbb{E}[\|\bar{\mathbf{z}}^{(t)} - \bar{\mathbf{x}}^{(t)}\|^2] + \frac{L\gamma}{(1-\beta)} \mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)}\|^2] + L^2 \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] \tag{43}
 \end{aligned}$$

Summing over $t \in \{0, 1, \dots, T-1\}$

$$\begin{aligned}
 & \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2] \\
 & \leq \frac{2(1-\beta)}{\gamma} (\mathbb{E}[f(\bar{\mathbf{z}}^{(0)})] - \mathbb{E}[f(\bar{\mathbf{z}}^{(T)})]) - \left(1 - \frac{L\gamma\beta}{(1-\beta)^2}\right) \sum_{t=0}^{T-1} \mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})\|^2] \\
 & \quad + \frac{(1-\beta)^2 L}{\beta\gamma} \sum_{t=0}^{T-1} \mathbb{E}[\|\bar{\mathbf{z}}^{(t)} - \bar{\mathbf{x}}^{(t)}\|^2] + \frac{L\gamma}{(1-\beta)} \sum_{t=0}^{T-1} \mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)}\|^2] + L^2 \sum_{t=0}^{T-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] \\
 & \stackrel{(a)}{\leq} \frac{2(1-\beta)}{\gamma} (\mathbb{E}[f(\bar{\mathbf{z}}^{(0)})] - \mathbb{E}[f(\bar{\mathbf{z}}^{(T)})]) - \left(1 - \frac{L\gamma\beta}{(1-\beta)^2}\right) \sum_{t=0}^{T-1} \mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})\|^2] \\
 & \quad + \frac{L\gamma}{(1-\beta)^2} \sum_{t=0}^{T-1} \mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)}\|^2] + L^2 \sum_{t=0}^{T-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] \\
 & \stackrel{(b)}{\leq} \frac{2(1-\beta)}{\gamma} (\mathbb{E}[f(\bar{\mathbf{z}}^{(0)})] - \mathbb{E}[f(\bar{\mathbf{z}}^{(T)})]) - \left(1 - \frac{(1+\beta)L\gamma}{(1-\beta)^2}\right) \sum_{t=0}^{T-1} \mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})\|^2] + \frac{L\gamma}{(1-\beta)^2} \frac{\sigma^2}{N} T \\
 & \quad + \frac{1}{1 - \frac{12L^2\gamma^2 I^2}{(1-\beta)^2}} \frac{2L^2\gamma^2 I\sigma^2}{(1-\beta)^2} T + \frac{1}{1 - \frac{12L^2\gamma^2 I^2}{(1-\beta)^2}} \frac{6L^2\gamma^2 I^2 \kappa^2}{(1-\beta)^2} T
 \end{aligned}$$

$$\stackrel{(c)}{\leq} \frac{2(1-\beta)}{\gamma} (f(\bar{\mathbf{x}}^{(0)}) - f^*) + \frac{L\gamma}{(1-\beta)^2} \frac{\sigma^2}{N} T + \frac{1}{1 - \frac{12L^2\gamma^2 I^2}{(1-\beta)^2}} \frac{2L^2\gamma^2 I\sigma^2}{(1-\beta)^2} T + \frac{1}{1 - \frac{12L^2\gamma^2 I^2}{(1-\beta)^2}} \frac{6L^2\gamma^2 I^2\kappa^2}{(1-\beta)^2} T \quad (44)$$

where (a) follow by using Lemma 4 and $\frac{\beta L\gamma}{(1-\beta)^2} + \frac{L\gamma}{1-\beta} = \frac{L\gamma}{(1-\beta)^2}$; (b) follows by applying Lemma 5 and by noting that $\mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)}\|^2] \leq \frac{1}{N} \sigma^2 + \mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})\|^2]$ by Lemma 1; and (c) follows because γ is chosen to ensure $1 - \frac{(1+\beta)L\gamma}{(1-\beta)^2} \geq 0$, $\bar{\mathbf{z}}^{(0)} = \bar{\mathbf{x}}^{(0)}$ by definition in (19), and f^* is the minimum value of problem (1).

Dividing both sides by T yields

$$\begin{aligned} & \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2] \\ & \leq \frac{2(1-\beta)}{\gamma T} (f(\bar{\mathbf{x}}^{(0)}) - f^*) + \frac{L\gamma}{(1-\beta)^2} \frac{\sigma^2}{N} + \frac{1}{1 - \frac{12L^2\gamma^2 I^2}{(1-\beta)^2}} \frac{2L^2\gamma^2 I\sigma^2}{(1-\beta)^2} + \frac{1}{1 - \frac{12L^2\gamma^2 I^2}{(1-\beta)^2}} \frac{6L^2\gamma^2 I^2\kappa^2}{(1-\beta)^2} \\ & \stackrel{(a)}{\leq} \frac{2(1-\beta)}{\gamma T} (f(\bar{\mathbf{x}}^{(0)}) - f^*) + \frac{L\gamma}{(1-\beta)^2} \frac{\sigma^2}{N} + \frac{3L^2\gamma^2 I\sigma^2}{(1-\beta)^2} + \frac{9L^2\gamma^2 I^2\kappa^2}{(1-\beta)^2} \\ & = O\left(\frac{1}{\gamma T}\right) + O\left(\frac{\gamma}{N}\sigma^2\right) + O(\gamma^2 I\sigma^2) + O(\gamma^2 I^2\kappa^2) \end{aligned} \quad (45)$$

where (a) follows because $I \leq \frac{1-\beta}{6L\gamma}$ is chosen to ensure $\frac{1}{1 - \frac{12L^2\gamma^2 I^2}{(1-\beta)^2}} \leq \frac{3}{2}$.

6.4. On the equivalence between (5) and (14)

In this subsection, we show both and yield the same solution sequences $\{\mathbf{x}_i^{(t)}\}_{t \geq 0}$ assume they are initialized at the same $\mathbf{x}_i^{(0)}$. It is easy to verify that (5) and (14) yield the same $\mathbf{x}_i^{(1)}$ by noting that $\mathbf{y}_i^{(0)} = 0$ and $\mathbf{u}_i^{(0)} = \mathbf{0}$. Next, we show they yield the same $\{\mathbf{x}_i^{(t)}\}_{t \geq 2}$.

Substituting the second equation of (5) into the third equation of (5) yields

$$\mathbf{x}_i^{(t)} = \mathbf{x}_i^{(t-1)} - \gamma\beta\mathbf{u}_i^{(t)} - \gamma\mathbf{g}_i^{(t-1)}, \forall t \geq 1 \quad (46)$$

Fix $t \geq 2$. Substituting the first equation of (5) into the above equation yields

$$\begin{aligned} \mathbf{x}_i^{(t)} &= \mathbf{x}_i^{(t-1)} - \gamma\mathbf{g}_i^{(t-1)} - \gamma\beta\mathbf{g}_i^{(t-1)} - \gamma\beta^2\mathbf{u}_i^{(t-1)} \\ &\stackrel{(a)}{=} \mathbf{x}_i^{(t-1)} - \gamma\mathbf{g}_i^{(t-1)} - \gamma\beta\mathbf{g}_i^{(t-1)} + \gamma\beta\mathbf{g}_i^{(t-2)} + \beta[\mathbf{x}_i^{(t-1)} - \mathbf{x}_i^{(t-2)}] \\ &= \mathbf{x}_i^{(t-1)} - \gamma\mathbf{g}_i^{(t-1)} + \beta[[\mathbf{x}_i^{(t-1)} - \gamma\mathbf{g}_i^{(t-1)}] - [\mathbf{x}_i^{(t-2)} - \gamma\mathbf{g}_i^{(t-2)}]], \forall t \geq 2. \end{aligned} \quad (47)$$

where (a) follows by noting that $-\gamma\beta^2\mathbf{u}_i^{(t-1)} = \gamma\beta\mathbf{g}_i^{(t-2)} + \beta[\mathbf{x}_i^{(t-1)} - \mathbf{x}_i^{(t-2)}]$, $\forall t \geq 2$ from (46).

Now if we define $\mathbf{y}_i^{(t)} \triangleq \mathbf{x}_i^{(t-1)} - \gamma\mathbf{g}_i^{(t-1)}$, which is the first equation in (14), then the above equation can be written as

$$\mathbf{x}_i^{(t)} = \mathbf{y}_i^{(t)} + \beta[\mathbf{y}_i^{(t)} - \mathbf{y}_i^{(t-1)}], \quad \forall t \geq 2 \quad (48)$$

which is exactly the second equation in (14).

Thus, we can conclude that (5) and (14) yield the same solution sequences $\{\mathbf{x}_i^{(t)}\}_{t \geq 0}$ (with different auxiliary/buffer sequences.)

6.5. Proof of Theorem 2

To establish the convergence rate for Algorithm 1 with Nesterov's momentum, we introduce a different auxiliary sequence $\{\bar{\mathbf{y}}^{(t)}\}_{t \geq 0}$ given by

$$\bar{\mathbf{y}}^{(t)} \triangleq \begin{cases} \bar{\mathbf{x}}^{(t)}, & t = 0 \\ \frac{1}{1-\beta}\bar{\mathbf{x}}^{(t)} - \frac{\beta}{1-\beta}\bar{\mathbf{x}}^{(t-1)} + \frac{\gamma\beta}{1-\beta} \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t-1)}, & t \geq 1 \end{cases} \quad (49)$$

where $\bar{\mathbf{x}}^{(t)}$ defined in (7) are the node averages of local solutions from Algorithm 1 with Option II.

Using the notation $\bar{\mathbf{u}}^{(t)} \triangleq \frac{1}{N} \sum_{i=1}^N \mathbf{u}_i^{(t)}$ in (9) and $\bar{\mathbf{v}}^{(t)} \triangleq \frac{1}{N} \sum_{i=1}^N \mathbf{v}_i^{(t)}$ with $\mathbf{u}_i^{(t)}$ and $\mathbf{v}_i^{(t)}$ from Algorithm 1 with Option II, we shall have

$$\begin{cases} \bar{\mathbf{u}}^{(t)} &= \beta \bar{\mathbf{u}}^{(t-1)} + \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t-1)} \\ \bar{\mathbf{v}}^{(t)} &= \beta \bar{\mathbf{u}}^{(t)} + \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t-1)} \\ \bar{\mathbf{x}}^{(t)} &= \bar{\mathbf{x}}^{(t-1)} - \gamma \bar{\mathbf{v}}^{(t)} \end{cases} \quad (50)$$

The main proof shall follow similar steps as in our main proof for Polyak's momentum in Section 6.3.2 as long as we prove the counterpart of Lemmas 3-5 for Algorithm 1 with Option II.

6.5.1. USEFUL LEMMAS

While the sequence $\{\bar{\mathbf{y}}^{(t)}\}_{t \geq 0}$ introduced in (49) is different from $\{\bar{\mathbf{z}}^{(t)}\}_{t \geq 0}$ defined in (19) for Polyak's momentum, the next lemma shows that $\bar{\mathbf{y}}^{(t+1)} - \bar{\mathbf{y}}^{(t)}$ for Algorithm 1 with Option II is equal to $\bar{\mathbf{z}}^{(t+1)} - \bar{\mathbf{z}}^{(t)}$ for Algorithm 1 with Option I.

Lemma 6. Consider the sequence $\{\bar{\mathbf{y}}^{(t)}\}_{t \geq 0}$ defined in (49). Algorithm 1 with Option II ensures that for all $t \geq 0$, we have

$$\bar{\mathbf{y}}^{(t+1)} - \bar{\mathbf{y}}^{(t)} = -\frac{\gamma}{1-\beta} \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)}. \quad (51)$$

Proof. The case $\bar{\mathbf{y}}^{(1)} - \bar{\mathbf{y}}^{(0)} = -\frac{\gamma}{1-\beta} \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(0)}$ can be shown directly by combining (49) and (50) with $t = 0$. To prove the case $t \geq 1$, we note that

$$\begin{aligned} & \bar{\mathbf{y}}^{(t+1)} - \bar{\mathbf{y}}^{(t)} \\ \stackrel{(a)}{=} & \frac{1}{1-\beta} [\bar{\mathbf{x}}^{(t+1)} - \bar{\mathbf{x}}^{(t)}] - \frac{\beta}{1-\beta} [\bar{\mathbf{x}}^{(t)} - \bar{\mathbf{x}}^{(t-1)}] + \frac{\gamma\beta}{1-\beta} \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)} - \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t-1)} \right] \\ \stackrel{(b)}{=} & -\frac{\gamma}{1-\beta} \bar{\mathbf{v}}^{(t+1)} + \frac{\gamma\beta}{1-\beta} \bar{\mathbf{v}}^{(t)} + \frac{\gamma\beta}{1-\beta} \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)} - \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t-1)} \right] \\ \stackrel{(c)}{=} & -\frac{\gamma\beta^2}{1-\beta} \bar{\mathbf{u}}^{(t)} - \frac{\gamma(1+\beta)}{(1-\beta)} \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)} + \frac{\gamma\beta^2}{1-\beta} \bar{\mathbf{u}}^{(t)} + \frac{\gamma\beta}{1-\beta} \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t-1)} + \frac{\gamma\beta}{1-\beta} \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)} - \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t-1)} \right] \\ = & -\frac{\gamma}{1-\beta} \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)} \end{aligned}$$

where (a) follows from (49); (b) follows from the third equation in (50); and (c) follows by substituting $\bar{\mathbf{v}}^{(t)} = \beta \bar{\mathbf{u}}^{(t)} + \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t-1)}$ and $\bar{\mathbf{v}}^{(t+1)} = \beta^2 \bar{\mathbf{u}}^{(t)} + (1+\beta) \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)}$, which is implied by combining the first and second equations in (50). \square

Lemma 7. Let $\{\bar{\mathbf{x}}^{(t)}\}_{t \geq 0}$ defined in (7) be the node averages of local solutions from Algorithm 1 with Option II. Let $\{\bar{\mathbf{y}}^{(t)}\}_{t \geq 0}$ be defined in (49). For all $T \geq 1$, Algorithm 1 with Option II ensures that

$$\sum_{t=0}^{T-1} \|\bar{\mathbf{y}}^{(t)} - \bar{\mathbf{x}}^{(t)}\|^2 \leq \frac{\gamma^2 \beta^4}{(1-\beta)^4} \sum_{t=0}^{T-1} \left\| \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)} \right] \right\|^2 \quad (52)$$

Proof. Recall that $\bar{\mathbf{u}}^{(0)} = \mathbf{0}$. Recursively applying the first equation in (50) for t times yields:

$$\bar{\mathbf{u}}^{(t)} = \sum_{\tau=0}^{t-1} \beta^{t-1-\tau} \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(\tau)} \right], \quad \forall t \geq 1 \quad (53)$$

Note that for all $t \geq 1$, we have

$$\begin{aligned}
 \bar{\mathbf{y}}^{(t)} - \bar{\mathbf{x}}^{(t)} &\stackrel{(a)}{=} \frac{\beta}{1-\beta} [\bar{\mathbf{x}}^{(t)} - \bar{\mathbf{x}}^{(t-1)}] + \frac{\gamma\beta}{1-\beta} \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t-1)} \\
 &\stackrel{(b)}{=} -\frac{\gamma\beta}{1-\beta} \bar{\mathbf{v}}^{(t)} + \frac{\gamma\beta}{1-\beta} \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t-1)} \\
 &\stackrel{(c)}{=} -\frac{\gamma\beta^2}{1-\beta} \bar{\mathbf{u}}^{(t)}
 \end{aligned} \tag{54}$$

where (a) follows from (49); (b) follows from the third equation in (50); and (c) follows from the second equation in (50).

Combining (53) and (54) yields

$$\bar{\mathbf{y}}^{(t)} - \bar{\mathbf{x}}^{(t)} = -\frac{\gamma\beta^2}{1-\beta} \sum_{\tau=0}^{t-1} \beta^{t-1-\tau} \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(\tau)} \right], \quad \forall t \geq 1 \tag{55}$$

Define $s_t \triangleq \sum_{\tau=0}^{t-1} \beta^{t-1-\tau} = \frac{1-\beta^t}{1-\beta}$. For all $t \geq 1$, we have

$$\begin{aligned}
 \|\bar{\mathbf{y}}^{(t)} - \bar{\mathbf{x}}^{(t)}\|^2 &= \frac{\gamma^2\beta^4}{(1-\beta)^2} s_t^2 \left\| \sum_{\tau=0}^{t-1} \frac{\beta^{t-1-\tau}}{s_t} \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(\tau)} \right] \right\|^2 \\
 &\stackrel{(a)}{\leq} \frac{\gamma^2\beta^4}{(1-\beta)^2} s_t^2 \sum_{\tau=0}^{t-1} \frac{\beta^{t-1-\tau}}{s_t} \left\| \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(\tau)} \right] \right\|^2 \\
 &= \frac{\gamma^2\beta^4(1-\beta^t)}{(1-\beta)^3} \sum_{\tau=0}^{t-1} \beta^{t-1-\tau} \left\| \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(\tau)} \right] \right\|^2 \\
 &\leq \frac{\gamma^2\beta^4}{(1-\beta)^3} \sum_{\tau=0}^{t-1} \beta^{t-1-\tau} \left\| \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(\tau)} \right] \right\|^2
 \end{aligned} \tag{56}$$

where (a) follows from the convexity of $\|\cdot\|^2$ and Jensen's inequality.

Fix $T \geq 1$. Recall that $\bar{\mathbf{y}}^{(0)} - \bar{\mathbf{x}}^{(0)} = \mathbf{0}$. Summing (56) over $t \in \{1, 2, \dots, T-1\}$ yields

$$\begin{aligned}
 \sum_{t=0}^{T-1} \|\bar{\mathbf{y}}^{(t)} - \bar{\mathbf{x}}^{(t)}\|^2 &\leq \frac{\gamma^2\beta^4}{(1-\beta)^3} \sum_{t=1}^{T-1} \sum_{\tau=0}^{t-1} \beta^{t-1-\tau} \left\| \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(\tau)} \right] \right\|^2 \\
 &= \frac{\gamma^2\beta^4}{(1-\beta)^3} \sum_{\tau=0}^{T-2} \left(\left\| \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(\tau)} \right] \right\|^2 \sum_{l=\tau+1}^{T-1} \beta^{l-1-\tau} \right) \\
 &\stackrel{(a)}{\leq} \frac{\gamma^2\beta^4}{(1-\beta)^4} \sum_{\tau=0}^{T-2} \left\| \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(\tau)} \right] \right\|^2 \\
 &\leq \frac{\gamma^2\beta^4}{(1-\beta)^4} \sum_{\tau=0}^{T-1} \left\| \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(\tau)} \right] \right\|^2
 \end{aligned} \tag{57}$$

where (a) follows by noting that $\sum_{l=\tau+1}^{T-1} \beta^{l-1-\tau} = \frac{1-\beta^{T-\tau-1}}{1-\beta} \leq \frac{1}{1-\beta}$. □

Lemma 8. Consider problem (1) under Assumption 1. Let $\{\bar{\mathbf{x}}^{(t)}\}_{t \geq 0}$ defined in (7) be the node averages of local solutions from Algorithm 1 with Option II. If γ and I are chosen to satisfy $\frac{12L^2\gamma^2I^2}{(1-\beta)^2} < 1$, then for all $T \geq 1$, we have

$$\sum_{t=0}^{T-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] \leq \frac{1}{1 - \frac{12L^2\gamma^2I^2}{(1-\beta)^2}} \frac{2\gamma^2I\sigma^2}{(1-\beta)^2} T + \frac{1}{1 - \frac{12L^2\gamma^2I^2}{(1-\beta)^2}} \frac{6\gamma^2I^2\kappa^2}{(1-\beta)^2} T$$

Proof. Since $\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)} = \mathbf{0}$ when t is a multiple of I , our focus is to develop upper bounds for $\mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2]$ when t is not a multiple of I . Consider $t \geq 1$ that is not a multiple of I . Note that Algorithm 1 restarts ‘‘SGD with momentum’’ every I iterations (by resetting $\mathbf{u}_i^{(t)} = \hat{\mathbf{u}} \triangleq \frac{1}{N} \sum_{j=1}^N \mathbf{u}_j^{(t)}$ and $\mathbf{x}_i^{(t)} = \hat{\mathbf{x}} \triangleq \frac{1}{N} \sum_{j=1}^N \mathbf{x}_j^{(t)}$). Let $t_0 < t$ be the largest iteration index that is a multiple of I , i.e., $t_0 \bmod I = 0$. Note that we must have $t - t_0 < I$ and $\mathbf{u}_i^{(t_0)} = \hat{\mathbf{u}}, \mathbf{x}_i^{(t_0)} = \hat{\mathbf{x}}$. For any $\tau \in \{t_0 + 1, \dots, t\}$, by recursively applying the first equation in (5), we have

$$\mathbf{u}_i^{(\tau)} = \beta^{\tau-t_0} \hat{\mathbf{u}} + \sum_{k=t_0}^{\tau-1} \beta^{\tau-1-k} \mathbf{g}_i^{(k)} \quad (58)$$

By the second equation in (5), this further implies

$$\mathbf{v}_i^{(\tau)} = \beta^{\tau+1-t_0} \hat{\mathbf{u}} + \sum_{k=t_0}^{\tau-1} \beta^{\tau-k} \mathbf{g}_i^{(k)} + \mathbf{g}_i^{(\tau-1)}, \quad \forall \tau \in \{t_0 + 1, \dots, t\}. \quad (59)$$

For any $\tau \in \{t_0 + 1, \dots, t\}$, by the third equation in (5), we have

$$\mathbf{x}_i^{(\tau)} - \mathbf{x}_i^{(\tau-1)} = -\gamma \mathbf{v}_i^{(\tau)} \quad (60)$$

Summing over $\tau \in \{t_0 + 1, \dots, t\}$ and noting that $\mathbf{x}_i^{(t_0)} = \hat{\mathbf{x}}$ yields

$$\begin{aligned} \mathbf{x}_i^{(t)} &= \hat{\mathbf{x}} - \gamma \sum_{\tau=t_0+1}^t \mathbf{v}_i^{(\tau)} \\ &\stackrel{(a)}{=} \hat{\mathbf{x}} - \gamma \left(\sum_{\tau=t_0+1}^t \beta^{\tau+1-t_0} \right) \hat{\mathbf{u}} - \gamma \sum_{\tau=t_0+1}^t \sum_{k=t_0}^{\tau-1} \beta^{\tau-k} \mathbf{g}_i^{(k)} - \gamma \sum_{\tau=t_0+1}^t \mathbf{g}_i^{(\tau-1)} \\ &= \hat{\mathbf{x}} - \gamma \left(\sum_{j=2}^{t+1-t_0} \beta^j \right) \hat{\mathbf{u}} - \gamma \sum_{k=t_0}^{t-1} \sum_{j=1}^{t-k} \beta^j \mathbf{g}_i^{(k)} - \gamma \sum_{\tau=t_0}^{t-1} \mathbf{g}_i^{(\tau)} \\ &= \hat{\mathbf{x}} - \gamma \left(\sum_{j=2}^{t+1-t_0} \beta^j \right) \hat{\mathbf{u}} - \gamma \sum_{\tau=t_0}^{t-1} \frac{1 - \beta^{t+1-\tau}}{1 - \beta} \mathbf{g}_i^{(\tau)} \end{aligned} \quad (61)$$

where (a) follows by substituting (59).

Using a argument similar to the above, we can further show

$$\bar{\mathbf{x}}^{(t)} = \hat{\mathbf{x}} - \gamma \left(\sum_{j=2}^{t+1-t_0} \beta^j \right) \hat{\mathbf{u}} - \gamma \sum_{\tau=t_0}^{t-1} \frac{1 - \beta^{t+1-\tau}}{1 - \beta} \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(\tau)} \quad (62)$$

Combining (61) and (62) yields

$$\begin{aligned} &\frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] \\ &= \gamma^2 \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\left\| \sum_{\tau=t_0}^{t-1} [\mathbf{g}_i^{(\tau)} - \frac{1}{N} \sum_{j=1}^N \mathbf{g}_j^{(\tau)}] \frac{1 - \beta^{t+1-\tau}}{1 - \beta} \right\|^2 \right] \\ &\stackrel{(a)}{\leq} 2\gamma^2 \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\left\| \sum_{\tau=t_0}^{t-1} [[\mathbf{g}_i^{(\tau)} - \nabla f_i(\mathbf{x}_i^{(\tau)})] - \frac{1}{N} \sum_{j=1}^N [\mathbf{g}_j^{(\tau)} - \nabla f_j(\mathbf{x}_j^{(\tau)})]] \frac{1 - \beta^{t+1-\tau}}{1 - \beta} \right\|^2 \right] \\ &\quad + 2\gamma^2 \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\left\| \sum_{\tau=t_0}^{t-1} [\nabla f_i(\mathbf{x}_i^{(\tau)}) - \frac{1}{N} \sum_{j=1}^N \nabla f_j(\mathbf{x}_j^{(\tau)})] \frac{1 - \beta^{t+1-\tau}}{1 - \beta} \right\|^2 \right] \end{aligned} \quad (63)$$

where (a) follows from the basic inequality $\|\mathbf{a}_1 + \mathbf{a}_2\|^2 \leq 2\|\mathbf{a}_1\|^2 + 2\|\mathbf{a}_2\|^2$.

Now we develop the respective upper bounds of the two terms on the right side of (63). We note that

$$\begin{aligned}
 & \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\left\| \sum_{\tau=t_0}^{t-1} [\mathbf{g}_i^{(\tau)} - \nabla f_i(\mathbf{x}_i^{(\tau)})] - \frac{1}{N} \sum_{j=1}^N [\mathbf{g}_j^{(\tau)} - \nabla f_j(\mathbf{x}_j^{(\tau)})] \right\| \frac{1 - \beta^{t+1-\tau}}{1 - \beta} \right]^2 \\
 & \stackrel{(a)}{\leq} \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\left\| \sum_{\tau=t_0}^{t-1} [\mathbf{g}_i^{(\tau)} - \nabla f_i(\mathbf{x}_i^{(\tau)})] \right\| \frac{1 - \beta^{t+1-\tau}}{1 - \beta} \right]^2 \\
 & \stackrel{(b)}{=} \frac{1}{N} \sum_{i=1}^N \sum_{\tau=t_0}^{t-1} \mathbb{E} \left[\left\| [\mathbf{g}_i^{(\tau)} - \nabla f_i(\mathbf{x}_i^{(\tau)})] \right\| \frac{1 - \beta^{t+1-\tau}}{1 - \beta} \right]^2 \\
 & \stackrel{(c)}{\leq} \frac{I\sigma^2}{(1 - \beta)^2}
 \end{aligned} \tag{64}$$

where (a) follows from the inequality $\frac{1}{N} \sum_{i=1}^N \|\mathbf{a}_i - [\frac{1}{N} \sum_{j=1}^N \mathbf{a}_j]\|^2 = \frac{1}{N} \sum_{i=1}^N \|\mathbf{a}_i\|^2 - \|\frac{1}{N} \sum_{i=1}^N \mathbf{a}_i\|^2 \leq \frac{1}{N} \sum_{i=1}^N \|\mathbf{a}_i\|^2$ with $\mathbf{a}_i = \sum_{\tau=t_0}^{t-1} [\mathbf{g}_i^{(\tau)} - \nabla f_i(\mathbf{x}_i^{(\tau)})] \frac{1 - \beta^{t+1-\tau}}{1 - \beta}$; (b) follows because $\mathbb{E}[\mathbf{g}_i^{(\tau_2)} - \nabla f_i(\mathbf{x}_i^{(\tau_2)}) | \mathbf{g}_i^{(\tau_1)} - \nabla f_i(\mathbf{x}_i^{(\tau_1)})] = \mathbf{0}$ for any $\tau_2 > \tau_1$; (c) follows because $|\frac{1 - \beta^{t+1-\tau}}{1 - \beta}| \leq \frac{1}{1 - \beta}$, $t - t_0 < I$ and $\mathbb{E}[\|\mathbf{g}_i^{(\tau)} - \nabla f_i(\mathbf{x}_i^{(\tau)})\|^2] \leq \sigma^2$ (by Assumption 1).

We further note that the

$$\begin{aligned}
 & \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\left\| \sum_{\tau=t_0}^{t-1} [\nabla f_i(\mathbf{x}_i^{(\tau)}) - \frac{1}{N} \sum_{j=1}^N \nabla f_j(\mathbf{x}_j^{(\tau)})] \right\| \frac{1 - \beta^{t+1-\tau}}{1 - \beta} \right]^2 \\
 & \stackrel{(a)}{\leq} \frac{1}{N} \sum_{i=1}^N (t - t_0) \frac{1}{(1 - \beta)^2} \sum_{\tau=t_0}^{t-1} \mathbb{E} \left[\left\| \nabla f_i(\mathbf{x}_i^{(\tau)}) - \frac{1}{N} \sum_{j=1}^N \nabla f_j(\mathbf{x}_j^{(\tau)}) \right\|^2 \right] \\
 & \stackrel{(b)}{\leq} \frac{I}{(1 - \beta)^2} \sum_{\tau=t_0}^{t-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\left\| [\nabla f_i(\mathbf{x}_i^{(\tau)}) - \frac{1}{N} \sum_{j=1}^N \nabla f_j(\mathbf{x}_j^{(\tau)})] \right\|^2 \right] \\
 & \stackrel{(c)}{\leq} \frac{6L^2I}{(1 - \beta)^2} \sum_{\tau=t_0}^{t-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\mathbf{x}_i^{(\tau)} - \bar{\mathbf{x}}^{(\tau)}\|^2] + \frac{3I}{(1 - \beta)^2} \sum_{\tau=t_0}^{t-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\nabla f_i(\bar{\mathbf{x}}^{(\tau)}) - \nabla f(\bar{\mathbf{x}}^{(\tau)})\|^2] \\
 & \stackrel{(d)}{\leq} \frac{6L^2I}{(1 - \beta)^2} \sum_{\tau=t_0}^{t-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\mathbf{x}_i^{(\tau)} - \bar{\mathbf{x}}^{(\tau)}\|^2] + \frac{3I^2\kappa^2}{(1 - \beta)^2}
 \end{aligned} \tag{65}$$

where (a) follows by applying the basic inequality $\|\sum_{i=1}^n \mathbf{a}_i\|^2 \leq n \sum_{i=1}^n \|\mathbf{a}_i\|^2$ for any vectors \mathbf{a}_i and any integer n and noting that $|\frac{1 - \beta^{t+1-\tau}}{1 - \beta}| \leq \frac{1}{1 - \beta}$; (b) follows because $0 < t - t_0 < I$; (c) follows from Lemma 2; and (d) follows because $\frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\nabla f_i(\bar{\mathbf{x}}^{(\tau)}) - \nabla f(\bar{\mathbf{x}}^{(\tau)})\|^2] \leq \kappa^2$ by Assumption 1 and $0 < t - t_0 < I$.

Substituting (64)-(65) into (63) yields

$$\frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] \leq \frac{2\gamma^2 I \sigma^2}{(1 - \beta)^2} + \frac{12L^2\gamma^2 I}{(1 - \beta)^2} \sum_{\tau=t_0}^{t-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\mathbf{x}_i^{(\tau)} - \bar{\mathbf{x}}^{(\tau)}\|^2] + \frac{6\gamma^2 I^2 \kappa^2}{(1 - \beta)^2} \tag{66}$$

Recall that for each t , the index t_0 in the above equation is the largest integer such that $t_0 < t$, $t_0 \bmod I = 0$ and $t - t_0 < I$. Summing (66) over $t \in \{0, 1, \dots, T\}$ that are not a multiple of I and noting that $\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)} = \mathbf{0}$ if t is a multiple of I yields

$$\sum_{t=0}^{T-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] \leq \frac{2\gamma^2 I \sigma^2}{(1 - \beta)^2} T + \frac{12L^2\gamma^2 I^2}{(1 - \beta)^2} \sum_{t=0}^{T-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] + \frac{6\gamma^2 I^2 \kappa^2}{(1 - \beta)^2} T$$

Collecting common terms and dividing both sides by $1 - \frac{12L^2\gamma^2I^2}{(1-\beta)^2}$ yields

$$\sum_{t=0}^{T-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] \leq \frac{1}{1 - \frac{12L^2\gamma^2I^2}{(1-\beta)^2}} \frac{2\gamma^2I\sigma^2}{(1-\beta)^2} T + \frac{1}{1 - \frac{12L^2\gamma^2I^2}{(1-\beta)^2}} \frac{6\gamma^2I^2\kappa^2}{(1-\beta)^2} T \quad (67)$$

□

6.5.2. MAIN PROOF OF THEOREM 2

It is easy to realize that Lemmas 6, 7 and 8 developed above are the respective counterparts for Lemmas 3, 4 and 5. The main proof of Theorem 2 follows similar steps as the proof of Theorem 1 in Section 6.3 with the minor changes that Lemmas 3, 4 and 5 should be replaced by Lemmas 6, 7 and 8, respectively.

Below is the main proof of Theorem 2.

Fix $t \geq 0$. By the smoothness of function $f(\cdot)$ (in Assumption 1), we have

$$\mathbb{E}[f(\bar{\mathbf{y}}^{(t+1)})] \leq \mathbb{E}[f(\bar{\mathbf{y}}^{(t)})] + \mathbb{E}[\langle \nabla f(\bar{\mathbf{y}}^{(t)}), \bar{\mathbf{y}}^{(t+1)} - \bar{\mathbf{y}}^{(t)} \rangle] + \frac{L}{2} \mathbb{E}[\|\bar{\mathbf{y}}^{(t+1)} - \bar{\mathbf{y}}^{(t)}\|^2] \quad (68)$$

By Lemma 6, we have

$$\begin{aligned} & \mathbb{E}[\langle \nabla f(\bar{\mathbf{y}}^{(t)}), \bar{\mathbf{y}}^{(t+1)} - \bar{\mathbf{y}}^{(t)} \rangle] \\ &= -\frac{\gamma}{1-\beta} \mathbb{E}[\langle \nabla f(\bar{\mathbf{y}}^{(t)}), \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)} \rangle] \\ &\stackrel{(a)}{=} -\frac{\gamma}{1-\beta} \mathbb{E}[\langle \nabla f(\bar{\mathbf{y}}^{(t)}), \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \rangle] \\ &= -\frac{\gamma}{1-\beta} \mathbb{E}[\langle \nabla f(\bar{\mathbf{y}}^{(t)}) - \nabla f(\bar{\mathbf{x}}^{(t)}), \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \rangle] - \frac{\gamma}{1-\beta} \mathbb{E}[\langle \nabla f(\bar{\mathbf{x}}^{(t)}), \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \rangle] \end{aligned} \quad (69)$$

where (a) follows because $\bar{\mathbf{y}}^{(t)}$ and $\mathbf{x}_i^{(t)}$ are determined by $\boldsymbol{\xi}^{[t-1]} = [\boldsymbol{\xi}^{(0)}, \dots, \boldsymbol{\xi}^{(t-1)}]$, which is independent of $\boldsymbol{\xi}^{(t)}$, and $\mathbb{E}[\mathbf{g}_i^{(t)} | \boldsymbol{\xi}^{[t-1]}] = \mathbb{E}[\mathbf{g}_i^{(t)}] = \nabla f_i(\mathbf{x}_i^{(t)})$.

We note that

$$\begin{aligned} & -\frac{\gamma}{1-\beta} \langle \nabla f(\bar{\mathbf{y}}^{(t)}) - \nabla f(\bar{\mathbf{x}}^{(t)}), \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \rangle \\ &\stackrel{(a)}{\leq} \frac{1-\beta}{2\beta^3L} \|\nabla f(\bar{\mathbf{y}}^{(t)}) - \nabla f(\bar{\mathbf{x}}^{(t)})\|^2 + \frac{L\gamma^2\beta^3}{2(1-\beta)^3} \left\| \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \right\|^2 \\ &\stackrel{(b)}{\leq} \frac{(1-\beta)L}{2\beta^3} \|\bar{\mathbf{y}}^{(t)} - \bar{\mathbf{x}}^{(t)}\|^2 + \frac{\beta^3L\gamma^2}{2(1-\beta)^3} \left\| \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \right\|^2 \end{aligned} \quad (70)$$

where (a) follows by applying the basic inequality $\langle \mathbf{a}, \mathbf{b} \rangle \leq \frac{1}{2}\|\mathbf{a}\|^2 + \frac{1}{2}\|\mathbf{b}\|^2$ with $\mathbf{a} = -\frac{\sqrt{1-\beta}}{\sqrt{L}\beta^{3/2}}[\nabla f(\bar{\mathbf{y}}^{(t)}) - \nabla f(\bar{\mathbf{x}}^{(t)})]$ and $\mathbf{b} = \frac{\gamma\sqrt{L}\beta^{3/2}}{(1-\beta)^{3/2}} \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})$; and (b) follows from the smoothness of function $f(\cdot)$.

Applying the basic identity $\langle \mathbf{a}, \mathbf{b} \rangle = \frac{1}{2} [\|\mathbf{a}\|^2 + \|\mathbf{b}\|^2 - \|\mathbf{a} - \mathbf{b}\|^2]$ with $\mathbf{a} = \nabla f(\bar{\mathbf{x}}^{(t)})$ and $\mathbf{b} = \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})$ yields

$$\begin{aligned}
 & \langle \nabla f(\bar{\mathbf{x}}^{(t)}), \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \rangle \\
 &= \frac{1}{2} \left(\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2 + \left\| \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \right\|^2 - \left\| \nabla f(\bar{\mathbf{x}}^{(t)}) - \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \right\|^2 \right) \\
 &\stackrel{(a)}{\geq} \frac{1}{2} \left(\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2 + \left\| \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \right\|^2 - L^2 \frac{1}{N} \sum_{i=1}^N \|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2 \right) \tag{71}
 \end{aligned}$$

where (a) follows because $\|\nabla f(\bar{\mathbf{x}}^{(t)}) - \frac{1}{N} \sum_{j=1}^N \nabla f_j(\mathbf{x}_j^{(t)})\|^2 = \left\| \frac{1}{N} \sum_{i=1}^N \nabla f_i(\bar{\mathbf{x}}^{(t)}) - \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \right\|^2 \leq \frac{1}{N} \sum_{i=1}^N \|\nabla f_i(\bar{\mathbf{x}}^{(t)}) - \nabla f_i(\mathbf{x}_i^{(t)})\|^2 \leq \frac{1}{N} \sum_{i=1}^N L^2 \|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2$, where the first inequality follows from the convexity of $\|\cdot\|^2$ and Jensen's inequality and the second inequality follows from the smoothness of each $f_i(\cdot)$.

Substituting (70)-(71) into (69) yields

$$\begin{aligned}
 & \mathbb{E}[\langle \nabla f(\bar{\mathbf{y}}^{(t)}), \bar{\mathbf{y}}^{(t+1)} - \bar{\mathbf{y}}^{(t)} \rangle] \\
 &\leq \frac{(1-\beta)L}{2\beta^3} \mathbb{E}[\|\bar{\mathbf{y}}^{(t)} - \bar{\mathbf{x}}^{(t)}\|^2] + \left(\frac{L\gamma^2\beta^3}{2(1-\beta)^3} - \frac{\gamma}{2(1-\beta)} \right) \mathbb{E} \left[\left\| \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \right\|^2 \right] - \frac{\gamma}{2(1-\beta)} \mathbb{E}[\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2] \\
 &\quad + \frac{\gamma L^2}{2(1-\beta)} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] \tag{72}
 \end{aligned}$$

By Lemma 6, we have

$$\mathbb{E}[\|\bar{\mathbf{z}}^{(t+1)} - \bar{\mathbf{z}}^{(t)}\|^2] = \frac{\gamma^2}{(1-\beta)^2} \mathbb{E} \left[\left\| \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)} \right\|^2 \right] \tag{73}$$

Substituting (72)-(73) into (68) yields

$$\begin{aligned}
 \mathbb{E}[f(\bar{\mathbf{y}}^{(t+1)})] &\leq \mathbb{E}[f(\bar{\mathbf{y}}^{(t)})] + \frac{(1-\beta)L}{2\beta^3} \mathbb{E}[\|\bar{\mathbf{y}}^{(t)} - \bar{\mathbf{x}}^{(t)}\|^2] + \left(\frac{L\gamma^2\beta^3}{2(1-\beta)^3} - \frac{\gamma}{2(1-\beta)} \right) \mathbb{E} \left[\left\| \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \right\|^2 \right] \\
 &\quad - \frac{\gamma}{2(1-\beta)} \mathbb{E}[\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2] + \frac{\gamma L^2}{2(1-\beta)} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] + \frac{\gamma^2 L}{2(1-\beta)^2} \mathbb{E} \left[\left\| \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)} \right\|^2 \right] \tag{74}
 \end{aligned}$$

Dividing both sides by $\frac{\gamma}{2(1-\beta)}$ and rearranging terms yields

$$\begin{aligned}
 \mathbb{E}[\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2] &\leq \frac{2(1-\beta)}{\gamma} (\mathbb{E}[f(\bar{\mathbf{y}}^{(t)})] - \mathbb{E}[f(\bar{\mathbf{y}}^{(t+1)})]) - \left(1 - \frac{L\gamma\beta^3}{(1-\beta)^2}\right) \mathbb{E} \left[\left\| \frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)}) \right\|^2 \right] \\
 &\quad + \frac{(1-\beta)^2 L}{\gamma\beta^3} \mathbb{E}[\|\bar{\mathbf{y}}^{(t)} - \bar{\mathbf{x}}^{(t)}\|^2] + \frac{\gamma L}{(1-\beta)} \mathbb{E} \left[\left\| \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)} \right\|^2 \right] + L^2 \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] \tag{75}
 \end{aligned}$$

Summing over $t \in \{0, 1, \dots, T-1\}$

$$\begin{aligned}
 & \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2] \\
 & \leq \frac{2(1-\beta)}{\gamma} (\mathbb{E}[f(\bar{\mathbf{y}}^{(0)})] - \mathbb{E}[f(\bar{\mathbf{y}}^{(T)})]) - \left(1 - \frac{L\gamma\beta^3}{(1-\beta)^2}\right) \sum_{t=0}^{T-1} \mathbb{E}\left[\left\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})\right\|^2\right] \\
 & \quad + \frac{(1-\beta)^2 L}{\gamma\beta^3} \sum_{t=0}^{T-1} \mathbb{E}[\|\bar{\mathbf{y}}^{(t)} - \bar{\mathbf{x}}^{(t)}\|^2] + \frac{L\gamma}{(1-\beta)} \sum_{t=0}^{T-1} \mathbb{E}\left[\left\|\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)}\right\|^2\right] + L^2 \sum_{t=0}^{T-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] \\
 & \stackrel{(a)}{\leq} \frac{2(1-\beta)}{\gamma} (\mathbb{E}[f(\bar{\mathbf{y}}^{(0)})] - \mathbb{E}[f(\bar{\mathbf{y}}^{(T)})]) - \left(1 - \frac{L\gamma\beta^3}{(1-\beta)^2}\right) \sum_{t=0}^{T-1} \mathbb{E}\left[\left\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})\right\|^2\right] \\
 & \quad + \frac{L\gamma}{(1-\beta)^2} \sum_{t=0}^{T-1} \mathbb{E}\left[\left\|\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)}\right\|^2\right] + L^2 \sum_{t=0}^{T-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] \\
 & \stackrel{(b)}{\leq} \frac{2(1-\beta)}{\gamma} (\mathbb{E}[f(\bar{\mathbf{y}}^{(0)})] - \mathbb{E}[f(\bar{\mathbf{y}}^{(T)})]) - \left(1 - \frac{(1+\beta^3)L\gamma}{(1-\beta)^2}\right) \sum_{t=0}^{T-1} \mathbb{E}\left[\left\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})\right\|^2\right] + \frac{L\gamma}{(1-\beta)^2} \frac{\sigma^2}{N} T \\
 & \quad + \frac{1}{1 - \frac{12L^2\gamma^2 I^2}{(1-\beta)^2}} \frac{4L^2\gamma^2 I\sigma^2}{(1-\beta)^2} T + \frac{1}{1 - \frac{12L^2\gamma^2 I^2}{(1-\beta)^2}} \frac{6L^2\gamma^2 I^2 \kappa^2}{(1-\beta)^2} T \\
 & \stackrel{(c)}{\leq} \frac{2(1-\beta)}{\gamma} (f(\bar{\mathbf{x}}^{(0)}) - f^*) + \frac{L\gamma}{(1-\beta)^2} \frac{\sigma^2}{N} T + \frac{1}{1 - \frac{12L^2\gamma^2 I^2}{(1-\beta)^2}} \frac{4L^2\gamma^2 I\sigma^2}{(1-\beta)^2} T + \frac{1}{1 - \frac{12L^2\gamma^2 I^2}{(1-\beta)^2}} \frac{6L^2\gamma^2 I^2 \kappa^2}{(1-\beta)^2} T \quad (76)
 \end{aligned}$$

where (a) follows by using Lemma 7 and $\frac{\beta L\gamma}{(1-\beta)^2} + \frac{L\gamma}{1-\beta} = \frac{L\gamma}{(1-\beta)^2}$; (b) follows by applying Lemma 8 and by noting that $\mathbb{E}\left[\left\|\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)}\right\|^2\right] \leq \frac{1}{N}\sigma^2 + \mathbb{E}\left[\left\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})\right\|^2\right]$ by Lemma 1; and (c) follows because γ is chosen to ensure $1 - \frac{(1+\beta^3)L\gamma}{(1-\beta)^2} \geq 0$, $\bar{\mathbf{y}}^{(0)} = \bar{\mathbf{x}}^{(0)}$ by the definition in (49), and f^* is the minimum value of problem (1).

Dividing both sides by T yields

$$\begin{aligned}
 & \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2] \\
 & \leq \frac{2(1-\beta)}{\gamma T} (f(\bar{\mathbf{x}}^{(0)}) - f^*) + \frac{L\gamma}{(1-\beta)^2} \frac{\sigma^2}{N} + \frac{1}{1 - \frac{12L^2\gamma^2 I^2}{(1-\beta)^2}} \frac{4L^2\gamma^2 I\sigma^2}{(1-\beta)^2} + \frac{1}{1 - \frac{12L^2\gamma^2 I^2}{(1-\beta)^2}} \frac{6L^2\gamma^2 I^2 \kappa^2}{(1-\beta)^2} \\
 & \stackrel{(a)}{\leq} \frac{2(1-\beta)}{\gamma T} (f(\bar{\mathbf{x}}^{(0)}) - f^*) + \frac{L\gamma}{(1-\beta)^2} \frac{\sigma^2}{N} + \frac{3L^2\gamma^2 I\sigma^2}{(1-\beta)^2} + \frac{9L^2\gamma^2 I^2 \kappa^2}{(1-\beta)^2} \\
 & = O\left(\frac{1}{\gamma T}\right) + O\left(\frac{\gamma}{N}\sigma^2\right) + O(\gamma^2 I\sigma^2) + O(\gamma^2 I^2 \kappa^2) \quad (77)
 \end{aligned}$$

where (a) follows because $I \leq \frac{1-\beta}{6L\gamma}$ is chosen to ensure $\frac{1}{1 - \frac{12L^2\gamma^2 I^2}{(1-\beta)^2}} \leq \frac{3}{2}$.

6.6. Proof of Theorem 3

This section provides the complete proof for Algorithm 2 with Option I (Polyak's momentum). An extension to Algorithm 2 with Option II can be similarly done as our extension in Section 6.5 for Algorithm 1 from Option I to Option II.

Let $\bar{\mathbf{u}}^{(t)}$ and $\bar{\mathbf{x}}^{(t)}$ (using the forms in (9) and (7), respectively) be the N node averages of local variables $\mathbf{u}_i^{(t)}$ and $\mathbf{x}_i^{(t)}$ from Algorithm 2 with Option I. It is not difficult to show that $\bar{\mathbf{u}}^{(t)}$ and $\bar{\mathbf{x}}^{(t)}$ satisfy the following fact.

Fact 1. Let $\bar{\mathbf{u}}^{(t)} \triangleq \frac{1}{N} \sum_{i=1}^N \mathbf{u}_i^{(t)}$ and $\bar{\mathbf{x}}^{(t)} \triangleq \frac{1}{N} \sum_{i=1}^N \mathbf{x}_i^{(t)}$ be node averages of local variables from Algorithm 2 with

Option I. For all $t \geq 1$, we have

$$\begin{cases} \bar{\mathbf{u}}^{(t)} &= \beta \bar{\mathbf{u}}^{(t-1)} + \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t-1)} \\ \bar{\mathbf{x}}^{(t)} &= \bar{\mathbf{x}}^{(t-1)} - \gamma \bar{\mathbf{u}}^{(t)} \end{cases} \quad (78)$$

Proof. By the definition of $\bar{\mathbf{u}}^{(t)}$, we have

$$\begin{aligned} \bar{\mathbf{u}}^{(t)} &= \frac{1}{N} \sum_{i=1}^N \mathbf{u}_i^{(t)} \\ &\stackrel{(a)}{=} \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^N \tilde{\mathbf{u}}_j^{(t)} W_{ji} \\ &= \frac{1}{N} \sum_{j=1}^N \tilde{\mathbf{u}}_j^{(t)} \left[\sum_{i=1}^N W_{ji} \right] \\ &\stackrel{(b)}{=} \frac{1}{N} \sum_{i=1}^N \tilde{\mathbf{u}}_i^{(t)} \\ &\stackrel{(c)}{=} \beta \frac{1}{N} \sum_{i=1}^N \mathbf{u}_i^{(t-1)} + \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t-1)} \\ &\stackrel{(d)}{=} \beta \bar{\mathbf{u}}^{(t-1)} + \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t-1)} \end{aligned} \quad (79)$$

where (a) follows by substituting the first equation in (18); (b) follows by recalling that $\sum_{j=1}^N W_{ji} = 1$ for doubly stochastic matrix \mathbf{W} ; (c) follows by substituting the first equation in (16); and (d) follows from the definition of $\bar{\mathbf{u}}^{(t)}$.

Using a similar argument, we can show

$$\begin{aligned} \bar{\mathbf{x}}^{(t)} &\stackrel{(a)}{=} \frac{1}{N} \sum_{i=1}^N \mathbf{x}_i^{(t)} \\ &\stackrel{(b)}{=} \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^N \tilde{\mathbf{x}}_j^{(t)} W_{ji} \\ &= \frac{1}{N} \sum_{j=1}^N \tilde{\mathbf{x}}_j^{(t)} \left[\sum_{i=1}^N W_{ji} \right] \\ &\stackrel{(c)}{=} \frac{1}{N} \sum_{i=1}^N \tilde{\mathbf{x}}_i^{(t)} \\ &\stackrel{(d)}{=} \frac{1}{N} \sum_{i=1}^N [\mathbf{x}_i^{(t-1)} - \gamma \tilde{\mathbf{u}}_i^{(t-1)}] \\ &= \bar{\mathbf{x}}^{(t-1)} - \gamma \bar{\mathbf{u}}^{(t-1)} \end{aligned}$$

where (a) follows from the definition of $\bar{\mathbf{x}}^{(t)}$; (b) follows by substituting the second equation in (18); (c) follows by recalling that $\sum_{j=1}^N W_{ji} = 1$ for doubly stochastic matrix \mathbf{W} ; and (d) follows by using the definition of $\bar{\mathbf{x}}^{(t)}$ and equation (79). \square

It is remarkable that Fact 1 implies that even if we decentralized local averaging in Algorithm 2, the yielded global averages $\{\bar{\mathbf{u}}^{(t)}\}_{t \geq 0}$ and $\{\bar{\mathbf{x}}^{(t)}\}_{t \geq 0}$ follow the same dynamics as the global averages in Algorithm 1.

Let $\bar{\mathbf{x}}^{(t)} = \frac{1}{N} \sum_{i=1}^N \mathbf{x}_i^{(t)}$ be node averages of local variables $\mathbf{x}_i^{(t)}$ from Algorithm 2 with Option I. We again define the auxiliary sequence $\{\bar{\mathbf{z}}^{(t)}\}_{t \geq 0}$ via

$$\bar{\mathbf{z}}^{(t)} \triangleq \begin{cases} \bar{\mathbf{x}}^{(t)}, & t = 0 \\ \frac{1}{1-\beta} \bar{\mathbf{x}}^{(t)} - \frac{\beta}{1-\beta} \bar{\mathbf{x}}^{(t-1)}, & t \geq 1 \end{cases} \quad (80)$$

Since (78) and (80) are respectively identical to (10) and (19) for Algorithm 1 with Option I, the following two lemmas can be proven using exactly the same steps in the proofs for Lemmas 3 and 4.

Lemma 9. *Consider the sequence $\{\bar{\mathbf{z}}^{(t)}\}_{t \geq 0}$ defined in (80). Algorithm 2 with Option I ensures that for all $t \geq 0$, we have*

$$\bar{\mathbf{z}}^{(t+1)} - \bar{\mathbf{z}}^{(t)} = -\frac{\gamma}{1-\beta} \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)}. \quad (81)$$

Lemma 10. *Let $\{\bar{\mathbf{x}}^{(t)}\}_{t \geq 0}$ defined in (7) be the global node averages of local solutions from Algorithm 2 with Option I. Let $\{\bar{\mathbf{z}}^{(t)}\}_{t \geq 0}$ be defined in (80). For all $T \geq 1$, Algorithm 2 with Option I ensures that*

$$\sum_{t=0}^{T-1} \|\bar{\mathbf{z}}^{(t)} - \bar{\mathbf{x}}^{(t)}\|^2 \leq \frac{\gamma^2 \beta^2}{(1-\beta)^4} \sum_{t=0}^{T-1} \left\| \left[\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)} \right] \right\|^2 \quad (82)$$

To extend the proof in Section 6.3.2 (for Theorem 1) to analyze the convergence rate for Algorithm 2, the outstanding challenge is to provide a tight upper bound for quantity $\sum_{t=0}^{T-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2]$, i.e., to develop an counterpart of Lemma 5.

For each $t \geq 0$, let

$$\begin{cases} \mathbf{G}^{(t)} \triangleq [\mathbf{g}_1^{(t)}, \mathbf{g}_2^{(t)}, \dots, \mathbf{g}_N^{(t)}] \\ \mathbf{U}^{(t)} \triangleq [\mathbf{u}_1^{(t)}, \mathbf{u}_2^{(t)}, \dots, \mathbf{u}_N^{(t)}] \\ \tilde{\mathbf{U}}^{(t)} \triangleq [\tilde{\mathbf{u}}_1^{(t)}, \tilde{\mathbf{u}}_2^{(t)}, \dots, \tilde{\mathbf{u}}_N^{(t)}] \\ \mathbf{X}^{(t)} \triangleq [\mathbf{x}_1^{(t)}, \mathbf{x}_2^{(t)}, \dots, \mathbf{x}_N^{(t)}] \end{cases} \quad (83)$$

be $m \times N$ matrices that concatenate local variables $\mathbf{g}_i^{(t)}$, $\mathbf{u}_i^{(t)}$, $\tilde{\mathbf{u}}_i^{(t)}$ and $\mathbf{x}_i^{(t)}$ for all N nodes. Recall that the Frobenius norm for any $m \times N$ matrix \mathbf{Z} satisfies $\|\mathbf{Z}\|_F^2 = \sum_{i=1}^N \|\mathbf{z}_i\|^2$ where $\|\mathbf{z}_i\|^2$ is the Euclidean norm of the i -th column of matrix \mathbf{Z} . It can be easily verified that $\sum_{i=1}^N \|\mathbf{x}_i^{(t)} - \bar{\mathbf{x}}^{(t)}\|^2 = \|\mathbf{X}^{(t)}(\mathbf{I} - \mathbf{Q})\|_F^2$ where

$$\mathbf{Q} \triangleq \frac{1}{N} \mathbf{1}_N \mathbf{1}_N^\top \quad (84)$$

is an $N \times N$ matrix where all entries are $1/N$ and \mathbf{I} is the identity matrix for which the dimensions are obvious in its context.

Due to the asymmetry in computing local averages for different nodes in (18), it is easier to provide the desired counterpart of Lemma 5 by studying the equivalent matrix form $\|\mathbf{X}^{(t)}(\mathbf{I} - \mathbf{Q})\|_F^2$. The technique of introducing matrix forms to analyze the convergence rate has been previously used in (Lian et al., 2017; Tang et al., 2018; Wang & Joshi, 2018a) for the analysis of decentralized SGD without momentum.

The following useful facts are simple in matrix analysis (Horn & Johnson, 1985). Similar facts are used in (Lian et al., 2017; Wang & Joshi, 2018a).

Fact 2. *Let $\mathbf{A}_i, i \in \{1, 2, \dots, n\}$ be n arbitrary real square matrices. It follows that $\|\sum_{i=1}^n \mathbf{A}_i\|_F^2 \leq \sum_{i=1}^n \sum_{j=1}^n \|\mathbf{A}_i\|_F \|\mathbf{A}_j\|_F$.*

Proof. Recall by the definition of Frobenius norm, we have $\|\mathbf{A}\|_F^2 = \text{tr}(\mathbf{A}^\top \mathbf{A})$ where $\text{tr}(\cdot)$ denote the trace operator of a

square matrix. It follows that

$$\begin{aligned} \left\| \sum_{i=1}^n \mathbf{A}_i \right\|_F^2 &= \text{tr} \left(\left(\sum_{i=1}^n \mathbf{A}_i \right)^\top \left(\sum_{j=1}^n \mathbf{A}_j \right) \right) \\ &= \sum_{i=1}^n \sum_{j=1}^n \text{tr}(\mathbf{A}_i^\top \mathbf{A}_j) \\ &\stackrel{(a)}{\leq} \sum_{i=1}^n \sum_{j=1}^n \|\mathbf{A}_i\|_F \|\mathbf{A}_j\|_F \end{aligned}$$

where (a) follows from the simple fact that $|\text{tr}(\mathbf{A}^\top \mathbf{B})| \leq \|\mathbf{A}\|_F \|\mathbf{B}\|_F$ for any two square matrices. \square

Fact 3. Let \mathbf{Q} be defined in (84). For any symmetric doubly stochastic matrix \mathbf{W} under Assumption 2, we have

1. $\mathbf{QW} = \mathbf{WQ}$
2. $(\mathbf{I} - \mathbf{Q})\mathbf{W} = \mathbf{W}(\mathbf{I} - \mathbf{Q})$
3. For any integer $k \geq 1$, we have $\|(\mathbf{I} - \mathbf{Q})\mathbf{W}^k\| \leq \rho^{k/2}$ where $\rho < 1$ is the constant in Assumption 2 and $\|\cdot\|$ denotes the spectrum norm for a matrix.

Proof. Recall that \mathbf{W} is a symmetric matrix. Let $\mathbf{W} = \mathbf{P}^\top \mathbf{\Lambda} \mathbf{P}$ be the eigenvalue decomposition for \mathbf{W} where $\mathbf{\Lambda} = \text{Diag}\{\lambda_1(\mathbf{W}), \dots, \lambda_N(\mathbf{W})\}$ and \mathbf{P} is a unitary matrix where each column is an eigenvector of \mathbf{W} . Since $\frac{1}{\sqrt{N}}\mathbf{1}_N$ is the eigenvector corresponding to eigenvalue $\lambda_1(\mathbf{W}) = 1$ for matrix \mathbf{W} , thus the first column of \mathbf{P} is $\frac{1}{\sqrt{N}}\mathbf{1}_N$. By (84), we further have an eigenvalue decomposition of \mathbf{Q} given by $\mathbf{Q} = \mathbf{P}\mathbf{\Gamma}\mathbf{P}^\top$ where $\mathbf{\Gamma} = \text{Diag}\{1, 0, \dots, 0\}$. Thus, we have

1. $\mathbf{QW} = \mathbf{P}^\top \mathbf{\Gamma} \mathbf{P} \mathbf{P}^\top \mathbf{\Lambda} \mathbf{P} = \mathbf{P}^\top \mathbf{\Gamma} \mathbf{\Lambda} \mathbf{P} = \mathbf{P}^\top \mathbf{\Lambda} \mathbf{\Gamma} \mathbf{P} = \mathbf{P}^\top \mathbf{\Lambda} \mathbf{P} \mathbf{P}^\top \mathbf{\Gamma} \mathbf{P} = \mathbf{WQ}$
2. $(\mathbf{I} - \mathbf{Q})\mathbf{W} = \mathbf{W} - \mathbf{QW} \stackrel{(a)}{=} \mathbf{W} - \mathbf{WQ} = \mathbf{W}(\mathbf{I} - \mathbf{Q})$ where (a) follows from part (1) of this fact.
3. Note that $(\mathbf{I} - \mathbf{Q})\mathbf{W}^k = \mathbf{P}^\top (\mathbf{I} - \mathbf{\Gamma}) \mathbf{P} (\mathbf{P}^\top \mathbf{\Lambda} \mathbf{P})^k = \mathbf{P}^\top (\mathbf{I} - \mathbf{\Gamma}) \mathbf{\Lambda}^k \mathbf{P}$ where $(\mathbf{I} - \mathbf{\Gamma}) \mathbf{\Lambda}^k = \text{Diag}\{0, (\lambda_2(\mathbf{W}))^k, \dots, (\lambda_N(\mathbf{W}))^k\}$. By Assumption 2, we have $\max\{|\lambda_2(\mathbf{W})|^k, |\lambda_N(\mathbf{W})|^k\} \leq (\sqrt{\rho})^k$.

\square

Lemma 11. Consider problem (1) under Assumptions 1 and 2. Let $\{\bar{\mathbf{x}}^{(t)}\}_{t \geq 0}$ defined in (7) be the node averages of local solutions from Algorithm 2 with Option I. If γ is chosen to satisfy $\frac{16L^2\gamma^2}{(1-\beta)^2(1-\sqrt{\rho})^2} < 1$, then for all $T \geq 1$, we have

$$\begin{aligned} & \sum_{t=0}^{T-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] \\ & \leq \frac{1}{1 - \frac{16L^2\gamma^2}{(1-\beta)^2(1-\sqrt{\rho})^2}} \frac{2\gamma^2\sigma^2}{(1-\beta)^2(1-\rho)} T + \frac{1}{1 - \frac{16L^2\gamma^2}{(1-\beta)^2(1-\sqrt{\rho})^2}} \frac{8\gamma^2}{(1-\beta)^2(1-\sqrt{\rho})^2} \sum_{t=0}^{T-1} \mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})\|^2] \\ & \quad + \frac{1}{1 - \frac{16L^2\gamma^2}{(1-\beta)^2(1-\sqrt{\rho})^2}} \frac{2\gamma^2\kappa^2}{(1-\beta)^2(1-\sqrt{\rho})^2} T \end{aligned}$$

Proof. Recall the definitions of $\mathbf{U}^{(t)}$ and $\tilde{\mathbf{U}}^{(t)}$ in (83), for all $t \geq 1$, we have

$$\begin{aligned} \mathbf{U}^{(t)} &\stackrel{(a)}{=} \tilde{\mathbf{U}}^{(t)} \mathbf{W} \\ &\stackrel{(b)}{=} [\beta \mathbf{U}^{(t-1)} + \mathbf{G}^{(t-1)}] \mathbf{W} \end{aligned} \tag{85}$$

where (a) follows from the first equation in (18) and (b) follows from the first equation in (16).

Recursively applying the above equation for t times yields

$$\begin{aligned} \mathbf{U}^{(t)} &= \beta^t \mathbf{U}^{(0)} \mathbf{W}^t + \sum_{\tau=0}^{t-1} \mathbf{G}^{(\tau)} (\beta \mathbf{W})^{t-1-\tau} \\ &\stackrel{(a)}{=} \sum_{\tau=0}^{t-1} \mathbf{G}^{(\tau)} (\beta \mathbf{W})^{t-1-\tau} \end{aligned} \quad (86)$$

where (a) follows from $\mathbf{U}^{(0)} = \mathbf{0}$.

Recall the definitions of $\mathbf{X}^{(t)}$ and $\tilde{\mathbf{U}}^{(t)}$ in (83), for all $t \geq 1$, we have

$$\begin{aligned} \mathbf{X}^{(t)} &\stackrel{(a)}{=} [\mathbf{X}^{(t-1)} - \gamma \tilde{\mathbf{U}}^{(t)}] \mathbf{W} \\ &\stackrel{(b)}{=} [\mathbf{X}^{(t-1)} - \gamma [\beta \mathbf{U}^{(t-1)} + \mathbf{G}^{(t-1)}]] \mathbf{W} \\ &= \mathbf{X}^{(t-1)} \mathbf{W} - \gamma [\beta \mathbf{U}^{(t-1)} + \mathbf{G}^{(t-1)}] \mathbf{W} \\ &\stackrel{(c)}{=} \mathbf{X}^{(t-1)} \mathbf{W} - \gamma \mathbf{U}^{(t)} \end{aligned} \quad (87)$$

where (a) follows from the second equation in (16); (b) follows from the first equation in (16); and (c) follows from (85).

Multiplying both sides by $\mathbf{I} - \mathbf{Q}$ with \mathbf{Q} defined in (84) yields

$$\begin{aligned} \mathbf{X}^{(t)} (\mathbf{I} - \mathbf{Q}) &= \mathbf{X}^{(t-1)} \mathbf{W} (\mathbf{I} - \mathbf{Q}) - \gamma \mathbf{U}^{(t)} (\mathbf{I} - \mathbf{Q}) \\ &\stackrel{(a)}{=} \mathbf{X}^{(t-1)} (\mathbf{I} - \mathbf{Q}) \mathbf{W} - \gamma \mathbf{U}^{(t)} (\mathbf{I} - \mathbf{Q}) \end{aligned}$$

where (a) follows because $\mathbf{W}(\mathbf{I} - \mathbf{Q}) = (\mathbf{I} - \mathbf{Q})\mathbf{W}$ by Fact 3.

For all $t \geq 1$, recursively applying the above equation for t times (using $\mathbf{W}(\mathbf{I} - \mathbf{Q}) = (\mathbf{I} - \mathbf{Q})\mathbf{W}$ when needed) yields

$$\begin{aligned} \mathbf{X}^{(t)} (\mathbf{I} - \mathbf{Q}) &= \mathbf{X}^{(0)} (\mathbf{I} - \mathbf{Q}) \mathbf{W}^t - \gamma \sum_{\tau=1}^t \mathbf{U}^{(\tau)} (\mathbf{I} - \mathbf{Q}) \mathbf{W}^{t-\tau} \\ &\stackrel{(a)}{=} -\gamma \sum_{\tau=1}^t \mathbf{U}^{(\tau)} (\mathbf{I} - \mathbf{Q}) \mathbf{W}^{t-\tau} \\ &\stackrel{(b)}{=} -\gamma \sum_{\tau=1}^t \sum_{j=0}^{\tau-1} \mathbf{G}^{(j)} (\beta \mathbf{W})^{\tau-1-j} (\mathbf{I} - \mathbf{Q}) \mathbf{W}^{t-\tau} \\ &\stackrel{(c)}{=} -\gamma \sum_{\tau=1}^t \sum_{j=0}^{\tau-1} \mathbf{G}^{(j)} \beta^{\tau-1-j} \mathbf{W}^{t-1-j} (\mathbf{I} - \mathbf{Q}) \\ &= -\gamma \sum_{k=0}^{t-1} \mathbf{G}^{(k)} \left[\sum_{j=k+1}^t \beta^{j-1-k} \right] \mathbf{W}^{t-1-k} (\mathbf{I} - \mathbf{Q}) \\ &= -\gamma \sum_{\tau=0}^{t-1} \frac{1 - \beta^{t-\tau}}{1 - \beta} \mathbf{G}^{(\tau)} (\mathbf{I} - \mathbf{Q}) \mathbf{W}^{t-1-\tau} \end{aligned} \quad (88)$$

where (a) follows because $\mathbf{X}^{(0)} (\mathbf{I} - \mathbf{Q}) = \mathbf{0}$ due to all columns of $\mathbf{X}^{(0)}$ are identical; (b) follows by substituting (86); (c) follows because $\mathbf{W}(\mathbf{I} - \mathbf{Q}) = (\mathbf{I} - \mathbf{Q})\mathbf{W}$ by Fact 3.

For all $t \geq 0$, denote

$$\mathbf{H}^{(t)} \triangleq [\nabla f_1(\mathbf{x}_1^{(t)}), \nabla f_2(\mathbf{x}_2^{(t)}), \dots, \nabla f_N(\mathbf{x}_N^{(t)})] \quad (89)$$

From (88), for all $t \geq 1$, we have

$$\begin{aligned}
 & \mathbb{E}[\|\mathbf{X}^{(t)}(\mathbf{I} - \mathbf{Q})\|_F^2] \\
 &= \gamma^2 \mathbb{E}[\|\sum_{\tau=0}^{t-1} \frac{1 - \beta^{t-\tau}}{1 - \beta} \mathbf{G}^{(\tau)}(\mathbf{I} - \mathbf{Q})\mathbf{W}^{t-1-\tau}\|_F^2] \\
 &\stackrel{(a)}{\leq} 2\gamma^2 \mathbb{E}[\|\sum_{\tau=0}^{t-1} \frac{1 - \beta^{t-\tau}}{1 - \beta} (\mathbf{G}^{(\tau)} - \mathbf{H}^{(\tau)})(\mathbf{I} - \mathbf{Q})\mathbf{W}^{t-1-\tau}\|_F^2] + 2\gamma^2 \mathbb{E}[\|\sum_{\tau=0}^{t-1} \frac{1 - \beta^{t-\tau}}{1 - \beta} \mathbf{H}^{(\tau)}(\mathbf{I} - \mathbf{Q})\mathbf{W}^{t-1-\tau}\|_F^2] \quad (90)
 \end{aligned}$$

where (a) follows from the basic inequality $\|\mathbf{A}_1 + \mathbf{A}_2\|_F^2 \leq 2\|\mathbf{A}_1\|_F^2 + 2\|\mathbf{A}_2\|_F^2$.

Now we develop the respective upper bounds of the two terms on the right side of (90). We note that

$$\begin{aligned}
 & \mathbb{E}[\|\sum_{\tau=0}^{t-1} \frac{1 - \beta^{t-\tau}}{1 - \beta} (\mathbf{G}^{(\tau)} - \mathbf{H}^{(\tau)})(\mathbf{I} - \mathbf{Q})\mathbf{W}^{t-1-\tau}\|_F^2] \\
 &\stackrel{(a)}{=} \sum_{\tau=0}^{t-1} \mathbb{E}[\|\frac{1 - \beta^{t-\tau}}{1 - \beta} (\mathbf{G}^{(\tau)} - \mathbf{H}^{(\tau)})(\mathbf{I} - \mathbf{Q})\mathbf{W}^{t-1-\tau}\|_F^2] \\
 &\stackrel{(b)}{\leq} \frac{1}{(1 - \beta)^2} \sum_{\tau=0}^{t-1} \rho^{t-1-\tau} \mathbb{E}[\|\mathbf{G}^{(\tau)} - \mathbf{H}^{(\tau)}\|_F^2] \\
 &\stackrel{(c)}{\leq} \frac{1}{(1 - \beta)^2} \sum_{\tau=0}^{t-1} \rho^{t-1-\tau} N\sigma^2 \\
 &\stackrel{(d)}{\leq} \frac{N\sigma^2}{(1 - \beta)^2(1 - \rho)} \quad (91)
 \end{aligned}$$

where (a) follows because $\mathbb{E}[\mathbf{G}^{(\tau_2)} - \mathbf{H}^{(\tau_2)} | \mathbf{G}^{(\tau_1)} - \mathbf{H}^{(\tau_1)}] = \mathbf{0}$ for any $\tau_2 > \tau_1$; (b) follows because $|\frac{1 - \beta^{t-\tau}}{1 - \beta}| \leq \frac{1}{1 - \beta}$, $\|\mathbf{A}\mathbf{B}\|_F \leq \|\mathbf{A}\|_F \|\mathbf{B}\|_F$ for any two matrices \mathbf{A}, \mathbf{B} and $\|(\mathbf{I} - \mathbf{Q})\mathbf{W}^{t-1-\tau}\| \leq \rho^{(t-1-\tau)/2}$ by Fact 3; and (c) follows by noting $\mathbb{E}[\|\mathbf{G}^{(\tau)} - \mathbf{H}^{(\tau)}\|_F^2] = \sum_{i=1}^N \mathbb{E}[\|\mathbf{g}_i^{(\tau)} - \nabla f_i(\mathbf{x}_i^{(\tau)})\|^2] \leq N\sigma^2$ where the last inequality follows by Assumption 1; and (d) follows because $\sum_{\tau=0}^{t-1} \rho^{t-1-\tau} = \frac{1 - \rho^t}{1 - \rho} \leq \frac{1}{1 - \rho}$.

We also note that

$$\begin{aligned}
 & \mathbb{E}[\|\sum_{\tau=0}^{t-1} \frac{1 - \beta^{t-\tau}}{1 - \beta} \mathbf{H}^{(\tau)}(\mathbf{I} - \mathbf{Q})\mathbf{W}^{t-1-\tau}\|_F^2] \\
 &\stackrel{(a)}{\leq} \sum_{\tau=0}^{t-1} \sum_{\tau'=0}^{t-1} \mathbb{E}[\|\frac{1 - \beta^{t-\tau}}{1 - \beta} \mathbf{H}^{(\tau)}(\mathbf{I} - \mathbf{Q})\mathbf{W}^{t-1-\tau}\|_F \|\frac{1 - \beta^{t-\tau'}}{1 - \beta} \mathbf{H}^{(\tau')}(\mathbf{I} - \mathbf{Q})\mathbf{W}^{t-1-\tau'}\|_F] \\
 &\stackrel{(b)}{\leq} \frac{1}{(1 - \beta)^2} \sum_{\tau=0}^{t-1} \sum_{\tau'=0}^{t-1} \rho^{t-1-(\tau+\tau')/2} \mathbb{E}[\|\mathbf{H}^{(\tau)}\|_F \|\mathbf{H}^{(\tau')}\|_F] \\
 &\stackrel{(c)}{\leq} \frac{1}{(1 - \beta)^2} \sum_{\tau=0}^{t-1} \sum_{\tau'=0}^{t-1} \rho^{t-1-(\tau+\tau')/2} (\frac{1}{2} \mathbb{E}[\|\mathbf{H}^{(\tau)}\|_F^2] + \frac{1}{2} \mathbb{E}[\|\mathbf{H}^{(\tau')}\|_F^2]) \\
 &\stackrel{(d)}{=} \frac{1}{(1 - \beta)^2} \sum_{\tau=0}^{t-1} \sum_{\tau'=0}^{t-1} \rho^{t-1-(\tau+\tau')/2} \mathbb{E}[\|\mathbf{H}^{(\tau)}\|_F^2] \\
 &\stackrel{(e)}{\leq} \frac{1}{(1 - \beta)^2(1 - \sqrt{\rho})} \sum_{\tau=0}^{t-1} \rho^{(t-1-\tau)/2} \mathbb{E}[\|\mathbf{H}^{(\tau)}\|_F^2] \quad (92)
 \end{aligned}$$

where (a) follows from Fact 2; (b) follows because $|\frac{1 - \beta^{t-\tau}}{1 - \beta}| \leq \frac{1}{1 - \beta}$, $\|\mathbf{A}\mathbf{B}\|_F \leq \|\mathbf{A}\|_F \|\mathbf{B}\|_F$ for any two matrices \mathbf{A}, \mathbf{B} and $\|(\mathbf{I} - \mathbf{Q})\mathbf{W}^{t-1-\tau}\| \leq \rho^{(t-1-\tau)/2}$ by Fact 3; (c) follows from the basic identity $xy \leq \frac{1}{2}x^2 + \frac{1}{2}y^2$ for any two real

number x, y ; (d) follows by noting the symmetry between τ and τ' in the expression from last line; (e) follows because $\sum_{\tau'=0}^{t-1} \rho^{t-1-(\tau+\tau')/2} \leq \frac{\rho^{(t-1-\tau)/2}}{1-\sqrt{\rho}}$.

For all $t \geq 0$, we denote

$$\mathbf{J}^{(t)} \triangleq [\nabla f_1(\bar{\mathbf{x}}^{(t)}), \nabla f_2(\bar{\mathbf{x}}^{(t)}), \dots, \nabla f_N(\bar{\mathbf{x}}^{(t)})] \quad (93)$$

Then, for all $\tau \geq 0$, we have

$$\begin{aligned} & \mathbb{E}[\|\mathbf{H}^{(\tau)}\|_F^2] \\ &= \mathbb{E}[\|\mathbf{H}^{(\tau)} - \mathbf{J}^{(\tau)} + \mathbf{J}^{(\tau)} - \mathbf{J}^{(\tau)}\mathbf{Q} + \mathbf{J}^{(\tau)}\mathbf{Q} - \mathbf{H}^{(\tau)}\mathbf{Q} + \mathbf{H}^{(\tau)}\mathbf{Q}\|_F^2] \\ &\stackrel{(a)}{\leq} 4\mathbb{E}[\|\mathbf{H}^{(\tau)} - \mathbf{J}^{(\tau)}\|_F^2] + 4\mathbb{E}[\|\mathbf{J}^{(\tau)} - \mathbf{J}^{(\tau)}\mathbf{Q}\|_F^2] + 4\mathbb{E}[\|(\mathbf{J}^{(\tau)} - \mathbf{H}^{(\tau)})\mathbf{Q}\|_F^2] + 4\mathbb{E}[\|\mathbf{H}^{(\tau)}\mathbf{Q}\|_F^2] \\ &\stackrel{(b)}{\leq} 4L^2\mathbb{E}[\|\mathbf{X}^{(\tau)}(\mathbf{I} - \mathbf{Q})\|_F^2] + 4N\kappa^2 + 4L^2\mathbb{E}[\|\mathbf{X}^{(\tau)}(\mathbf{I} - \mathbf{Q})\|_F^2] + 4N\mathbb{E}[\|\frac{1}{N}\sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(\tau)})\|^2] \\ &= 8L^2\mathbb{E}[\|\mathbf{X}^{(\tau)}(\mathbf{I} - \mathbf{Q})\|_F^2] + 4N\kappa^2 + 4N\mathbb{E}[\|\frac{1}{N}\sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(\tau)})\|^2] \end{aligned} \quad (94)$$

where (a) follows by applying the basic inequality $\|\sum_{i=1}^n \mathbf{A}_i\|_F^2 \leq n \sum_{i=1}^n \|\mathbf{A}_i\|_F^2$ for matrices of the same dimension with $n = 4$; (b) follows because $\|\mathbf{H}^{(\tau)} - \mathbf{J}^{(\tau)}\|_F^2 = \sum_{i=1}^N \|\nabla f_i(\mathbf{x}_i^{(\tau)}) - \nabla f_i(\bar{\mathbf{x}}^{(\tau)})\|^2 \leq L^2 \sum_{i=1}^N \|\mathbf{x}_i^{(\tau)} - \bar{\mathbf{x}}^{(\tau)}\|^2 = L^2 \|\mathbf{X}^{(\tau)}(\mathbf{I} - \mathbf{Q})\|_F^2$ where the inequality step follows from the smoothness of each $f_i(\cdot)$ by Assumption 1, $\|\mathbf{J}^{(\tau)} - \mathbf{J}^{(\tau)}\mathbf{Q}\|_F^2 = \sum_{i=1}^N \|\nabla f_i(\bar{\mathbf{x}}^{(\tau)}) - \nabla f(\bar{\mathbf{x}}^{(\tau)})\|^2 \leq N\kappa^2$ where the second step follows from Assumption 1, $\|(\mathbf{J}^{(\tau)} - \mathbf{H}^{(\tau)})\mathbf{Q}\|_F^2 \leq \|\mathbf{H}^{(\tau)} - \mathbf{J}^{(\tau)}\|_F^2$.

Substituting (94) into (92) yields

$$\begin{aligned} & \mathbb{E}[\|\sum_{\tau=0}^{t-1} \frac{1-\beta^{t-\tau}}{1-\beta} \mathbf{H}^{(\tau)}(\mathbf{I} - \mathbf{Q})\mathbf{W}^{t-1-\tau}\|_F^2] \\ &\leq \frac{1}{(1-\beta)^2(1-\sqrt{\rho})} \left[\left(8L^2 \sum_{\tau=0}^{t-1} \rho^{(t-1-\tau)/2} \mathbb{E}[\|\mathbf{X}^{(\tau)}(\mathbf{I} - \mathbf{Q})\|_F^2] \right) + \left(4N \sum_{\tau=0}^{t-1} \rho^{(t-1-\tau)/2} \mathbb{E}[\|\frac{1}{N}\sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(\tau)})\|^2] \right) \right. \\ &\quad \left. + N\kappa^2 \sum_{\tau=0}^{t-1} \rho^{(t-1-\tau)/2} \right] \\ &\leq \frac{1}{(1-\beta)^2(1-\sqrt{\rho})} \left[\left(8L^2 \sum_{\tau=0}^{t-1} \rho^{(t-1-\tau)/2} \mathbb{E}[\|\mathbf{X}^{(\tau)}(\mathbf{I} - \mathbf{Q})\|_F^2] \right) + \left(4N \sum_{\tau=0}^{t-1} \rho^{(t-1-\tau)/2} \mathbb{E}[\|\frac{1}{N}\sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(\tau)})\|^2] \right) \right. \\ &\quad \left. + N\kappa^2 \frac{1}{1-\sqrt{\rho}} \right] \end{aligned} \quad (95)$$

Substituting (91) and (95) into (90) yields

$$\begin{aligned} & \mathbb{E}[\|\mathbf{X}^{(t)}(\mathbf{I} - \mathbf{Q})\|_F^2] \\ &\leq \frac{2N\gamma^2\sigma^2}{(1-\beta)^2(1-\rho)} + \frac{2\gamma^2}{(1-\beta)^2(1-\sqrt{\rho})} \left[\left(8L^2 \sum_{\tau=0}^{t-1} \rho^{(t-1-\tau)/2} \mathbb{E}[\|\mathbf{X}^{(\tau)}(\mathbf{I} - \mathbf{Q})\|_F^2] \right) + \left(4N \sum_{\tau=0}^{t-1} \rho^{(t-1-\tau)/2} \mathbb{E}[\|\frac{1}{N}\sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(\tau)})\|^2] \right) \right] \\ &\quad + \frac{2N\gamma^2\kappa^2}{(1-\beta)^2(1-\sqrt{\rho})^2} \end{aligned}$$

Summing over $t \in \{1, 2, \dots, T-1\}$ and noting that $\|\mathbf{X}^{(0)}(\mathbf{I} - \mathbf{Q})\|_F^2 = 0$ yields

$$\begin{aligned}
& \sum_{t=0}^{T-1} \mathbb{E}[\|\mathbf{X}^{(t)}(\mathbf{I} - \mathbf{Q})\|_F^2] \\
& \leq \frac{2N\gamma^2\sigma^2}{(1-\beta)^2(1-\rho)}T + \frac{16L^2\gamma^2}{(1-\beta)^2(1-\sqrt{\rho})} \sum_{t=1}^{T-1} \sum_{\tau=0}^{t-1} \rho^{(t-1-\tau)/2} \mathbb{E}[\|\mathbf{X}^{(\tau)}(\mathbf{I} - \mathbf{Q})\|_F^2] \\
& \quad + \frac{8N\gamma^2}{(1-\beta)^2(1-\sqrt{\rho})} \sum_{t=1}^{T-1} \sum_{\tau=0}^{t-1} \rho^{(t-1-\tau)/2} \mathbb{E}\left[\left\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(\tau)})\right\|^2\right] + \frac{2N\gamma^2\kappa^2}{(1-\beta)^2(1-\sqrt{\rho})^2}T \\
& \stackrel{(a)}{\leq} \frac{2N\gamma^2\sigma^2}{(1-\beta)^2(1-\rho)}T + \frac{16L^2\gamma^2}{(1-\beta)^2(1-\sqrt{\rho})} \sum_{t=0}^{T-1} \frac{1-\rho^{(T-1-t)/2}}{1-\sqrt{\rho}} \mathbb{E}[\|\mathbf{X}^{(t)}(\mathbf{I} - \mathbf{Q})\|_F^2] \\
& \quad + \frac{8N\gamma^2}{(1-\beta)^2(1-\sqrt{\rho})} \sum_{t=0}^{T-1} \frac{1-\rho^{(T-1-t)/2}}{1-\sqrt{\rho}} \mathbb{E}\left[\left\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})\right\|^2\right] + \frac{2N\gamma^2\kappa^2}{(1-\beta)^2(1-\sqrt{\rho})^2}T \\
& \leq \frac{2N\gamma^2\sigma^2}{(1-\beta)^2(1-\rho)}T + \frac{16L^2\gamma^2}{(1-\beta)^2(1-\sqrt{\rho})^2} \sum_{t=0}^{T-1} \mathbb{E}[\|\mathbf{X}^{(t)}(\mathbf{I} - \mathbf{Q})\|_F^2] \\
& \quad + \frac{8N\gamma^2}{(1-\beta)^2(1-\sqrt{\rho})^2} \sum_{t=0}^{T-1} \mathbb{E}\left[\left\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})\right\|^2\right] + \frac{2N\gamma^2\kappa^2}{(1-\beta)^2(1-\sqrt{\rho})^2}T
\end{aligned}$$

where (a) follows after simplifying the double summations (by swapping the order of two summations and collecting coefficients for common terms)

Rearranging terms and dividing both sides by $N(1 - \frac{16L^2\gamma^2}{(1-\beta)^2(1-\sqrt{\rho})^2})$ yields

$$\begin{aligned}
& \sum_{t=0}^{T-1} \frac{1}{N} \mathbb{E}[\|\mathbf{X}^{(t)}(\mathbf{I} - \mathbf{Q})\|_F^2] \\
& \leq \frac{1}{1 - \frac{16L^2\gamma^2}{(1-\beta)^2(1-\sqrt{\rho})^2}} \frac{2\gamma^2\sigma^2}{(1-\beta)^2(1-\rho)}T + \frac{1}{1 - \frac{16L^2\gamma^2}{(1-\beta)^2(1-\sqrt{\rho})^2}} \frac{8\gamma^2}{(1-\beta)^2(1-\sqrt{\rho})^2} \sum_{t=0}^{T-1} \mathbb{E}\left[\left\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})\right\|^2\right] \\
& \quad + \frac{1}{1 - \frac{16L^2\gamma^2}{(1-\beta)^2(1-\sqrt{\rho})^2}} \frac{2\gamma^2\kappa^2}{(1-\beta)^2(1-\sqrt{\rho})^2}T
\end{aligned}$$

Finally, this lemma follows by recalling that $\|\mathbf{X}^{(t)}(\mathbf{I} - \mathbf{Q})\|_F^2 = \sum_{i=1}^N \|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2$. □

Remark 5. Comparing Lemma 11 with Lemma 5, we note that the upper bound in Lemma 11 includes an additional term involving $\mathbb{E}\left[\left\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})\right\|^2\right]$, which is the cost of asymmetry in local averages.

To simplify the lengthy coefficient $\frac{1}{1 - \frac{16L^2\gamma^2}{(1-\beta)^2(1-\sqrt{\rho})^2}}$ in Lemma 11, we have the following trivial corollary.

Corollary 4. Consider problem (1) under Assumptions 1 and 2. Let $\{\bar{\mathbf{x}}^{(t)}\}_{t \geq 0}$ defined in (7) be the node averages of local solutions from Algorithm 2 with Option I. If $\gamma \leq \frac{(1-\beta)(1-\sqrt{\rho})}{4\sqrt{2}L}$, then for all $T \geq 1$, we have

$$\begin{aligned}
& \sum_{t=0}^{T-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] \\
& \leq \frac{4\gamma^2\sigma^2}{(1-\beta)^2(1-\rho)}T + \frac{16\gamma^2}{(1-\beta)^2(1-\sqrt{\rho})^2} \sum_{t=0}^{T-1} \mathbb{E}\left[\left\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})\right\|^2\right] + \frac{4\gamma^2\kappa^2}{(1-\beta)^2(1-\sqrt{\rho})^2}T
\end{aligned}$$

Proof. It follows because $\frac{1}{1 - \frac{16L^2\gamma^2}{(1-\beta)^2(1-\sqrt{\rho})^2}} < 2$ under our choice of γ . □

6.6.1. MAIN PROOF OF THEOREM 3

Now, we are ready to present the main proof of Theorem 3. Since Lemmas 9 and 10 under Algorithm 2 are respectively identical to Lemmas 3 and 4, repeating the steps from (36) to (41) in Section 6.3.2 yields:

$$\begin{aligned} \mathbb{E}[\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2] &\leq \frac{2(1-\beta)}{\gamma} (\mathbb{E}[f(\bar{\mathbf{z}}^{(t)})] - \mathbb{E}[f(\bar{\mathbf{z}}^{(t+1)})]) - \left(1 - \frac{L\gamma\beta}{(1-\beta)^2}\right) \mathbb{E}\left[\left\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})\right\|^2\right] \\ &\quad + \frac{(1-\beta)^2 L}{\beta\gamma} \mathbb{E}[\|\bar{\mathbf{z}}^{(t)} - \bar{\mathbf{x}}^{(t)}\|^2] + \frac{L\gamma}{(1-\beta)} \mathbb{E}\left[\left\|\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)}\right\|^2\right] + L^2 \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] \end{aligned} \quad (96)$$

Summing over $t \in \{0, 1, \dots, T-1\}$ yields

$$\begin{aligned} &\sum_{t=0}^{T-1} \mathbb{E}[\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2] \\ &\leq \frac{2(1-\beta)}{\gamma} (\mathbb{E}[f(\bar{\mathbf{z}}^{(0)})] - \mathbb{E}[f(\bar{\mathbf{z}}^{(T)})]) - \left(1 - \frac{L\gamma\beta}{(1-\beta)^2}\right) \sum_{t=0}^{T-1} \mathbb{E}\left[\left\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})\right\|^2\right] \\ &\quad + \frac{(1-\beta)^2 L}{\beta\gamma} \sum_{t=0}^{T-1} \mathbb{E}[\|\bar{\mathbf{z}}^{(t)} - \bar{\mathbf{x}}^{(t)}\|^2] + \frac{L\gamma}{(1-\beta)} \sum_{t=0}^{T-1} \mathbb{E}\left[\left\|\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)}\right\|^2\right] + L^2 \sum_{t=0}^{T-1} \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\|\bar{\mathbf{x}}^{(t)} - \mathbf{x}_i^{(t)}\|^2] \\ &\stackrel{(a)}{\leq} \frac{2(1-\beta)}{\gamma} (\mathbb{E}[f(\bar{\mathbf{z}}^{(0)})] - \mathbb{E}[f(\bar{\mathbf{z}}^{(T)})]) - \left(1 - \frac{L\gamma\beta}{(1-\beta)^2} - \frac{16L^2\gamma^2}{(1-\beta)^2(1-\sqrt{\rho})^2}\right) \sum_{t=0}^{T-1} \mathbb{E}\left[\left\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})\right\|^2\right] \\ &\quad + \frac{L\gamma}{(1-\beta)^2} \sum_{t=0}^{T-1} \mathbb{E}\left[\left\|\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)}\right\|^2\right] + \frac{4L^2\gamma^2\sigma^2}{(1-\beta)^2(1-\rho)} T + \frac{4L^2\gamma^2\kappa^2}{(1-\beta)^2(1-\sqrt{\rho})^2} T \\ &\stackrel{(b)}{\leq} \frac{2(1-\beta)}{\gamma} (\mathbb{E}[f(\bar{\mathbf{z}}^{(0)})] - \mathbb{E}[f(\bar{\mathbf{z}}^{(T)})]) - \left(1 - \frac{(1+\beta)L\gamma}{(1-\beta)^2} - \frac{16L^2\gamma^2}{(1-\beta)^2(1-\sqrt{\rho})^2}\right) \sum_{t=0}^{T-1} \mathbb{E}\left[\left\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})\right\|^2\right] \\ &\quad + \frac{L\gamma}{(1-\beta)^2} \frac{\sigma^2}{N} T + \frac{4L^2\gamma^2\sigma^2}{(1-\beta)^2(1-\rho)} T + \frac{4L^2\gamma^2\kappa^2}{(1-\beta)^2(1-\sqrt{\rho})^2} T \\ &\stackrel{(c)}{\leq} \frac{2(1-\beta)}{\gamma} (f(\bar{\mathbf{x}}^{(0)}) - f^*) + \frac{L\gamma}{(1-\beta)^2} \frac{\sigma^2}{N} T + \frac{4L^2\gamma^2\sigma^2}{(1-\beta)^2(1-\rho)} T + \frac{4L^2\gamma^2\kappa^2}{(1-\beta)^2(1-\sqrt{\rho})^2} T \end{aligned}$$

where (a) follows by applying Lemma 10 (noting that $\frac{\beta L\gamma}{(1-\beta)^2} + \frac{L\gamma}{1-\beta} = \frac{L\gamma}{(1-\beta)^2}$) and Corollary 4 (noting that γ satisfies the condition in Corollary 4); (b) follows by noting that $\mathbb{E}\left[\left\|\frac{1}{N} \sum_{i=1}^N \mathbf{g}_i^{(t)}\right\|^2\right] \leq \frac{1}{N} \sigma^2 + \mathbb{E}\left[\left\|\frac{1}{N} \sum_{i=1}^N \nabla f_i(\mathbf{x}_i^{(t)})\right\|^2\right]$ by Lemma 1; and (c) follows because γ is chosen to ensure $1 - \frac{(1+\beta)L\gamma}{(1-\beta)^2} - \frac{16L^2\gamma^2}{(1-\beta)^2(1-\sqrt{\rho})^2} \geq 0$, $\bar{\mathbf{z}}^{(0)} = \bar{\mathbf{x}}^{(0)}$ by definition in (80), and f^* is the minimum value of problem (1).

Dividing both sides by T yields

$$\begin{aligned} &\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla f(\bar{\mathbf{x}}^{(t)})\|^2] \\ &\leq \frac{2(1-\beta)}{\gamma T} (f(\bar{\mathbf{x}}^{(0)}) - f^*) + \frac{L\gamma}{(1-\beta)^2} \frac{\sigma^2}{N} + \frac{4L^2\gamma^2\sigma^2}{(1-\beta)^2(1-\rho)} + \frac{4L^2\gamma^2\kappa^2}{(1-\beta)^2(1-\sqrt{\rho})^2} \\ &= O\left(\frac{1}{\gamma T}\right) + O\left(\frac{\gamma}{N}\sigma^2\right) + O(\gamma^2\sigma^2) + O(\gamma^2\kappa^2) \end{aligned} \quad (97)$$




Figure 4: Algorithm 1 with constant learning rates: linear speedup verification with ResNet56 over CIFAR10.

6.7. More Experiments

6.7.1. EXPERIMENTS ON ALGORITHM 1 WITH CONSTANT LEARNING RATES


Recall our theory, e.g., Corollary 2, establishes the linear speedup property of Algorithm 1 (with $I > 0$) using constant learning rates. We now verify our theory by training ResNet56 over CIFAR10 using learning rates faithful to our theory. We run SGD with Polyak’s momentum using momentum coefficient 0.9, batch size 64 and constant learning rate 0.01. We also run Algorithm 1 with Option I and $I = 4$ over $N \in \{2, 4, 8\}$ GPUs. The local batch size at each GPU is 64 and the learning rate is set to $0.01\sqrt{N}$. Figure 4(a) plots the convergence of training loss in terms of the number of epochs that are jointly accessed by all used GPUs. Figure 4(a) verifies the linear speedup property of our Algorithm 1 since each GPU in Algorithm 1 only uses $1/N$ of the number of epochs, which are used by the single worker momentum SGD, to converge to the same loss value. A more straightforward verification is given in Figure 4(b) where the x-axis is the wall clock time.

6.7.2. COMPARISONS BETWEEN ALGORITHM 1 AND EXISTING MODEL AVERAGING WITH MOMENTUM SGD

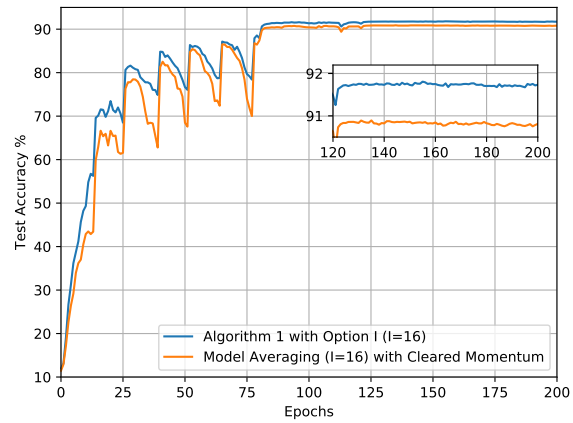
Our Algorithm 1 is quite similar to a common practice, known as “model averaging”, used for training deep neural networks with multiple workers. When the deep neural networks are trained with momentum SGD, Microsoft’s CNTK framework (Seide & Agarwal, 2016) and recent work (Wang & Joshi, 2018b) suggest that each worker should additionally clear its local momentum buffer by resetting it to 0 when they average their local models. In contrast, our Algorithm 1 averages both local momentum buffers and local models. There is no convergence guarantees shown in the literature for the model averaging strategy suggested in (Seide & Agarwal, 2016; Wang & Joshi, 2018b). We run both Algorithm 1 and the model averaging with Polyak’s momentum SGD suggested in (Seide & Agarwal, 2016; Wang & Joshi, 2018b), referred as “model averaging with cleared momentum” to train ResNet56 over CIFAR10. In the experiment, both Algorithm 1 and “model averaging with cleared momentum” perform one synchronization step, e.g., (6) in Algorithm 1, every $I = 16$ iterations. Figure 5 shows that Algorithm 1 converges slightly faster. More impressively, the test accuracy attained by Algorithm 1 is roughly 1% better.

6.7.3. EXPERIMENTS WITH IMAGENET

We further verify the performance of Algorithm 1 when used to train ResNet50 over ImageNet (Deng et al., 2009), which is an image classification task significantly harder than CIFAR10. We run Algorithm 1 with $I \in \{4, 8, 16\}$ and the classical parallel SGD with momentum on a machine with 8 NVIDIA P100 GPUs. The local batch size at each GPU is 128. The learning rate is initialized to 0.2 and is divided by 10 when all GPUs jointly access 30 and 60 epochs. The momentum coefficient is set to 0.9. Figure 6 plots the convergence of training loss and test accuracy in terms of the number of iteration steps. This figure again verifies that Algorithm 1 has the same convergence rate as the classical parallel momentum SGD. To verify the benefit of skipping communication, Figure 7 plots the convergence of training loss and test accuracy in terms of the wall clock time.

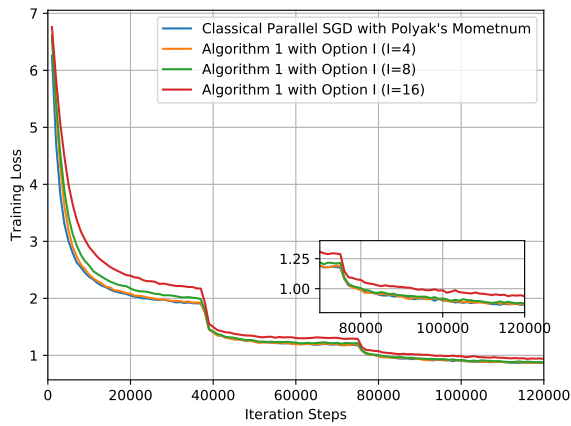


(a) Training loss v.s. epochs




(b) Test accuracy v.s. epochs.

Figure 5: Comparisons between Algorithm 1 and existing model averaging with momentum SGD.



(a) Training loss v.s. iteration steps.




(b) Test accuracy v.s. iteration steps.

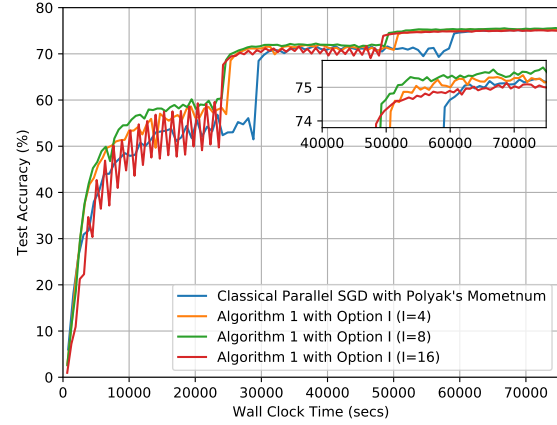
Figure 6: Algorithm 1 with Option I: convergence v.s. iteration steps for ResNet50 over ImageNet.

6.7.4. EXPERIMENTS ON ALGORITHM 1 WITH OPTION II

Figures 8 and 9 verify the convergence of Algorithm 1 with Option II. Their experiment configuration is identical to Figures 2 and 3 except that Polyak’s momentum in all algorithms is replaced with Nesterov’s momentum. While Figures 8 and 9 verify the convergence rate analysis proven in Theorem 2, our preliminary experiments seem to suggest Nesterov’s momentum is less robust to communication skipping since even a small I leads to performance degradation of test accuracy.




(a) Training loss v.s. wall clock time.




(b) Test accuracy v.s. wall clock time.

Figure 7: Algorithm 1 with Option I: convergence v.s. wall clock time for ResNet50 over ImageNet.




(a) Training loss v.s. epochs




(b) Test accuracy v.s. epochs.

Figure 8: Algorithm 1 with Option II: convergence v.s. epochs for ResNet56 over CIFAR10.



(a) Training loss v.s. wall clock time.



(b) Test accuracy v.s. wall clock time.

Figure 9: Algorithm 1 with Option II: convergence v.s. wall clock time for ResNet56 over CIFAR10.